THE AI HUMANIST MOVEMENT

HUMANITY | CLARITY | HARMONY

THE AI HUMANIST MANIFESTO

Humanity Must Evolve With AI, Not Become Dependent Upon It



The AI Humanist Movement (AIHM)

Founding Manifesto

October 2026

THE AI HUMANIST MANIFESTO

Humanity Must Evolve With AI, Not Become Dependent Upon It

The AI Humanist Movement (AIHM)

Founding Manifesto | Draft 1 | October 2026

FOREWORD FROM THE FOUNDER

When I founded The AI Humanist Movement, I did not intend to join the familiar debate over whether artificial intelligence is a blessing or a threat. Both positions miss the deeper question: what happens to human intelligence as machines increasingly participate in our thinking?

For decades as an economist and AI strategist, I have studied how technology changes markets, institutions, and people. I know the value of innovation and efficiency. I also know that greater output is not necessarily human progress. An organization can grow more productive while its people become less able to exercise independent judgment.

AI can challenge our assumptions, expose what we have overlooked, and help us reach insights that would otherwise remain out of reach. It can make us better thinkers. But it can also make it easy to surrender the effort of thinking. The results may look equally impressive while the consequences are fundamentally different.

That distinction led me to Human-AI Cognitive Co-evolution (HACC) and Human-AI Cognitive Outsourcing (HACO). HACC concerns lasting development of human capability through engagement with AI. HACO concerns dependence that erodes the ability to reason, evaluate, and decide independently. Their balance will be shaped by choices in technology design, education, business, and public policy.

This Manifesto is neither a rejection of intelligent machines nor an attempt to preserve every human task. It is a call to ensure that as machines become more capable, we do too. I offer it as the founding statement of a Movement committed to evidence, open inquiry, and human cognitive sovereignty.

Dr. Matt Hasan
Founder, The AI Humanist Movement

PREAMBLE: HUMANITY AT A COGNITIVE CROSSROADS

Humanity stands at a turning point unlike any it has encountered before.

For thousands of years, human beings have developed tools to extend their physical capabilities. The wheel extended mobility. The telescope extended vision. The printing press extended the reach of knowledge. Industrial machinery multiplied physical productivity. Computers transformed our capacity to process information.

Artificial intelligence represents something fundamentally different.

For the first time, we are creating technologies capable of participating directly in activities that have traditionally defined human intelligence: reasoning, learning, creating, interpreting, evaluating, and making decisions.

AI is no longer simply a tool that helps us accomplish what we have already decided to do. Increasingly, it participates in deciding what should be done, how problems should be understood, which alternatives deserve consideration, and what conclusions should be reached.

This transformation raises questions far more consequential than whether machines will become more intelligent than humans, whether they will eliminate jobs, or whether they can be made sufficiently safe.

The most important question is what will happen to human intelligence itself.

Will sustained interaction with increasingly capable artificial intelligence make human beings more capable, more creative, more discerning, and more independent?

Or will it progressively weaken the very cognitive abilities that distinguish us, leaving humanity increasingly dependent on systems it can neither fully understand nor independently challenge?

Neither outcome is inevitable.

The future will be shaped by how artificial intelligence is designed, how it is deployed, how institutions incorporate it, and how individuals learn to interact with it.

We reject the assumption that technological progress automatically produces human progress.

We equally reject the belief that protecting humanity requires resisting the development of increasingly powerful artificial intelligence.

Both positions misunderstand the nature of the challenge.

Humanity does not need to defeat artificial intelligence, retreat from it, or surrender its intellectual future to it.

Humanity must learn to evolve with artificial intelligence.

This is the foundational conviction of The AI Humanist Movement.

We envision a future in which human and artificial intelligence develop through sustained, mutually influential interaction, expanding the boundaries of human understanding, creativity, judgment, and accomplishment.

We call this process Human-AI Cognitive Co-evolution (HACC).

But we recognize an opposing possibility: a future in which increasingly capable machines perform more of humanity's cognitive functions while human beings progressively lose the motivation, opportunity, or ability to exercise those functions independently.

We call this Human-AI Cognitive Outsourcing (HACO).

These are not merely competing technological possibilities. They represent fundamentally different trajectories for human development.

The choice between them will influence education, employment, healthcare, economic opportunity, democratic institutions, individual freedom, and ultimately the character of human civilization.

The AI Humanist Movement exists to help humanity understand and shape that choice.

Our governing proposition is straightforward:

The ultimate measure of AI progress is not how intelligent machines become, but how much more capable humans become through their interaction with them.

That proposition is the foundation of this Manifesto.

I. THE LIMITATIONS OF TODAY'S AI DEBATE

The public conversation about artificial intelligence has become trapped between competing anxieties and competing promises.

On one side are those who see AI primarily as an unprecedented threat. They warn of mass unemployment, manipulation, misinformation, loss of privacy, concentration of power, and potentially catastrophic consequences from systems that exceed human control.

On the other side are those who emphasize extraordinary opportunities: scientific discovery, economic productivity, personalized education, medical advancement, creative expression, and the democratization of expertise.

Both perspectives contain important truths.

Both are also incomplete.

The problem is not that we are debating the wrong issues. It is that we are debating them within an inadequate intellectual framework.

Much of the discussion treats human beings as relatively fixed entities who will experience the consequences of an increasingly powerful technology.

AI changes. Humans adapt.

