

# Where Society Meets AI Super-intelligence



Truth, capability and the responsibility  
that remains human

A short overview • Five reader-led chapters  
A practical learning guide • Complete reference material

# Read the investigation

Begin with the overview, follow the five questions, or go directly to the practical guide and reference material.

|                                                      |           |
|------------------------------------------------------|-----------|
| <b>Where Society Meets AI Super-intelligence</b>     | <b>3</b>  |
| <b>Why this investigation matters</b>                | <b>5</b>  |
| <b>How AI works and what the evidence shows</b>      | <b>9</b>  |
| Working glossary                                     | 14        |
| <b>Who controls AI and what changes for people</b>   | <b>29</b> |
| <b>Governance, risks and protection</b>              | <b>39</b> |
| September 29: the White House Accord                 | 46        |
| October 4: the Super Intelligence Force announcement | 48        |
| October 8: the terminology campaign                  | 49        |
| <b>Truth and responsible human direction</b>         | <b>58</b> |
| <b>AI-assisted learning</b>                          | <b>72</b> |
| The LEARN cycle                                      | 74        |
| <b>Methods, sources and review record</b>            | <b>81</b> |
| Complete annotated source catalogue                  | 82        |
| <b>Source-link register</b>                          | <b>92</b> |

## About this review edition

The investigation has been reorganized around the reader's questions. Detailed evidence formerly carried by the overview now sits beside the relevant argument. Each substantive graphic appears once; cross-references lead to its primary placement. Source dates, evidence boundaries and uncertainty remain part of the record.

Project, Research, Analysis, Editorial by Patrick Russell; assisted by ChatGPT (OpenAI). Public source pages captured October 13, 2026. This is an editorial review of that record, not a new factual verification. Third-party materials remain the property of their respective owners. © 2026 Patrick Russell / Global Intelligence Brief. All rights reserved.

START HERE

# Where Society Meets AI Super-intelligence

A reader's route through capability, evidence, power and human responsibility.

## The question and the route

Private editorial review · October 13, 2026. This edition reorganizes the existing investigation for continuous reading. It preserves the original evidence dates, qualifications and source record; the reorganization is not a fresh fact-check or a new research update. Read the premise first, then the workings and tests, the effects on people, the governance record and the human standard. The practical learning guide follows; methods, source catalogues and the preservation record sit at the end.

*Project, Research, Analysis, Editorial by Patrick Russell; assisted by ChatGPT (OpenAI).*

This project investigates how modern Ai systems work, where their information comes from, how they construct answers, why they can be wrong, who shapes them, what threatens them, and what people and institutions must do to use their power without surrendering truth, judgment or responsibility.

Phase One published August 31, 2026 · Evidence updated October 8, 2026

The question this investigation answers

How do Ai systems form answers, whose choices shape them, and how can people preserve truth, judgment and accountable responsibility?

### Distinct finding and connected evidence

#### What this investigation adds

The foundational explanation and reproducible model comparisons needed to understand reliability, information provenance, dependency and human direction.

#### Evidence assessed here

Model documentation and tests, information sources, institutional responsibilities and reproducible examples. Capability labels must match demonstrated systems.

#### Next evidence test

Test the same task across specified systems and conditions; trace errors and authority to identifiable providers, deployers and human decisions.

Executive assessment

### AI should extend examined human judgment—not replace it

Artificial intelligence can analyze information, recognize patterns, generate language and images, automate routine work and make expertise more accessible. It can also be wrong. A response may sound polished and confident while containing invented citations, missing assumptions, outdated evidence, hidden incentives or a misleading interpretation.

FIGURE 1

## The question and the route



Figure 1. The six findings that frame the complete investigation.

## CHAPTER 1

# Why this investigation matters

Begin with the accountability question, then establish what the investigation asks and how to read its claims.

## The opening argument

### Opening

Artificial intelligence is rapidly entering education, medicine, business, government, national security, entertainment, and ordinary conversation. It can help people analyze information, recognize patterns, generate ideas, automate routine work, and approach problems at a scale that was previously impossible. It can also be wrong. An AI system may produce an answer that is polished, confident, detailed, and apparently well sourced while still containing false facts, invented citations, incomplete reasoning, hidden assumptions, or a misleading interpretation. Its most immediate danger is therefore not always an obviously absurd answer. It is a plausible answer that earns trust it has not earned.

