

AI for Human Flourishing

Promise, Proof & Public Value



People directing technology toward health, education, independence and productive work. Editorial illustration.

Every person deserves the chance to develop their abilities, exercise sound judgment, care for others and build a life of purpose. Health, education, useful work and secure communities help make that possible. Artificial intelligence should be judged by what it contributes to those human ends. Its promise includes easing burdens that consume our time, widening access to knowledge and helping us solve difficult problems. Realizing that promise requires people to remain responsible for their decisions, families and communities to remain strong, and institutions to respect human dignity and freedom.

The case for AI therefore rests on the kind of human development it makes possible. Greater productive power deserves support when people can use it to live well, provide for those who depend on them and contribute something worthwhile. Making that connection is the central task.

THE ARGUMENT

Why human flourishing should guide AI

Human flourishing concerns the development of the whole person. It includes physical and mental health, understanding, character, meaningful relationships and the freedom to take responsibility for a life. Material security matters because hunger, untreated illness and constant insecurity can narrow a person's choices. Yet a comfortable life can still be aimless, lonely or dependent. A serious account of progress must make room for both material conditions and the qualities people develop through living.

This understanding has deep roots. Aristotle described flourishing as an active life shaped by reason and practiced virtue, with friendship and sound judgment among its essential goods. The modern capability approach asks what people are genuinely able to do and become, directing attention to the conditions that make meaningful choice possible. These traditions give the AI question a human focus: what abilities are being developed, and what opportunities can people actually exercise? [Aristotle's Ethics and The Capability Approach](#), Stanford Encyclopedia of Philosophy.

This is why productivity matters, and why it cannot carry the whole argument. Producing useful goods and services with fewer resources can make essentials more affordable and give people room to pursue other ends. Whether those possibilities become real depends on what happens to prices, earnings, working conditions and time. A faster process brings little relief to a family if its only consequence is a higher work quota.

Human dignity also reaches beyond productive capacity. A person's worth does not rise or fall with a wage, a test score or an ability to work independently. Children, people with disabilities and those who need care belong fully within the purpose of progress. Assistance can be an expression of dignity, even when independence is limited.

Ethics, agency and productivity belong together. Ethics concerns the ends we should pursue and the means we may use. Agency is the practical ability to understand, choose and act. Productivity expands what our efforts can accomplish. AI should support all three. The people who build and use it remain answerable for the purposes it serves. No institution needs authority to prescribe a single good life for everyone; it does need to respect the freedom and responsibility through which people pursue their own.

The reading route

Obstacles and costs
Putting human purposes in charge
How AI expands human potential
Benefits that reach people
Reference survey and sources

OBSTACLES

What can hold human development back

Many obstacles are familiar: preventable suffering, poor education, insecure livelihoods and unreliable access to essential services. Others concern how we live together. Isolation can deprive a person of belonging. Manipulation can undermine informed choice. Work that offers no opportunity to learn can leave people with fewer options when circumstances change. AI deserves attention where it can help remove such barriers, with careful examination of any new dependence it creates.

Learning makes that danger concrete. In a randomized study involving nearly 1,000 high-school mathematics students in Turkey, a general GPT-4 interface improved performance during assisted practice. When the assistance was removed, those students scored 17 percent below the control group on subsequent exams. A teacher-designed version that supplied tutoring guidance avoided the detected harm, but did not establish an improvement in unaided exam performance. The study covered four sessions in one school, so its findings should not be turned into a verdict on every educational use of AI. It does show why a completed assignment can be a misleading measure of learning. [Bastani and colleagues, PNAS, 2025](#)

The same question belongs in workplaces and public services. Does assistance help people acquire competence, or leave them unable to judge the answer? Can they challenge a consequential decision, or must they accept what the system produces? Convenience can conceal a loss of skill or authority that becomes visible only when something fails.

Economic and institutional choices create further risks. An employer might use saved time to improve service and train staff, or intensify work while reducing opportunities to learn. A tool might help a small enterprise compete, or bind it to a supplier it cannot readily leave. These are different possible uses of the technology, not predetermined outcomes. Refusing a useful tool can also carry costs when it would reduce avoidable suffering or relieve a genuine burden. Prudence requires examining both choices.

Costs that affect the whole person

Four kinds of cost deserve attention wherever AI is introduced. Each can affect the conditions in which people live and the choices they retain.

Resources

Count the electricity, water, land, materials and upkeep a system requires, including costs borne by local households and communities.