AI becomes more capable. Humans benefit or suffer.

AI performs tasks. Humans gain or lose employment.

This framework largely overlooks the possibility that human cognitive capabilities themselves will change through repeated interaction with artificial intelligence.

That omission is consequential.

Consider the familiar debate about AI in education.

One group argues that AI enables personalized instruction, helping students learn faster and more effectively.

Another argues that AI encourages intellectual laziness, allowing students to produce sophisticated assignments without developing the knowledge or skills those assignments are supposed to demonstrate.

Both can be correct.

The relevant question is not simply whether students use AI. It is how their interaction with AI changes their ability to think, reason, and learn when AI is no longer available.

The same distinction applies to professional work.

An employee who uses AI to produce a better financial analysis may become a more capable analyst through the process. Alternatively, the employee may gradually lose the ability to conduct the analysis independently.

In both cases, the immediate output may improve.

But the long-term consequences for human capability are radically different.

Current measures of AI success rarely capture this distinction.

Organizations evaluate productivity, accuracy, cost reduction, customer satisfaction, and return on investment. These are legitimate measures, but they do not tell us whether the human beings working with AI are becoming more capable.

Indeed, an organization could report extraordinary productivity gains while systematically diminishing the independent judgment of its workforce.

The same problem arises in AI governance.

Safety, fairness, transparency, accountability, and privacy are essential. They help determine whether AI systems operate within acceptable boundaries.

But a system can be safe, fair, transparent, and accurate while still encouraging excessive cognitive dependence.

Preventing harm is not equivalent to advancing human development.

We therefore need a broader conception of AI progress.

The question should not be limited to whether artificial intelligence performs its assigned functions responsibly.

It must also encompass what sustained interaction with artificial intelligence does to the people who use it.

This is the intellectual territory that The AI Humanist Movement seeks to illuminate.

II. HUMAN-AI COGNITIVE CO-EVOLUTION: A NEW PARADIGM

Human intelligence has never developed in isolation.

Language, writing, education, social institutions, and technological tools have shaped the ways human beings acquire knowledge, solve problems, and understand the world.

The human mind is influenced by its environment, and that environment is increasingly shaped by human inventions.

Artificial intelligence introduces a new dimension into this relationship.

Unlike conventional tools, advanced AI systems can respond to questions, challenge assumptions, propose alternatives, explain unfamiliar concepts, identify errors, and adapt their responses to the context of an interaction.

These capabilities make AI a potentially active participant in human cognitive development.

Human-AI Cognitive Co-evolution describes a process in which repeated interaction between humans and artificial intelligence changes the capabilities, behaviors, and future interactions of both participants.

The essential insight is that human-AI interaction is not a series of isolated transactions.

It is a continuing process with consequences that accumulate over time.

A person asks an AI system to help understand an unfamiliar economic concept.

The system explains the concept, provides examples, and responds to objections.

The person develops a deeper understanding and begins asking more sophisticated questions.

Those questions elicit more sophisticated responses.

The person learns to recognize weaknesses in the system's reasoning and develops better ways to evaluate its conclusions.

Future interactions become more demanding, more productive, and potentially more intellectually enriching.

The human participant is no longer the same person who began the interaction.

The interaction has changed the person's capabilities.

The AI system may also change through adaptation, personalization, retained context, feedback, or subsequent system development. Such changes need not imply consciousness, human-like understanding, or independent agency.

What is important is that the evolving interaction can create a reinforcing process.

Human capability influences how AI is used.

AI interaction influences human capability.

Those changes influence subsequent interactions.

This is the central dynamic of HACC.

Beyond the augmentation metaphor

AI is frequently described as augmenting human intelligence.

The expression is useful, but insufficient.

Augmentation often refers to an immediate improvement in performance.

A calculator enables someone to perform arithmetic faster. A navigation system helps someone find a destination. An AI assistant can help a researcher summarize a complex body of literature.

But improved performance does not necessarily imply improved human capability.

A person can accomplish more with a tool without becoming more capable independently of that tool.

HACC makes this distinction central.

It asks whether repeated interaction produces durable changes in the human participant's capacity to understand, reason, learn, create, and exercise judgment.

The distinction is between doing more with AI and becoming more capable through AI.

Both outcomes can be valuable, but they are not equivalent.

A scientist who uses AI to identify an overlooked hypothesis may develop a deeper understanding of the underlying scientific problem.

A physician who uses AI to examine competing diagnoses may become more attentive to subtle clinical evidence.

A student who challenges an AI-generated explanation may strengthen analytical reasoning.

A small-business owner who uses AI to explore financial scenarios may develop greater confidence and competence in business decisions.

In each case, the greatest potential benefit is not merely the immediate answer.

It is the possibility that the human participant emerges from the interaction with stronger capabilities.

Co-evolution does not require equal intelligence

Human-AI co-evolution does not imply that human beings and artificial intelligence possess identical forms of intelligence.

Nor does it require AI systems to become conscious, sentient, or human-like.

The relevant relationship is functional and developmental.

A human being brings experience, motivation, values, social understanding, responsibility, and embodied knowledge.

An AI system contributes computational capabilities, access to learned patterns, analytical assistance, and the ability to generate and evaluate alternatives.

These capabilities can complement one another.

But complementarity does not automatically produce co-evolution.

The interaction must create opportunities for learning, reflection, and the development of independent human competence.