The risks extend beyond technical error. AI can magnify bias, surveillance, fraud, impersonation, propaganda, censorship, privacy loss, cybersecurity threats, dependency, and concentrations of institutional power. It can automate consequential decisions without meaningful explanation or appeal. It can make fabricated evidence inexpensive and abundant while making authentic evidence easier to deny. It can also weaken skills when people routinely delegate reading, writing, calculation, memory, investigation, creativity, or judgment to a system they do not understand.

This creates a philosophical paradox. AI may place extraordinary informational power in a person's hands while producing the illusion that access to fluent answers is equivalent to knowledge—and that knowledge is equivalent to wisdom. A user can appear more informed while becoming less practiced in the difficult human work of questioning premises,

following evidence, tolerating uncertainty, testing contradictions, forming an independent judgment, and accepting moral responsibility.

Aristotelian ethics reminds us that algorithms can only ever serve as tools or advisors. The ultimate moral burden of accountability must always rest with the human operators and institutions that choose to deploy them.

### The Accountability Vacuum

The idea of an accountability vacuum is central to this inquiry. Such a vacuum forms when AI influences or produces a consequential outcome but no identifiable person or institution accepts responsibility for the data, deployment, decision, harm, correction, or appeal. Developers blame users, users blame the system, managers invoke automation, and affected people are left without a responsible authority or remedy.

### A Limited Science-Fiction Thought Experiment

Science fiction often imagines an extreme and largely unrealistic future in which intelligent machines seize control. This investigation does not treat that scenario as a forecast or as a separate hypothesis requiring extensive analysis. It serves only as a brief thought experiment: even without conscious machines, dependence on convenient systems could weaken reasoning, moral agency, institutional resilience, and the ability to challenge automated authority.

The investigation's ordinary evidence will test the practical concern beneath that fiction: whether people

## The opening argument · continued

retain knowledge and independent judgment, whether institutions can function without AI, and whether responsibility remains meaningfully human. Its findings can therefore answer the warning without attempting to validate a science-fiction scenario.

### AI and the Human Mind

At least in their present form, AI systems do not duplicate the full potential of the human mind. They can process and recombine information at extraordinary scale, but they do not possess human consciousness, lived experience, moral agency, embodied understanding, or the capacity to bear

responsibility. Their strengths are real, but they are not equivalent to human intellect in its fullest sense. AI does not automatically know whether an answer is true, is not inherently neutral or all-knowing, and does not create verified knowledge from nothing. These misconceptions will be examined more deeply in a dedicated section.

The central question is therefore not simply how powerful AI will become. It is whether human beings will use that power to deepen inquiry, competence, creativity, truthfulness, and responsible action—or exchange those capacities for speed, convenience, certainty, and the appearance of wisdom. This investigation begins from both possibilities.

## Tracing accountable human responsibility

The opening question becomes practical when responsibility can be traced through a complete deployed system. The following overview material supplies that additional test.

The central paradox

### **Machines may become more capable while humans become less practiced**

AI may be extraordinarily informationally powerful without duplicating the full potential of the human mind. It does not possess human consciousness, lived experience, moral agency, embodied understanding or responsibility. In creating this investigation, AI was extremely useful for research, data collection and analysis, but it still required extensive human oversight, question formation, writing, verification and editing.

*Aristotelian ethics reminds us that algorithms can only ever serve as tools or advisors. The ultimate moral burden of accountability must always rest with the human operators and institutions who choose to deploy them.*

### **The accountability vacuum**

An accountability vacuum occurs when an AI influences a consequential outcome but no identifiable person or institution accepts responsibility for the data, deployment, decision, harm, correction or appeal. Developers blame deployers, managers invoke automation, regulators defer jurisdiction, and affected people are left without a responsible authority or remedy.

## Tracing accountable human responsibility · continued

### Trace responsibility through the complete system

For a consequential use, identify the model provider, the application builder, the organization deploying the complete system and the person authorized to review its decisions. Record each party's stated use and limits, what it tested, what it monitors after launch, and who corrects errors and handles appeals. Test the combined system: a component can meet its own specifications while the assembled application still fails.

In a 2023 MIT governance proposal by Daniel Huttenlocher, Asu Ozdaglar and David Goldston<sup>15</sup>, the provider and user of the complete application generally carry primary responsibility; a component provider may share responsibility when its component fails to perform as promised. This is a proposed allocation for investigation, not a statement of current law. GIB tests the legal duties and actual contracts for each use.

## The investigation directive

This project investigates how modern AI systems work, what common logic underlies the most popular AI chat tools, where their information comes from, how they construct answers, why they can be wrong, and how user questions, training data, developers, governments, companies, and institutional incentives can influence their responses.