Human capacity

Consider skills gained or lost, time, working conditions, relationships and the ability to function when the tool is unavailable.

Power

Ask who controls the system and its data, what choices people retain, and whether they can leave a provider or challenge a decision.

Security

Examine exposure to deception, cyberattack, unreliable data and failure, together with the means of recovering safely.

These considerations require judgment. A gain in one area does not automatically justify a loss of dignity or freedom in another.

PRACTICAL RESPONSIBILITY

How to put human purposes in charge

The practical starting point is a human need that someone is responsible for meeting. A school should decide what students must learn before choosing its AI tools. A clinic should identify the barrier preventing people from receiving care. An employer should explain what improvement it seeks and how employees will be prepared for the change. Starting this way keeps the institution's purpose visible when a new capability becomes available.

For education, preserving effort is part of the design. Teachers can ask students to attempt a problem, explain their reasoning and then use assistance to identify a misunderstanding. Independent work remains necessary to establish what students can do themselves. Families also need space for conversation, reading and shared activities in which attention is given to another person. AI may help with preparation or routine administration; decisions about children's formation should remain with responsible adults who know them.

A randomized field trial of Tutor CoPilot tested another approach: giving AI suggestions to human tutors while they worked with students. The revised working paper reports a four-percentage-point increase in passing immediate lesson assessments, but no significant improvement on end-of-year standardized tests. Assistance supported the tutor's teaching, while the immediate assessments measured work students completed without the tutor's help. Wang and colleagues, Tutor CoPilot working paper, revised 2025



A teacher helps a student think through a problem, with AI supporting the learning. Editorial illustration; not a scene from the cited study.

PRACTICAL RESPONSIBILITY · CONTINUED

Build capability and preserve authority

At work, adoption should include a plan for developing people. Employees need time to understand a tool's limitations, practise checking its work and learn tasks that become more valuable as routine work changes. Businesses should consider how gains can support better service, competitive prices and sustainable employment. Where displacement occurs, credible preparation and routes into new work matter more than a general assurance that other jobs will eventually appear. Enterprise and personal initiative need room to respond, including practical access for smaller firms.

The same practical attention should extend to farms, local utilities and rural communities. Useful advice needs to fit local circumstances, and adoption must include people who can maintain the system and act on its results. Strengthening productive capacity means building that ability in ordinary institutions, as well as developing more powerful technology.

Human oversight must involve real authority. The person responsible for a consequential decision needs enough knowledge and time to question the recommendation, access to relevant evidence and the ability to intervene. People affected by mistakes need a usable way to obtain correction. Privacy, freedom from arbitrary coercion and lawful treatment cannot simply be traded away because a system produces an attractive total return.

Public policy should support those responsibilities while leaving room for experimentation and competition. It should address concrete risks, protect legitimate choices and avoid making useful tools accessible only to institutions able to absorb heavy costs. Providers and purchasers should also account for the energy, water, infrastructure and security demands of the systems they deploy. Local communities should be able to understand costs placed on them. A durable improvement must be supportable by the people and resources on which it depends.


CONCEPTUAL PATHWAYS

How AI expands human potential

LEARNING

From one question to a personalized path.

WITHOUT AI: expertise within available time and attention



More useful practice. Shorter feedback loops.

Human check: solve without AI; teacher verifies.

Conceptual pathways. Benefits depend on good design, appropriate use and human judgment. These are possible benefits, not measured study results.


CONCEPTUAL PATHWAYS

How AI expands human potential

BUSINESS WORK

From one need to a wider field of possibilities.

WITHOUT AI: expertise within available time and attention



More tested options. More time for high-value work.

Human check: validate, decide and measure customer results.

Conceptual pathways. Benefits depend on good design, appropriate use and human judgment. These are possible benefits, not measured study results.

THE EVIDENCE IN PRACTICE

The benefits that would make progress real

There is a practical basis for this ambition. A study of a staggered AI rollout to 5,172 customer-support agents found an average 15 percent increase in issues resolved per hour. Less experienced and initially less proficient workers gained most; the most skilled workers saw small declines in some quality measures. The study suggests a way that assistance can put useful expertise within reach of people still learning a job. It does not establish that workers received higher wages or that the gains became broadly shared prosperity. Those further outcomes depend on decisions the productivity figure cannot settle. [Brynjolfsson, Li and Raymond, The Quarterly Journal of Economics, 2025](#)

For an employer, this creates a concrete opportunity: connect assistance with training so that a beginner becomes a more capable worker. For the worker, the deeper benefit would be greater competence and better prospects, alongside useful help with today's task. Looking beyond the immediate output changes how a tool should be introduced and what its users should expect from it.