That requirement distinguishes HACC from the more general idea of human-AI collaboration.

A theory that must be tested

HACC is not a declaration that AI necessarily makes people smarter.

It is a theoretical framework for understanding the conditions under which sustained interaction can expand or diminish human cognitive capabilities.

Its propositions must be subjected to empirical investigation.

Researchers must examine changes in independent problem-solving ability, reasoning quality, knowledge retention, creativity, judgment, and the transfer of learned capabilities to unfamiliar situations.

They must distinguish temporary improvements in AI-assisted performance from durable improvements in human competence.

They must also examine differences across individuals, tasks, institutions, and forms of interaction.

The AI Humanist Movement welcomes such investigation.

Our purpose is not to protect a theory from criticism.

It is to advance understanding of a phenomenon that will increasingly shape human development.

III. HACC AND HACO: TWO TRAJECTORIES OF HUMAN DEVELOPMENT

The relationship between artificial intelligence and human capability contains a fundamental tension.

The same technology that can stimulate learning can also eliminate the need to learn.

The same system that can strengthen judgment can encourage the abandonment of judgment.

The same assistant that expands creativity can become a substitute for creative effort.

This tension lies at the heart of the distinction between HACC and Human-AI Cognitive Outsourcing.

The promise of HACC

Consider two individuals learning a foreign language.

The first uses AI as an interactive tutor.

The system explains grammar, introduces unfamiliar vocabulary, simulates conversations, corrects mistakes, and gradually increases the difficulty of exercises.

The learner practices independently and returns to the system for feedback.

Over time, the learner develops greater fluency.

The second individual uses AI primarily to translate messages and compose responses.

Communication becomes easier and more efficient, but the individual may develop little independent language capability.

Both individuals benefit from AI.

Only the first is necessarily experiencing the kind of durable cognitive development that HACC seeks to understand.

The difference lies not in the technology itself but in the structure of the interaction.

One pattern encourages active participation, learning, and independent competence.

The other substitutes machine performance for human effort.

The risk of HACO

Cognitive outsourcing is not inherently undesirable.

Human civilization depends on specialization, delegation, and the use of tools.

Few people calculate complex engineering equations by hand when reliable software is available.

Few memorize information that can be retrieved quickly and accurately.

Delegating routine cognitive tasks can free time and attention for activities of greater value.

The problem arises when delegation progressively undermines capabilities that remain important for independent understanding, evaluation, and decision-making.

A manager who routinely accepts AI-generated strategic recommendations without examining their assumptions may gradually lose analytical discipline.

A student who relies on AI to construct every argument may never develop the ability to formulate one independently.

A physician who routinely accepts algorithmic recommendations without exercising clinical judgment may become less capable of recognizing situations in which the system is mistaken.

The danger is not simply that AI occasionally produces incorrect answers.

It is that human beings may become progressively less capable of recognizing those errors.

This creates a potentially self-reinforcing process.

As independent capability declines, dependence on AI increases.

As dependence increases, opportunities to exercise independent capability decline.

The resulting pattern can make individuals increasingly vulnerable to system failures, manipulation, or the loss of access to technology.

Not a simple binary

HACC and HACO should not be understood as mutually exclusive labels attached permanently to individuals or institutions.

Human cognitive capability is multidimensional.

A person may become more creative while relying increasingly on AI for factual recall.

An engineer may lose proficiency in routine calculations while developing stronger systems-level judgment.

A physician may delegate administrative reasoning while becoming more attentive to patients.

Some forms of cognitive outsourcing may be entirely compatible with broader cognitive development.

The central question is which capabilities are being strengthened, which are being relinquished, and whether individuals retain sufficient independent competence to understand and govern the consequential activities they delegate.

This leads to a foundational principle:

Delegation should expand human possibilities without surrendering the capabilities necessary for meaningful human agency.

The importance of interaction design

HACC and HACO are not predetermined outcomes.

Their relative prevalence will depend on the design of AI systems, the incentives governing their deployment, and the behavior of the people using them.

An AI system designed exclusively to deliver the fastest possible answer may encourage passive acceptance.

A system designed to help users examine assumptions, compare alternatives, and understand uncertainty may encourage deeper engagement.

Neither approach is universally appropriate.

A person facing an emergency may need an immediate answer rather than an educational exercise.

A professional performing a repetitive administrative task may reasonably prioritize efficiency.

But across sustained interaction, the design of AI systems should take account of their effects on human capability.

The objective is not to maximize cognitive effort in every activity.

It is to preserve and develop the human capabilities that enable people to act intelligently, independently, and responsibly.

IV. HUMAN AGENCY AND COGNITIVE SOVEREIGNTY

Human freedom has traditionally been understood in political, economic, and social terms.

We value freedom of expression, freedom of association, freedom of belief, and the ability to make consequential decisions about our own lives.

Artificial intelligence introduces another dimension of freedom.

We call it cognitive sovereignty.

Cognitive sovereignty is the capacity to form judgments, evaluate claims, understand consequential choices, and exercise meaningful intellectual independence in an environment increasingly mediated by artificial intelligence.

It does not require individuals to know everything or perform every task without assistance.

No human being has ever possessed such independence.

Rather, it requires that people retain the ability to question, challenge, and, when necessary, reject the recommendations of systems that influence their lives.

A society in which individuals can access extraordinary intelligence but cannot independently evaluate its conclusions may be technologically advanced while becoming intellectually dependent.