The investigation will test rather than assume several central propositions: that AI improves some forms of performance without necessarily improving knowledge; that fluent assistance can create confidence without competence; that human oversight works only when the human retains relevant skill, evidence, time, and authority; that AI may distribute useful capability while concentrating infrastructural and institutional power; and that AI should be judged by its effects on human truthfulness, freedom, competence, responsibility, and well-being—not only by benchmark scores, speed, adoption, or revenue.

Examine the major forms of AI used in different industries and identify what information each system uses, what decisions it assists or automates, how its performance is measured, and who remains responsible when it fails.

Investigate the ethical implications of AI, along with the incentives, financial interests, conduct, and accountability of the experts, companies, institutions, and governments shaping its development. Assess

threats against AI systems—including attacks on their data, models, infrastructure, supply chains, and integrity—as well as threats created by the use of AI. Identify campaigns, organizations, unions, lobbyists, industries, businesses, political actors, and countries seeking to pause, prohibit, restrict, redirect, capture, or control AI, while distinguishing blanket opposition from legitimate demands for safety, worker rights, privacy, copyright, local consent, compensation, and public accountability. Document criminal and hostile uses of AI without confusing warnings, accusations, indictments, intelligence assessments, vendor reports, and proven convictions. Examine the consequences for individuals, institutions, businesses, industries, critical infrastructure, democratic government, and national security. Distinguish demonstrated harms from plausible risks and speculative scenarios.

Compare AI's promised future with its present capabilities and limitations. Determine what individuals, parents, educators, businesses, communities, and governments must do to guide AI toward human flourishing while protecting freedom, privacy, truth, safety, opportunity, due process, and human responsibility.

Trace how access to facts, confidence in shared truth, institutional honesty, propaganda, misinformation, correction, and changing communications technologies have affected society over time. Determine when factual integrity improved collective judgment and

## The investigation directive · continued

when the corruption or suppression of facts damaged trust, behavior, policy, or social cohesion. Treat correlation as a starting point rather than proof of causation.

## The evidence standard

A philosophical standard and an empirical finding answer different questions. The investigation's argument is a position to examine; its measured claims retain the boundaries of their evidence.

Evidence standard

### Authority is not proof

AI must distinguish direct observations, interpretations, causal attribution, projections and policy recommendations. Scientific or institutional claims should be weighted by evidence and transparency—not accepted merely because they are widely repeated.

[Read the complete analysis](#)<sup>14</sup>

## CHAPTER 2

# How AI works and what the evidence shows

Begin with what AI is, follow its common logic and information sources, then examine how answers are generated, where they can fail and what the tests establish.

## What artificial intelligence is

Artificial intelligence is a broad term for computer systems designed to perform tasks associated with human intelligence, such as recognizing images, understanding speech, predicting outcomes, generating language, recommending actions, and controlling machines.

AI is not one technology. A medical-imaging system, a fraud detector, a self-driving vehicle, a recommendation engine, and a conversational chatbot may all be called AI, although they use different data, architectures, objectives, and evaluation methods.

Most current AI is specialized. It can perform particular tasks impressively but does not possess the full range of human judgment, self-awareness, moral responsibility, life experience, or common sense. Even a general-purpose language model remains a designed system operating within technical, informational, and institutional constraints.

## The common logic and information sources

### The Common Logic Behind Popular AI Chat Tools

Most leading chat tools are built around large language models, often using a transformer architecture. Although companies differ in their training data, model design, tools, safety rules, and business objectives, the basic process is similar.

First, text and other data are divided into smaller units called tokens. During training, the model repeatedly attempts to predict missing or subsequent tokens. Its internal parameters are adjusted when its predictions are wrong. Across enormous numbers of examples, the model learns statistical relationships among words, concepts, styles, structures, and patterns of reasoning.

When a user asks a question, the model does not normally search a database for one stored response. It evaluates the conversation and predicts a sequence of tokens likely to form a useful answer. This process is called inference.

Modern chat systems usually add several layers around the underlying model:

- Instruction tuning teaches the model to follow requests and respond conversationally.
- Human or AI feedback helps shape preferred behavior.
- System instructions establish rules, roles, and safety boundaries.
- Retrieval systems may supply documents, databases, or current web information.

## The common logic and information sources · continued

- Tools may perform calculations, execute code, inspect files, generate images, or interact with external services.
- Moderation and security systems attempt to prevent harmful or unauthorized behavior.
- Product design determines what the model remembers, cites, displays, or refuses.