An architect uses digital assistance while retaining professional judgment. Editorial illustration; the cited workplace study involved customer-support agents.

THE EVIDENCE IN PRACTICE

Care that reaches people

Healthcare offers another example of what becomes possible when technology fits a real need. In the ACCESS randomized trial of 164 young people with diabetes, offering AI retinal screening during an existing diabetes-clinic visit produced 100 percent screening completion in the intervention group. Completion was 22 percent in the group referred for a separate eye examination over six months. The intervention combined AI with convenient, on-site care, so the improvement cannot be assigned to the algorithm alone. This was an investigational use in that age group. It established greater screening access, not reduced vision loss; only 16 of the 25 participants referred after AI screening completed their eye-care follow-up. Completing the care pathway remained essential. [Wolf and colleagues, Nature Communications, 2024](#)



Clinical judgment remains a human responsibility. This MRI scene is an editorial illustration, not a depiction of the retinal-screening trial.

PUBLIC AND PERSONAL BENEFIT

From capability to a changed life

Other benefits should be described with the same care. A more accurate forecast matters to a family when a warning reaches them in time and they have somewhere safe to go. Better agricultural advice matters when it helps a farmer make a sound decision at an affordable cost. Scientific discoveries can open possibilities for treatment or more reliable energy, but practical development, delivery and access still stand between a discovery and a changed life. These are pathways to human benefit, not promises that every application will succeed.

Consider an emergency crew responding to a flood. A useful system might help the team interpret changing conditions and identify a route, while trained responders judge whether it is safe and decide whom to assist first. This is an illustrative pathway, not a reported result from a particular deployment. Its value would lie in timely help reaching people, with reliable communications and accountable command supporting the crew.

Some of the most valuable outcomes may be ordinary. A caregiver could spend less time on paperwork and more time with the person who needs care. A small business owner could have more time to serve customers and train an apprentice. Such possibilities deserve investigation alongside large scientific advances. Time saved becomes part of flourishing when people have genuine freedom to use it well.



A firefighter uses decision support while assessing conditions and choosing how to help. Editorial illustration of a possible use.

The human measure of progress

AI can contribute to a society in which more people are healthy, capable and able to take responsibility for their lives. Achieving that requires deliberate choices about education, work, institutions and the uses to which power is put. Families and communities have purposes that technology should serve. Human beings still have to practise honesty, develop judgment, accept responsibility and care for one another.

The strongest case for AI is the opportunity to support that fuller life. We should pursue useful advances with confidence where the evidence warrants it, correct failures and preserve the freedom to choose a better approach. Success will be found in people who can do more of what matters and in communities better able to sustain them.

APPENDIX · WIDER APPLICATIONS

Reference survey of twelve benefit domains

This survey separates current applications, prospective gains and the evidence needed to establish a human benefit. A deployed tool does not by itself establish better lives. The entries offer a wider field of inquiry alongside the worked cases in the main argument.

How to read the evidence

These categories describe evidence maturity for a specific use or claim, not a blanket grade for a domain or proof of net benefit.

A Operational result

Used in real settings with an observable output or outcome. Deployment alone does not prove net benefit.

B Demonstrated pilot

Tested with credible evidence, but scale, generalizability, cost or long-term performance remains uncertain.

C Funded development

An identified program, investment or research path exists, but the claimed benefit is not yet delivered.

D Speculative claim

Plausible in principle but unsupported by a validated system, accountable plan or measurable result.

Each domain separates what is used now, what might improve next and what evidence would establish public value. Read individual claims against the A–D evidence scale above: “In use now” can include operational results or pilots; “Credible next gains” are possibilities until measured. A tool being deployed does not by itself prove a net benefit.

01 Human suffering & health

In use now

Clinical imaging support, risk detection, workflow assistance, outbreak surveillance, drug-target discovery, translation and assistive interfaces are already deployed in bounded settings.

Credible next gains

Earlier diagnosis, individualized treatment, faster therapeutic design, wider access to scarce expertise and more responsive public-health systems.

GIB verification test

Compare patient outcomes—not model accuracy alone. Require representative data, clinician authority, informed consent, error reporting, appeal and proof that access expands rather than stratifies.

02 Quality of life & disability

In use now

Speech recognition, live captions, visual description, language translation, navigation, prosthetic control and administrative assistance can increase independence and reduce routine burdens.