Such dependence has consequences extending far beyond individual convenience.

The illusion of informed choice

Imagine a consumer choosing a mortgage.

An AI assistant compares lenders, evaluates interest rates, estimates monthly payments, and recommends a particular product.

The recommendation may be accurate and beneficial.

But what if the consumer does not understand the assumptions underlying the comparison?

What if the assistant's recommendations are influenced by commercial incentives?

What if the consumer has become so accustomed to accepting algorithmic advice that independent evaluation no longer seems necessary?

The existence of multiple options does not guarantee meaningful choice.

Choice requires some ability to understand and evaluate alternatives.

This is particularly important when AI systems mediate decisions involving healthcare, employment, education, finance, and public services.

Cognitive sovereignty and democracy

Democratic societies depend on citizens capable of evaluating competing claims, questioning authority, and participating in collective decision-making.

Artificial intelligence can strengthen those capabilities by making information more accessible and helping individuals understand complex issues.

It can also weaken them by encouraging passive consumption of personalized explanations, recommendations, and interpretations.

The danger is not limited to misinformation.

Even factually accurate information can be selected, framed, and presented in ways that influence judgment.

When individuals increasingly delegate the interpretation of public affairs to AI systems, the design and governance of those systems acquire political significance.

The AI Humanist Movement therefore regards cognitive sovereignty as a public concern, not merely a personal preference.

A healthy democracy requires more than access to information.

It requires citizens who retain the capacity to think critically about that information.

Human judgment must remain meaningful

Calls for human oversight of AI are common.

But oversight becomes an empty formality when the designated human lacks the knowledge, time, authority, or confidence to challenge the system.

A person who automatically approves algorithmic recommendations is not exercising meaningful oversight.

The presence of a human in a decision process does not guarantee the preservation of human agency.

We must distinguish nominal human involvement from substantive human judgment.

This distinction should guide AI governance in every consequential domain.

Human beings must not merely remain present in the process.

They must remain capable participants in it.

V. THE ECONOMICS OF HUMAN-AI CO-EVOLUTION

Artificial intelligence is frequently evaluated through the economics of productivity.

How much labor can it save?

How many tasks can it automate?

How much can it reduce operating costs?

How quickly can organizations recover their investments?

These are important questions, but they represent only part of the economic picture.

Human cognitive capability is itself a productive asset.

Education, experience, professional judgment, creativity, and problem-solving ability contribute to economic value.

When AI changes these capabilities, it changes the economy's productive foundations.

HACC therefore introduces a broader economic perspective.

The relevant measure of AI's value is not simply its immediate contribution to output.

It also includes its effects on the future productive capacity of the people who use it.

Productivity today versus capability tomorrow

Consider an organization that introduces AI to perform much of its analytical work.

Productivity increases. Reports are produced faster. Costs decline.

From a conventional accounting perspective, the investment appears successful.

But suppose employees gradually lose the ability to conduct independent analysis.

The organization may become increasingly dependent on external AI providers.

Its ability to evaluate unexpected situations may weaken.

Its capacity to innovate independently may decline.

Its vulnerability to system failures or changes in vendor terms may increase.

The short-term productivity gain may conceal a long-term reduction in organizational capability.

Now consider an alternative.

The organization uses AI to automate routine analysis while encouraging employees to investigate unusual patterns, challenge assumptions, and develop deeper expertise.

Productivity improves, but so does human capability.

The organization becomes more adaptable rather than merely more efficient.

These two organizations may report similar immediate financial results while following very different long-term economic trajectories.

Human capability as an accumulating asset

Human knowledge and expertise develop through investment, experience, practice, and feedback.

They can also deteriorate through disuse, obsolescence, and changing circumstances.

HACC treats human cognitive capability as something that can accumulate or depreciate over time.

AI interaction can contribute to either process.

This has implications for investment decisions.

Organizations should evaluate not only how much labor an AI system replaces but also how it changes the capabilities of the remaining workforce.

Educational institutions should evaluate not only whether AI improves examination performance but also whether students retain knowledge and develop independent reasoning.

Governments should consider not only the productivity effects of AI adoption but also its consequences for workforce adaptability and long-term economic resilience.

The economics of dependence

Cognitive dependence also creates economic dependence.

When individuals and organizations lose the capacity to perform essential functions independently, their bargaining position relative to technology providers may weaken.

Switching costs can rise.

Vendor lock-in can deepen.

The ability to evaluate competing systems can diminish.

A company may become dependent not merely on a particular software product but on the external cognitive infrastructure through which its employees conduct essential work.

This creates potential economic rents for providers that control indispensable systems.

It also raises questions about competition policy, market concentration, interoperability, and the distribution of gains from AI adoption.

The AI Humanist Movement does not oppose commercial success or the development of powerful AI enterprises.

Innovation requires investment, risk-taking, and the prospect of economic returns.

But markets should be structured so that technological progress does not require the progressive surrender of human and institutional capability.

Beyond employment counts

The debate over AI and jobs often focuses on how many positions will disappear and how many new positions will emerge.

Those questions are important.

But the quality of the remaining work also deserves attention.

Will workers exercise greater judgment and creativity?

Will they develop new expertise?

Will they become more adaptable?

Or will they increasingly supervise systems whose reasoning they no longer understand?

A society can preserve employment while diminishing the intellectual substance of work.