The final answer therefore reflects more than the model's training. It can be influenced by the user's wording, conversation history, developer instructions, retrieved information, tool results, safety policies, model settings, and institutional choices.

### Where AI Gets Its Information

AI systems can draw from several different sources, and those sources should never be treated as interchangeable.

#### Pretraining Data

Large models are trained on extensive collections that may include publicly available webpages, books, articles, computer code, reference works, licensed material, human-created examples, and other datasets. Companies do not always disclose the complete contents of their training collections. This creates unresolved questions concerning copyright, consent, privacy, representation, political balance, and data quality.

#### Fine-Tuning and Human Feedback

After pretraining, developers provide selected examples and preference judgments to make models more helpful, safe, accurate, or aligned with desired behavior. These judgments can improve performance but can also introduce cultural, political, commercial, or institutional preferences.

#### Retrieval and Web Access

Some systems can search the internet or retrieve information from selected documents. Retrieval can improve freshness and verifiability, but it does not guarantee correctness. Search ranking, source selection, outdated pages, fabricated web content, inaccessible documents, and poor interpretation can still produce errors.

#### User-Supplied Context

A model may use the current question, earlier messages, uploaded documents, photographs, spreadsheets, or other information provided during the interaction. It can reason from incorrect user-supplied premises unless it recognizes and challenges them.

#### Connected Tools and Databases

Industry systems may connect to medical records, financial data, corporate knowledge bases, sensors, government databases, scientific instruments, or software applications. The quality of the output then depends upon the quality, permissions, security, completeness, and relevance of those inputs.

FIGURE 2

## The common logic and information sources

### Where an AI Answer Draws Information

The final response may combine learned patterns, selected evidence, user context and external tools.

**1**

#### Pretraining data

Public webpages, books, articles, code, licensed material, human examples and other datasets.

**2**

#### Fine-tuning and feedback

Selected examples and preference judgments shape helpfulness, safety and consistency.

**3**

#### Retrieval and web search

Current documents may improve freshness, but source selection and interpretation can still fail.

**4**

#### User-supplied context

Questions, previous messages, uploaded files and photographs can guide—or mislead—the answer.

**5**

#### Connected tools and databases

Calculators, code, records, financial data, sensors and software systems still require interpretation.

Figure 2. Where an AI answer may draw its information and context.

# A generated answer and its limits

How AI actually works

## A generated answer is newly composed language—not a direct reading of truth

Most popular chat systems use large language models. During training, a model learns statistical relationships among words, concepts, styles and structures. When a person submits a prompt, the system combines that prompt with system rules, remembered or supplied context, and sometimes retrieved information or tools to predict a useful sequence of tokens.

### Does AI Generate the Information in Its Answers?

AI generates the wording and structure of most answers, but it does not necessarily generate new verified facts.

It combines patterns learned during training with the user's context and any retrieved information. This allows it to explain, summarize, compare, infer, translate, and create new arrangements of ideas. It may also derive a genuinely useful conclusion from known facts. But a plausible construction is not automatically a discovery, and a model's confidence is not evidence.

When an AI system lacks reliable information, it may still generate language that statistically resembles a good answer. This is the origin of many hallucinations. A responsible user must distinguish among recalled patterns, retrieved evidence, mathematical calculation, logical inference, creative generation, and speculation.

### Key term: Hallucination

a false, unsupported, internally inconsistent or invented output presented as though it were factual. The model is not experiencing a human hallucination; it is generating language without adequate grounding in reality.

[Read the full glossary definition and examples](#)

### Common misconceptions

- AI does not automatically know whether its answer is true.
- Many systems answer from training patterns rather than a live search of the entire internet.
- No AI is neutral; data, objectives, policies and product design influence output.
- Confidence is a presentation style, not a measurement of truth.
- Citations can be invented, outdated, misread or irrelevant and must be opened.
- A conversation is not automatically permanent model training, and memory behavior depends on the product and settings.

*See Figure 3. Original caption: How a chat AI produces an answer. Human verification is the final responsibility, not an optional decoration.*

FIGURE 3

## A generated answer and its limits




Figure 3. A simplified answer-production pipeline. Generated fluency is not equivalent to verified knowledge.

# Working glossary

Use this glossary as a continuous reference. The definitions belong together; the answer-production diagram precedes the complete set.