Credible next gains

More natural assistive devices, home safety monitoring, personalized rehabilitation and lower-cost access to education, legal information and government services.

GIB verification test

Measure retained independence, affordability, reliability and user control. A tool that creates surveillance, dependency or an inaccessible failure mode may shift rather than reduce suffering.

APPENDIX · WIDER APPLICATIONS

The twelve benefit domains

03 Earth, ecosystems & biodiversity

In use now

Machine learning helps classify satellite and sensor data, detect land-use change, track species, locate pollution and direct limited conservation resources.

Credible next gains

Near-real-time ecosystem accounting, earlier detection of illegal extraction, better restoration targeting and faster discovery of lower-impact materials.

GIB verification test

Verify observations on the ground, disclose blind spots and compare environmental benefit with computing, mineral, land and water costs.

05 Water access & security

In use now

Utilities use analytics and AI to detect leaks, anticipate equipment failure, optimize pressure and treatment, forecast demand and identify contamination anomalies.

Credible next gains

Lower non-revenue water, more efficient irrigation, adaptive treatment and desalination, drought allocation support and affordable service expansion.

GIB verification test

Report actual water saved, service reliability, energy use, maintenance cost and household affordability. Strong data governance must precede automation.

04 Climate resilience

In use now

AI-enhanced forecasting and nowcasting can improve the speed and local detail of warnings for storms, floods, wildfire conditions and extreme heat.

Credible next gains

Faster climate-risk analysis, resilient infrastructure planning, methane detection and better integration of weather-dependent energy.

GIB verification test

Forecast skill must translate into warnings people receive, trust and can act on. Keep physical models, expert forecasters and public warning authorities in the loop.

06 Science & discovery

In use now

AI predicts protein structures, searches chemical and materials spaces, controls experiments, analyzes massive datasets and helps scientists find patterns faster.

Credible next gains

New medicines, catalysts, batteries, resilient crops and materials; automated laboratories may compress discovery cycles from years to months in selected fields.

GIB verification test

Predictions become knowledge only after reproducible experiments. Track failed hypotheses, provenance, independent replication and whether researchers can interrogate the model's reasoning.

APPENDIX · WIDER APPLICATIONS

The twelve benefit domains

07 Food quality & supply

In use now

Computer vision and sensor models monitor crops, soil, livestock, pests, disease, storage and supply chains; decision tools can time irrigation and inputs more precisely.

Credible next gains

Higher yields with fewer inputs, earlier crop-disease response, less spoilage, improved traceability and locally adapted advice for small farmers.

GIB verification test

Measure nutrition, safety, yield stability, farmer income, soil and water—not output volume alone. Protect farm data and avoid systems that deepen vendor dependence.

09 Energy production & reliability

In use now

AI improves demand forecasts, equipment maintenance, subsurface analysis, grid monitoring, dispatch support and renewable-generation forecasts.

Credible next gains

Faster materials discovery, better storage, autonomous plant operations, improved nuclear and geothermal development and more efficient industrial processes.

GIB verification test

Count delivered energy, reliability, safety, total-system cost and emissions. Net claims must include the electricity and infrastructure consumed by AI itself.

08 Economic prosperity & work

In use now

AI assists software, design, customer service, logistics, research, documentation and many professional tasks; gains are strongest where work is reorganized around the tool.

Credible next gains

Higher productivity, new products, faster business formation and broader access to advanced capabilities—alongside significant job redesign and displacement.

GIB verification test

Separate task speed from economy-wide productivity. Track wages, job quality, new-firm formation, concentration, regional distribution and whether workers share the gains.

10 Industry, manufacturing & logistics

In use now

Machine vision detects defects; predictive maintenance reduces downtime; digital twins, robotics, scheduling and route optimization improve factories and supply chains.

Credible next gains

More flexible domestic production, faster engineering, resilient supply networks, mass customization and safer high-risk work.

GIB verification test

Verify throughput, defect reduction, energy use, worker safety and resilience. Preserve manual fallback and defend connected industrial systems from cyberattack.

APPENDIX · WIDER APPLICATIONS

Security and accountable command

11 U.S. & global security

In use now

AI supports intelligence analysis, cyber defense, anomaly detection, logistics, translation, disaster response and monitoring of illicit networks.

Credible next gains

Faster warning, stronger infrastructure defense, better supply-chain visibility and improved coordination among allies—while adversaries use the same tools.