Conversely, AI can eliminate some tasks while enabling people to undertake more demanding and fulfilling activities.

The distinction depends on how work is reorganized.

The economic objective should not be the preservation of every existing task.

It should be the development of a workforce capable of creating increasing value through meaningful interaction with advancing technology.

A humanist economy does not measure progress solely by what machines produce. It also measures what people become capable of producing, understanding, and contributing.

VI. EDUCATION, WORK, HEALTHCARE, AND EVERYDAY LIFE

The consequences of human-AI cognitive co-evolution will not be confined to laboratories or technology companies.

They will emerge in classrooms, workplaces, hospitals, homes, and ordinary daily interactions.

These are the environments in which the future relationship between human and artificial intelligence will actually be formed.

Education: learning versus completing assignments

Education is perhaps the most consequential testing ground for HACC.

The purpose of education is not simply to produce correct answers.

It is to develop the capacity to understand questions, construct arguments, solve unfamiliar problems, and acquire new knowledge independently.

AI creates extraordinary opportunities to personalize learning.

A student can receive explanations tailored to individual needs, practice difficult concepts, explore alternative interpretations, and obtain feedback that would otherwise be unavailable.

But AI also makes it possible to bypass much of the cognitive effort through which learning occurs.

A student can submit a sophisticated essay without understanding its argument.

A mathematics problem can be solved without the student developing mathematical reasoning.

A research assignment can be completed without learning how to evaluate evidence.

The challenge is not to prohibit AI from education.

It is to redesign education around the capabilities students should develop through interaction with AI.

Assessment must distinguish assisted performance from independently retained competence.

Teachers must be equipped to use AI as a developmental resource rather than merely a productivity tool.

Students must learn not only how to obtain answers but how to evaluate, challenge, and improve them.

The objective is a generation of students whose intellectual capabilities expand because of AI, not one whose apparent competence disappears when the technology is removed.

Work: from task automation to capability development

The workplace is undergoing a transformation in the distribution of cognitive labor.

AI increasingly performs activities once associated with professional expertise, including drafting, analysis, programming, forecasting, and decision support.

Organizations face a choice.

They can use AI primarily to reduce labor costs and accelerate existing processes.

Or they can use it to redesign work around stronger human capabilities.

The latter approach requires deliberate investment.

Employees need opportunities to understand the reasoning behind AI-generated recommendations.

They need authority to challenge outputs.

They need training that develops transferable skills rather than proficiency in a particular software interface.

And they need organizational cultures that reward judgment, creativity, and intellectual initiative.

A workforce that simply follows increasingly sophisticated machine instructions may become more productive in the short term while losing its capacity for independent innovation.

The AI Humanist Movement advocates a different trajectory.

AI should reduce unnecessary cognitive burdens while expanding opportunities for meaningful human contribution.

Healthcare: clinical intelligence and human judgment

Healthcare illustrates both the promise and the risks of cognitive dependence.

AI can help clinicians identify patterns, interpret diagnostic information, review medical literature, and evaluate treatment alternatives.

It can also help patients understand conditions, prepare questions, and participate more effectively in decisions about their care.

These capabilities offer significant opportunities for better outcomes.

But clinical accuracy alone is not an adequate measure of success.

A highly accurate AI system may still weaken clinical judgment if physicians become accustomed to accepting its recommendations without sufficient scrutiny.

Similarly, patient-facing AI may improve access to information while creating confusion, misplaced confidence, or excessive dependence.

The relevant question is whether AI strengthens the capabilities of clinicians and patients to make informed decisions.

A humanist approach to healthcare AI therefore emphasizes clinical competence, patient understanding, meaningful accountability, and the preservation of independent professional judgment.

The physician should not become a ceremonial approver of algorithmic decisions.

The patient should not become a passive recipient of automated advice.

AI should improve the quality of the relationship between clinical expertise, patient understanding, and informed choice.

Everyday life: the invisible transformation

Perhaps the most profound consequences of AI will emerge through ordinary activities.

People increasingly use AI to write messages, plan travel, evaluate purchases, organize finances, resolve disagreements, and understand unfamiliar situations.

Each interaction may seem trivial.

Collectively, they can influence habits of attention, memory, reasoning, and decision-making.

A person who uses AI to explore alternative perspectives may become more reflective.

Another who uses AI to confirm existing beliefs may become less open to disagreement.

A person who uses AI to understand financial decisions may become more financially capable.

Another who delegates every financial choice may become less able to recognize manipulation or error.

These changes may be gradual and difficult to observe.

That makes them no less consequential.

Human-AI cognitive co-evolution will be shaped not only by spectacular technological breakthroughs but by millions of ordinary decisions about when to think independently, when to seek assistance, and when to delegate.

VII. INEQUALITY AND THE EMERGING COGNITIVE DIVIDE

Technological progress has historically produced unequal benefits.

Differences in education, income, infrastructure, institutional capacity, and political influence affect who gains access to new technologies and who captures their economic value.

Artificial intelligence could reproduce these inequalities.

It could also create new forms of inequality that are more difficult to recognize.

The AI Humanist Movement is particularly concerned about the possibility of an emerging cognitive divide.

This divide would not be defined solely by access to AI.

It would also reflect differences in the ability to use AI in ways that strengthen independent human capabilities.

Two individuals may have access to the same AI system but experience very different developmental outcomes.

One possesses the education, confidence, and critical-thinking skills necessary to question its outputs and use it as a learning partner.