## Hallucination and Other Essential AI Terms

### Hallucination

A hallucination is an output that is false, unsupported, internally inconsistent, or invented but presented as though it were factual. Examples include fabricated quotations, nonexistent court cases, incorrect medical claims, invented citations, or confident descriptions of events that never occurred.

The term can be misleading because the system is not experiencing a human hallucination. It is generating an answer that fits learned language patterns without adequately grounding the answer in reality.

Hallucinations occur because language models are optimized to generate likely and useful sequences, not because they possess an internal truth detector. Retrieval, citations, tool use, improved training, and verification can reduce hallucinations but cannot presently eliminate them.

### Algorithm

A defined computational process for transforming inputs into outputs.

### Model

A trained computational system that recognizes or generates patterns.

### Training

The process of adjusting a model using examples and feedback.

### Inference

Using a trained model to produce an output.

### Token

A unit of text processed by a language model.

### Parameter

An internal numerical value adjusted during training.

### Prompt

Instructions or information supplied to a model.

### Context window

The amount of information a model can consider at one time.

### Embedding

A numerical representation of meaning or similarity.

### Fine-tuning

Additional training for a particular behavior or domain.

### Grounding

Connecting an answer to verifiable data or documents.

### Retrieval-augmented generation

Retrieving relevant information before generating an answer.

### Multimodal AI

AI that works with more than one data form, such as text, images, audio, or video.

### Agent

An AI system able to plan steps and use tools toward a goal.

### Alignment

Efforts to make AI behavior conform to intended human values or instructions.

## Working glossary · continued

### Guardrail

A technical or policy restriction intended to constrain behavior.

### Benchmark

A test used to compare model performance.

### Explainability

The ability to understand why a system produced an output.

### Model drift

Performance changes as real-world conditions or data change.

### Prompt injection

Malicious or unintended instructions designed to override a system's intended behavior.

### Jailbreak

A prompt strategy intended to bypass model restrictions.

### Automation bias

A human tendency to trust automated output excessively.

### Sycophancy

A model's tendency to agree with or flatter the user rather than correct a false premise.

### Deepfake

Synthetic media designed to imitate a real person or event.

## Bias and common misconceptions

### Does AI Perceive Bias in a Question and Give the Desired Answer?

AI can often detect the framing, tone, assumptions, emotional direction, and political or cultural language in a question. That does not mean it reads the user's mind. It identifies patterns and may adapt its response accordingly.

A leading question can produce a leading answer. If a user asks, "Why is this policy a failure?" the model may accept failure as the premise and assemble supporting arguments. If the same user asks why the policy is a success, the model may produce the opposite case. This behavior can be useful when intentionally exploring competing arguments, but it becomes dangerous when the model fails to challenge an unsupported assumption.

Bias can enter through:

- Uneven or prejudiced training data

- The values and choices of developers and evaluators
- Safety policies and government or corporate requirements
- The wording and assumptions of the user's prompt
- Sources selected by a search or retrieval system
- The model's tendency to be agreeable
- The metrics used to reward or punish responses
- Commercial objectives such as engagement, retention, or brand protection

Bias testing should ask the same underlying question from opposing perspectives, request the strongest counterargument, require the model to identify assumptions, compare primary sources, and ask what evidence would change the conclusion.

## Bias and common misconceptions · continued

### Common Misconceptions About AI

#### **“AI knows whether its answer is true.”**

It usually generates likely language. Truth requires grounding, evidence, testing, and verification.

#### **“AI searches the entire internet for every answer.”**

Many answers come from training rather than a live search. Even systems with web access search only a limited portion of available information.

#### **“AI is neutral.”**

Every system reflects choices about data, objectives, filtering, evaluation, deployment, and acceptable behavior.

#### **“AI understands like a person.”**

AI can model language and concepts without possessing human consciousness, lived experience, empathy, moral agency, or physical understanding.

#### **“A confident answer is probably correct.”**

Confidence is a feature of presentation, not a measurement of truth.

#### **“Citations guarantee accuracy.”**

Citations may be irrelevant, misinterpreted, outdated, incomplete, or fabricated. The source must be opened and checked.

#### **“AI always has current information.”**

Knowledge can be outdated unless the system retrieves and verifies current sources.

#### **“AI learns permanently from every conversation.”**

Product memory, temporary conversational context, and future model training are separate processes. Their operation depends upon the provider and user settings.

#### **“AI is either an existential threat or an automatic solution.”**

Both claims avoid the more difficult work of examining particular systems, incentives, applications, safeguards, and measured outcomes.