GIB verification test

Demand lawful authority, audit trails, red-teaming, counter-deception, human accountability and resilience when communications, models or data are compromised.

12 Military capability

In use now

AI assists sensing, planning, maintenance, logistics, cyber operations, uncrewed systems, battlefield medicine and decision support.

Credible next gains

Human-machine teams may compress decision cycles, increase precision, protect service members and accelerate capability development and industrial production.

GIB verification test

Strategic advantage cannot erase the laws of armed conflict. Lethal force, escalation and target validation require accountable human command, secure systems and rigorous testing under adversarial conditions.

Sources for the four worked cases

Bastani and colleagues — Generative AI without guardrails can harm learning

PNAS, 2025. Classroom-randomized study distinguishing assisted mathematics performance from subsequent unaided learning.

Wang and colleagues — Tutor CoPilot

Revised working paper, 2025. Tutor-facing assistance and immediate student mastery; no significant end-of-year test improvement.

Brynjolfsson, Li and Raymond — Generative AI at Work

The Quarterly Journal of Economics, 2025. Staggered rollout to 5,172 customer-support agents; measured productivity and quality outcomes.

Wolf and colleagues — The ACCESS randomized trial

Nature Communications, 2024. AI retinal screening integrated into diabetes care; screening and follow-up outcomes.

APPENDIX · SOURCES

Primary and authoritative starting evidence

Sources for the worked cases are listed first, followed by references for the broader domain survey. Institutional forecasts remain claims; operational results require outcome evidence and independent verification where available.

01 Stanford Encyclopedia of Philosophy — Aristotle's Ethics

Aristotle's account of flourishing as reason-guided activity, virtue, practical wisdom, friendship and participation in a well-lived community.

02 Stanford Encyclopedia of Philosophy — The Capability Approach

The modern framework associated with Amartya Sen and Martha Nussbaum for judging well-being by people's substantive freedoms and opportunities.

03 World Health Organization — Harnessing AI for health

Current uses in diagnosis, clinical care, drug development, surveillance and health-system management; equity warning.

04 WHO — Governance of large multimodal models for health

Risk, evidence and governance requirements for generative and multimodal systems in health.

05 World Meteorological Organization — AI action plan

Operational integration, early warnings, global access and remaining forecasting gaps.

06 WMO — AI-powered nowcasting

Evidence and pilots for rapid local weather prediction and warning.

07 Food and Agriculture Organization — Digital agriculture in action

Crop, soil and livestock monitoring and decision support.

08 FAO — AI can be a game-changing solution for farmers

Prospective agricultural foundation models, farmer advice and climate strategy.

09 World Bank — Scaling digital water solutions

AI-enabled utility operations, leak detection, data governance and measured scaling.

10 OECD — AI, productivity, distribution and growth

AI as a possible general-purpose technology and the uncertainty around broad economic gains.

11 OECD — Macroeconomic productivity gains in G7 economies

Scenario-based estimates connecting task gains, exposure and adoption.

12 U.S. Department of Defense — 2026 AI strategy

Current U.S. military strategy for AI-enabled warfare and capability development.

13 U.S. Department of Defense — AI's role in national security

Defense adoption priorities and strategic competition context.

14 Government Accountability Office — AI accountability framework

Governance, data, performance and monitoring tests applicable across public uses.

15 MIT — Generative AI Impact Consortium

Daniel Huttenlocher's research questions on human-AI collaboration, behavioral change and improved life; a framework for testing, not evidence of achieved outcomes.

16 NIST — AI Risk Management Framework

Voluntary framework for trustworthy design, deployment and monitoring.

17 International Energy Agency — Energy and AI

AI's energy demand, grid implications and opportunities across energy systems.

APPENDIX · FURTHER READING

Related GIB investigations

U.S. Leadership in Artificial Intelligence

Connect measurable human benefit to the national objective, implementation record and unresolved constraints.

Fear, Power & the Future of AI

Which AI fears are rational, who advances them and how human control can be preserved.

Global AI Power & Governance

Which countries lead in capability, infrastructure, law and deployment—and what falling behind would mean.

Where Society Meets AI Super-intelligence

How AI works, who directs it, and how truth and human responsibility are preserved.

Data Centers & Public Value

The power, water, land, jobs, security and ratepayer costs beneath AI.

The Well-Formed Mind in the AI Age

Whether AI strengthens human judgment and learning—or substitutes for them.

Nuclear Got A Bad Wrap?

Firm power, deployment evidence and the energy foundation for an AI-intensive economy.