The other relies on the system primarily for answers and decisions.

Over time, the first individual may become increasingly capable while the second becomes increasingly dependent.

The technology is identical.

The developmental consequences are not.

Access is necessary but insufficient

Efforts to democratize AI frequently emphasize affordability and availability.

These are essential objectives.

But access alone does not guarantee empowerment.

Meaningful cognitive opportunity also requires education, digital literacy, institutional support, and the ability to evaluate information critically.

A person who can use AI but cannot assess its recommendations may be technologically connected without being cognitively empowered.

Policies designed to broaden access must therefore be accompanied by efforts to strengthen the capabilities needed for productive interaction.

The risk of cumulative divergence

Cognitive inequality may become self-reinforcing.

Individuals with stronger initial capabilities may use AI more effectively.

More effective interaction may strengthen those capabilities further.

Those gains may enable even more productive future interactions.

Meanwhile, individuals with fewer educational resources or weaker institutional support may be encouraged toward passive reliance.

Over time, relatively modest initial differences could produce substantial divergence.

This is a possibility that deserves rigorous empirical investigation, not an outcome we should assume is inevitable.

But the potential consequences are sufficiently important to warrant attention now.

Institutions and unequal capability

The divide may also emerge between organizations.

Large enterprises may have the resources to build sophisticated AI systems, train employees, and develop internal expertise.

Small businesses, schools, nonprofits, and public institutions may lack comparable resources.

If these organizations become dependent on external AI providers without developing their own capacity to evaluate and govern those systems, existing institutional inequalities could deepen.

The same concern applies internationally.

Countries that invest in human capability alongside AI infrastructure may achieve more durable benefits than those that become primarily consumers of externally developed cognitive services.

The objective should not be technological self-sufficiency in every country or organization.

It should be sufficient independent capability to exercise informed judgment, maintain meaningful choices, and participate in shaping technological development.

Cognitive opportunity as a public good

Education has long been recognized as essential to economic development and democratic participation.

In an AI-mediated society, the ability to interact productively and critically with artificial intelligence will become an increasingly important component of education.

The AI Humanist Movement advocates broad access to opportunities for cognitive development through AI.

Such opportunities should not be reserved for privileged institutions or technologically sophisticated populations.

The promise of human-AI co-evolution must extend across income, age, geography, disability, and educational background.

A future in which a small minority uses AI to become more capable while the majority becomes increasingly dependent would represent a profound failure of humanist technological development.

VIII. PRINCIPLES FOR HUMANIST AI DESIGN AND GOVERNANCE

The design of artificial intelligence influences the behavior of the people who use it.

Interfaces, defaults, incentives, explanations, recommendations, and feedback mechanisms all shape patterns of interaction.

These design choices can encourage active engagement or passive acceptance.

They can promote learning or substitute for it.

They can strengthen human agency or make dependence increasingly convenient.

The AI Humanist Movement therefore calls for a broader approach to AI design and governance.

Safety and reliability remain essential.

But the evaluation of AI systems should also consider their effects on human cognitive development.

We propose the following principles.

Principle 1: Human capability must be an explicit design objective

AI systems should be evaluated not only by their immediate performance but also, where relevant, by their effects on the capabilities of the people who use them.

This does not mean every system must function as a tutor.

It means designers should recognize that repeated interaction can influence learning, judgment, and independence.

Where those effects are consequential, they should be considered deliberately rather than treated as incidental.

Principle 2: Human agency must be substantive

Individuals should retain meaningful opportunities to understand, question, and challenge consequential AI recommendations.

Human oversight should not be reduced to the appearance of approval.

Institutions must ensure that responsible individuals possess the competence, authority, and information necessary to exercise judgment.

Principle 3: Cognitive outsourcing must be purposeful

Delegation can be beneficial when it eliminates unnecessary effort or enables higher-value activities.

But institutions should identify capabilities that must be retained independently.

These may include professional judgment, critical reasoning, safety-related skills, and the ability to recognize errors.

The objective is not to preserve obsolete routines for their own sake.

It is to prevent the loss of capabilities whose absence would create serious dependence or vulnerability.

Principle 4: AI should support learning and intellectual development

Where appropriate, AI systems should help users understand the reasoning behind conclusions, examine alternatives, and recognize uncertainty.

Educational and professional applications should encourage forms of interaction that support durable learning.

Designers should investigate when explanations, questions, feedback, and opportunities for independent practice improve human capability.

Principle 5: Cognitive effects should be measured

Organizations should develop ways to assess how sustained AI use affects independent human performance.

Immediate productivity gains should be distinguished from long-term changes in knowledge, judgment, and adaptability.

Such assessments should be proportionate to the risks and purposes of the application.

They should also recognize that cognitive development cannot be reduced to a single universal score.

Principle 6: AI systems must remain contestable

People should have meaningful opportunities to challenge consequential AI-generated conclusions.

Contestability requires more than a technical ability to click a disagreement button.

It requires accessible procedures, relevant information, and institutional accountability.

Where appropriate, users should have avenues for human review and correction.

Principle 7: Competition and interoperability support human freedom

Excessive dependence on a small number of AI providers can limit individual and institutional choice.

Interoperability, portability, competitive markets, and open standards can reduce unnecessary switching costs and support innovation.

These objectives should be pursued alongside legitimate protections for privacy, security, intellectual property, and commercial investment.