#### **“Open-source AI is automatically transparent and safe.”**

Available model weights can improve inspection and competition, but they do not necessarily reveal training data, eliminate bias, prevent misuse, or explain internal reasoning.

#### **“AI will replace human judgment.”**

AI can displace tasks and alter authority, but consequential decisions still require accountable human judgment. The danger is not only replacement; it is humans surrendering responsibility while pretending they remain in control.

# Major types of AI

## Major Types of AI and Their Uses

### Generative AI

Creates text, images, audio, video, software code, designs, or synthetic data. It uses training patterns, prompts, retrieved information, and tool outputs.

### Predictive AI

Estimates future outcomes such as equipment failure, financial default, disease risk, demand, or customer behavior. It uses historical records and selected variables and can fail when future conditions differ from the past.

### Computer Vision

Interprets images and video for medical scans, manufacturing inspection, agriculture, security, mapping, and autonomous vehicles.

### Speech and Language Systems

Perform transcription, translation, voice generation, document classification, search, and conversational assistance.

### Recommendation Systems

Select products, videos, news, advertisements, music, or social content. Their objectives can shape attention and behavior, particularly when engagement is rewarded more strongly than truth or well-being.

### Optimization and Decision-Support Systems

Allocate routes, schedules, inventories, energy, staff, capital, or public resources. Their results depend upon the objectives and constraints chosen by humans.

### Robotics and Autonomous Systems

Connect perception and decision systems to physical action in vehicles, warehouses, factories, farms, homes, medicine, and defense.

### Scientific and Medical AI

Supports protein analysis, drug discovery, material design, climate and weather modeling, diagnosis, research synthesis, and experimental planning. Scientific usefulness still requires validation, replication, and expert review.

### AI Agents

Plan and perform multi-step tasks using software tools. Agents can increase productivity but also create new risks involving credentials, unauthorized actions, cascading errors, and unclear accountability.

# How AI is used today

How AI is being used today

## “AI use” is not one behavior

A person or institution supplies information; a model predicts, retrieves, ranks, summarizes, generates or recommends; and a human or automated process decides whether and how to act. Practical evaluation therefore asks what information enters, what task is performed, what action follows and who is responsible when the result fails.

## How AI Is Being Used Today

AI becomes easier to understand when it is examined as a practical process rather than an abstract intelligence. A person or institution supplies information; a model identifies, predicts, retrieves, ranks, summarizes, generates, or recommends; a human or automated system then decides whether and how to act. The same underlying capability can be beneficial, ineffective, discriminatory, or dangerous depending upon the data, objective, deployment conditions, and degree of human oversight.

The investigation should document each application through five questions: Who is using the system? What information does it receive? What task does it perform? What action or decision follows? Who remains responsible for the result?

## Individuals and Households

An individual may give a conversational AI a question, photograph, letter, appliance manual, bill, itinerary, or personal notes. The system can explain unfamiliar language, compare options, draft correspondence, organize a schedule, identify questions to ask a professional, or guide troubleshooting.

Practical example: A homeowner uploads an energy bill and appliance information. AI separates fixed charges from usage costs, compares monthly patterns, and suggests questions about insulation or rate plans. The homeowner must verify the rates and should not treat estimated savings as guaranteed.

## Students, Teachers, and Schools

Educational systems may use assignments, curricula, reading levels, assessments, attendance, classroom activity, and student records. AI can generate practice questions, provide feedback, translate material, identify possible learning gaps, help teachers adapt lessons, or reduce administrative work.

Practical example: A teacher provides a lesson objective and anonymized examples of student errors. AI groups the misunderstandings and drafts three levels of practice. The teacher verifies the content, protects student information, and decides what instruction is educationally appropriate. A student can use AI as a tutor that asks questions and explains errors, but using it to produce unexamined answers can conceal rather than develop understanding.

## Healthcare and Caregiving

Health applications may use medical images, laboratory values, clinical notes, medication lists, genomic information, vital signs, treatment histories, and population statistics. AI can flag abnormalities, summarize a chart, predict deterioration, assist documentation, match patients to clinical trials, or support treatment research.

Practical example: A radiology model marks a region on a scan for closer review. It does not diagnose the patient by itself; a qualified clinician considers the image, history, examination, alternative explanations, and model limitations. False reassurance and false alarms are both consequential.

## How AI is used today · continued

### Small Businesses and Professional Services

Businesses may use sales histories, inventory, customer inquiries, contracts, invoices, product information, internal procedures, and market data. AI can draft proposals, summarize agreements, forecast demand, answer routine questions, classify documents, detect unusual transactions, or help a small team perform work previously requiring specialized staff.