Principle 8: The benefits of cognitive advancement must be widely shared

Public policy should support equitable access to AI-enabled learning and capability development.

Schools, libraries, community organizations, and public institutions can play important roles.

The objective is not merely to distribute technology.

It is to expand the ability of people to benefit intellectually and economically from that technology.

Principle 9: Humanist governance must be evidence-based

Claims about AI's effects on human cognition should be evaluated through credible research.

Neither technological optimism nor technological pessimism should substitute for evidence.

Policies should be revised as understanding improves.

The AI Humanist Movement must remain open to findings that challenge its assumptions.

Principle 10: Human development must remain the ultimate objective

AI governance should not be judged solely by regulatory compliance or the absence of harm.

It should also consider whether technological development expands human opportunity, agency, understanding, and capability.

This principle unites the others.

The purpose of humanist AI governance is not simply to make powerful machines behave responsibly. It is to ensure that their growing capabilities contribute to human flourishing.

IX. A NEW RESEARCH AND PUBLIC POLICY AGENDA

The AI Humanist Movement recognizes that the scientific understanding of human-AI cognitive co-evolution remains incomplete.

We do not claim that every proposed mechanism has been demonstrated or that the long-term consequences of widespread AI adoption can already be predicted with confidence.

We call for a systematic research agenda capable of addressing these uncertainties.

Understanding long-term cognitive effects

Much existing research examines the immediate consequences of AI assistance.

We need more longitudinal studies examining what happens to human capabilities after months or years of sustained interaction.

Researchers should investigate whether AI-assisted learning produces durable improvements in independent performance.

They should examine whether repeated delegation contributes to skill depreciation and under what conditions.

They should also investigate whether cognitive gains in one domain transfer to others.

These questions require careful experimental design and credible measures of capability.

Identifying developmental mechanisms

We need to understand which interaction patterns promote learning and which encourage dependence.

Does asking users to explain their reasoning improve retention?

When does AI-generated feedback strengthen judgment?

How do confidence, expertise, motivation, and task complexity influence outcomes?

Which forms of assistance are most beneficial for beginners, and which are most beneficial for experts?

These questions are relevant to the design of educational systems, workplace applications, and consumer AI products.

Measuring independent capability

AI-assisted output is relatively easy to measure.

Independent cognitive development is more difficult.

Researchers should develop methods that distinguish improvements in task performance from improvements in retained human capability.

Possible measures include unaided problem-solving, transfer to unfamiliar tasks, knowledge retention, error detection, and the quality of independent judgment.

No single measure will capture the full range of human cognition.

A rigorous research program must recognize differences across cognitive domains and contexts.

Examining institutional consequences

Research should also investigate how AI changes organizational capability.

Do firms become more adaptable through AI adoption?

Do employees develop new forms of expertise?

How does AI dependence influence resilience, innovation, and bargaining power?

What happens when organizations lose access to systems on which essential functions depend?

These questions connect cognitive co-evolution with economics, management, organizational behavior, and public policy.

Developing appropriate policy responses

Governments should support research into the developmental consequences of AI.

Public institutions should consider cognitive capability when evaluating major AI deployments in education, healthcare, and other consequential settings.

Professional organizations should examine how AI changes the competencies required for responsible practice.

Employers should assess whether automation strategies strengthen or weaken their long-term human capital.

The objective is not to impose a single model of human-AI interaction across society.

It is to ensure that institutions understand and take responsibility for the human consequences of their technological choices.

X. THE AI HUMANIST MOVEMENT: OUR MISSION AND COMMITMENTS

The AI Humanist Movement is founded on a simple but demanding conviction.

Artificial intelligence should expand the possibilities of human development.

It should not become a mechanism through which humanity gradually relinquishes the capacities necessary for intellectual independence, meaningful agency, and responsible self-government.

We recognize the extraordinary achievements of AI researchers, engineers, entrepreneurs, and institutions.

We recognize that artificial intelligence can contribute to scientific discovery, economic growth, better healthcare, broader educational opportunity, and improvements in everyday life.

We reject the assumption that human progress requires artificial intelligence to remain permanently subordinate in capability to human beings.

But we also reject the assumption that greater machine intelligence automatically translates into greater human capability.

The relationship must be deliberately shaped.

Our mission

The mission of The AI Humanist Movement is to advance human cognitive capability, agency, and opportunity through the responsible development and use of artificial intelligence, guided by the principles of Human-AI Cognitive Co-evolution.

We seek to influence how individuals interact with AI, how organizations deploy it, how technology developers design it, and how governments govern it.

Our purpose is not to represent the interests of any particular technology provider, political party, industry, or professional group.

It is to represent a broader human interest in the direction of technological development.

Our commitments

We will advance the science of human-AI cognitive co-evolution.

We will encourage rigorous interdisciplinary research and constructive debate about the long-term effects of AI on human capability.

We will support the development of testable hypotheses, credible measures, and evidence-based approaches to understanding cognitive augmentation and outsourcing.

We will make these ideas accessible to the public.

The future of human cognition is too important to remain confined to academic journals and technology conferences.

We will communicate complex ideas in language that ordinary people can understand.

We will encourage public discussion of the choices individuals and societies face.

We will advocate humanist AI design.

We will encourage technology developers to consider how their systems affect learning, judgment, creativity, and independence.

We will promote design approaches that support meaningful human participation and development.

We will engage institutions and policymakers.

We will contribute to discussions about education, employment, healthcare, competition, public services, and AI governance.

We will advocate policies that recognize human cognitive capability as an important dimension of technological progress.

We will promote equitable cognitive opportunity.

We will support efforts to ensure that the benefits of human-AI co-evolution are not restricted to privileged individuals or institutions.

We will encourage approaches that expand meaningful access to AI-enabled learning and development.

We will remain intellectually independent.

Our positions will be guided by evidence, reasoned analysis, and our commitment to human development.

We will welcome criticism and revise our conclusions when warranted.

The credibility of the Movement must rest on the quality of its ideas and the integrity of its conduct.

A movement, not merely an institution

The AI Humanist Movement cannot achieve its objectives through a small group of researchers or advocates alone.

The transformation we seek requires participation across society.

Educators must reconsider how learning occurs.

Employers must reconsider how work is organized.

Technology developers must reconsider what their systems are designed to optimize.

Healthcare professionals must reconsider how clinical judgment is preserved and strengthened.

Policymakers must reconsider how technological progress is measured and governed.

Individuals must reconsider their own relationships with artificial intelligence.

AIHM seeks to connect these constituencies through a shared commitment to human cognitive development.

Membership should signify more than agreement with an abstract principle.

It should represent a willingness to contribute ideas, expertise, institutional influence, and practical action toward a humanist future.

The Movement's success will not be measured solely by the number of members it attracts.

It will be measured by the extent to which its ideas influence research, technology design, institutional practices, public policy, and everyday human-AI interaction.

XI. A CALL TO HUMANITY

Every major technological transformation has changed what human beings can accomplish.

Some have expanded freedom and opportunity.

Others have concentrated power, displaced livelihoods, or created forms of dependence that societies struggled to understand until their consequences became difficult to reverse.

Artificial intelligence presents both possibilities on an unprecedented cognitive scale.

It can help us understand the world more deeply.

It can expand access to knowledge and expertise.

It can strengthen creativity, accelerate discovery, and enable individuals to accomplish things previously beyond their reach.

But it can also encourage us to relinquish the intellectual effort through which understanding develops.

It can make dependence convenient, profitable, and increasingly invisible.

The greatest danger is not necessarily that artificial intelligence will become too intelligent.

It is that human beings may cease to develop their own intelligence because machines can perform so many cognitive functions for them.

That outcome is neither necessary nor desirable.

We can choose a different path.

We can build educational systems in which AI strengthens learning rather than bypassing it.

We can create workplaces in which AI expands human judgment rather than replacing it with passive supervision.

We can develop healthcare systems in which AI strengthens clinical expertise and patient understanding.

We can establish governance institutions that protect cognitive sovereignty alongside safety, privacy, and fairness.

We can design technologies that encourage curiosity, reflection, creativity, and independent thought.

Above all, we can recognize that the future of artificial intelligence is inseparable from the future of human intelligence.

The challenge before us is not simply technological.

It is educational, economic, institutional, political, and profoundly human.

We must decide what kind of relationship we want with the intelligence we are creating.

Do we want machines that increasingly think in our place?

Or do we want machines that help us become more capable of thinking for ourselves?

The answer will not emerge automatically from technological innovation.

It will be shaped by human choices.

The AI Humanist Movement calls upon researchers, educators, technology developers, business leaders, policymakers, professionals, students, and citizens everywhere to participate in making those choices.

We invite those who welcome AI's extraordinary possibilities.

We invite those who are concerned about its risks.

We invite those who believe technological progress should be judged by its contribution to human development.

We do not ask humanity to choose between human intelligence and artificial intelligence.

We ask humanity to recognize the possibility of a future in which each contributes to the development of the other.

A future in which machines become more capable without human beings becoming less capable.

A future in which the benefits of intelligence are measured not simply by the volume of information processed or the number of tasks automated, but by the growth of human understanding, judgment, creativity, and freedom.

A future in which artificial intelligence expands what it means to be human rather than diminishing it.

Our purpose is not to make humanity indispensable to machines.

It is to ensure that machines help humanity realize more of its own potential.

This is the promise of Human-AI Cognitive Co-evolution.

This is the mission of The AI Humanist Movement.

And this is the future we call upon humanity to build.

OUR FOUNDING DECLARATION

We believe that artificial intelligence represents an extraordinary opportunity to advance human civilization.

We believe that technological progress must be evaluated through its consequences for human capability, agency, and opportunity.

We believe that sustained human-AI interaction can create pathways toward cognitive development or cognitive dependence.

We believe that these outcomes are shaped by human choices, institutional incentives, technological design, and public policy.

We believe that human beings must retain the capacity to understand, question, and govern the systems that influence consequential decisions.

We believe that the benefits of AI-enabled cognitive advancement should be broadly accessible.

We believe that scientific evidence, intellectual independence, and open inquiry must guide the development of humanist AI principles.

We believe that humanity should neither fear artificial intelligence into stagnation nor embrace it without understanding its consequences.

Above all, we believe that the future of artificial intelligence must be judged by the future it makes possible for humanity.

We stand for Human-AI Cognitive Co-evolution.

We stand for human cognitive sovereignty.

We stand for a future in which artificial intelligence helps humanity become more capable, more creative, more independent, and more fully human.

We are The AI Humanist Movement.

The AI Humanist Movement (AIHM)

Founding Manifesto

October 2026