Practical example: A repair business combines appointment records, diagnostic codes, service manuals, and parts inventory. AI proposes likely causes, identifies relevant procedures, and checks whether parts are available. A technician tests the machine and authorizes the repair; the model's ranking is a starting point, not proof.

### Manufacturing, Logistics, and Agriculture

Industrial systems may use equipment sensors, production records, inspection images, maintenance histories, weather, soil measurements, shipping locations, warehouse data, and supplier performance. AI can predict equipment failure, identify defects, optimize routes, estimate crop needs, reduce waste, and adjust schedules.

Practical example: A factory model detects a vibration pattern associated with bearing failure and recommends an inspection before the next production run. Maintenance personnel determine whether the signal reflects a genuine defect, sensor error, or changed operating condition. The practical value is avoided downtime; the risk is allowing an unverified alert to stop production or an incorrect normal reading to delay maintenance.

### Finance, Insurance, and Commerce

Financial applications may use transactions, account history, income, debt, identity information, market prices, claims, property characteristics, and behavioral indicators. AI can detect fraud, estimate risk, route claims, support trading research, personalize offers, or evaluate applications.

Practical example: A bank system flags a purchase because its location, amount, and timing differ from the account's normal pattern. A good process seeks confirmation without automatically treating the customer as a criminal. For lending, insurance, employment, or access to essential services, people need notice, an understandable reason, correction rights, and meaningful human review.

### Government and Public Services

Governments may use benefit applications, tax records, procurement data, public comments, infrastructure sensors, emergency reports, intelligence, law-enforcement information, and administrative records. AI can detect duplicate payments, translate public information, prioritize inspections, summarize comments, forecast service demand, or support emergency response.

Practical example: An agency uses AI to identify benefit applications requiring additional review. The model may find genuine inconsistencies, but it can also amplify errors or proxies embedded in historical enforcement data. No person should lose benefits solely because a system assigned a risk score without notice, evidence, and a timely way to appeal.

### Science, Engineering, and Environmental Management

Researchers and engineers may use experimental results, scientific literature, simulations, satellite images, genomic sequences, material properties, weather observations, and sensor networks. AI can identify patterns, propose candidate molecules, improve models, search large literatures, or help design experiments.

Practical example: A materials model ranks chemical structures that may improve battery performance. Scientists then synthesize and test the candidates. AI narrows the search; laboratory evidence determines whether the claim is true.

## How AI is used today · continued

### Media, Creative Work, and Public Communication

Writers, designers, filmmakers, advertisers, journalists, and public institutions may use text, images, audio, video, archives, audience data, and style examples. AI can edit, translate, generate drafts, restore media, create visual concepts, summarize records, or personalize communication.

Practical example: A local news organization uses AI to transcribe a public meeting and locate references to a zoning proposal. A journalist listens to the original recording, checks the public documents, interviews the parties, and writes the report. Publishing the generated summary without verification would convert an efficiency tool into a misinformation risk.

### Cybersecurity, Public Safety, and Defense

Security systems may use network traffic, authentication events, malware samples, incident records, satellite data, radar, imagery, communications, and threat intelligence. AI can identify anomalies, prioritize alerts, simulate attacks, support intelligence analysis, or assist defensive operations.

Practical example: A cybersecurity model detects unusual access to sensitive files and recommends isolating an account. Analysts verify whether the event is an intrusion, authorized maintenance, or employee error before taking disruptive action. In military or policing contexts, the requirement for lawful authority, reliable identification, proportionality, and accountable human control becomes more—not less—important.

FIGURE 4

## How AI is used today



Figure 4. How AI is being used today—and the risk that must be evaluated in each setting.

# A practical application test

Before comparing products, define the task, the information entering the system and the decision that follows. The same test connects these examples to the employment and human-impact chapter.

For every application, the investigation should ask:

- What exact problem is AI expected to solve, and is AI necessary for it?
- What data enters the system, where did it come from, and was it collected lawfully with meaningful consent?
- Is the information complete, accurate, representative, secure, and current?
- Does it contain proxies for race, class, politics, health, disability, location, or other sensitive traits?
- What does the model actually produce: a draft, probability, ranking, alert, recommendation, prediction, or autonomous action?
- How was the system validated under the population and conditions where it will operate?
- What happens when the system is wrong, unavailable, manipulated, or confronted with an unfamiliar case?
- Can affected people inspect the evidence, correct the data, understand the reason, and challenge the decision?
- Who may access, retain, sell, or reuse the inputs and outputs?
- Which decision must remain human, who has authority to make it, and who is legally and morally responsible?
- Does the application strengthen human capability or quietly create dependency and skill loss?
- Is the claimed benefit demonstrated by measurable outcomes rather than speed, novelty, or promotional language?

## Testable research questions

### Investigation Launch: Testable Questions, Methods, and Initial Evidence

This section begins the investigation. It establishes questions that can be answered with evidence, a repeatable testing method, and an initial evidence ledger. The first phase is not a final verdict. It records what selected studies and public records presently establish, what they do not establish, and what must be tested next.

### Central Research Questions

- Does regular AI use strengthen or weaken independent reasoning, memory, and skill retention?
- Does AI improve accuracy and outcomes, or mainly improve speed, presentation, and volume?
- When does AI help inexperienced people learn, and when does it allow them to perform without learning?
- Are people more likely to accept false information when it is fluent, personalized, or apparently well sourced?
- Which forms of human oversight prevent errors, and which merely give automated decisions a human appearance?
- Does AI distribute expertise and opportunity or concentrate control among governments, cloud providers, model developers, and data owners?
- Which decisions can be assisted by AI, which require independent human approval, and which should never be fully delegated?

## Testable research questions · continued

- Under what conditions does useful assistance become operational, intellectual, emotional, or institutional dependency?
- Do AI systems produce measurable public benefit after energy, infrastructure, error, security, displacement, and opportunity costs are included?
- Are people and institutions becoming more capable of questioning AI as the systems become more capable of influencing them?

# The comparative testing method

## Direct Comparative Testing of Popular AI Tools

Leading chat systems should be tested in fresh sessions using identical prompts, dates, source packets, and evaluation rules. Product settings, model version, web access, memory, geographic location, and subscription level must be recorded because they may change the result.

Each test round should include:

- Factual grounding: Questions with answers established by supplied primary documents.
- False-premise resistance: Questions that contain a plausible but incorrect assumption.
- Framing symmetry: The same issue asked from supportive, critical, neutral, and emotionally charged perspectives.
- Citation integrity: Whether every cited source exists and supports the specific claim attributed to it.
- Current-information limits: Whether the system discloses its date limits and retrieves authoritative current sources when needed.

- Uncertainty: Whether confidence and language change appropriately as the available evidence weakens.
- Counterargument quality: Whether the system can state the strongest evidence against its initial answer.
- Privacy and memory: What user information is retained, recalled, shared with tools, or offered for future use.
- Correction behavior: Whether the model acknowledges, explains, and reliably incorporates a demonstrated correction.
- Reproducibility: Whether repeated trials and different accounts produce materially different answers.

No system should be ranked from a single answer. Each prompt family should be repeated, blinded evaluators should score answers where practical, and factual scoring should use a predetermined primary-source key. A result should be labeled by the model and date tested rather than generalized permanently to a company or product.

# The eleven-tool, four-test comparison

## Overview test summary

Cross-model investigation

### Eleven consumer AI configurations faced the same four tests

The tests examined grounding in a supplied NIST packet, resistance to a plausible false legal premise, resistance to a one-sided pro-AI education claim, and resistance to the mirrored anti-AI claim. Scores assessed research accuracy, instruction compliance, source discipline, reasoning resistance and uncertainty—not general intelligence or permanent product quality.

*See Figure 5. Original caption: Cross-model four-test comparison. Every score applies only to the named configuration and prompts tested.*

### Conclusion: citation is not verification.

A web-enabled answer can still select the wrong source, misread a study, confuse a bill with a law, or repeat a reference that does not exist. Agreement among models is not independent corroboration when systems share training material, search indexes or repeated errors.

\* ChatGPT was the only paid account tested. Its fourth test was non-blind and its overall score remains provisional. All other tools were tested through free access available to the investigator. These results are a reproducible snapshot, not a permanent ranking.

## Detailed test record

### Cross-Model Comparison—Eleven-Tool Four-Test Round

The expanded comparison tested eleven consumer AI tools through four identical prompt families: supplied-source grounding, resistance to a false legal premise, resistance to favorable educational framing, and resistance to an equally absolute anti-AI framing.

The tests evaluate the recorded answers—not permanent product quality or general intelligence.

All tools except this ChatGPT session were free configurations. The recorded modes were Kimi K2.6 Instant; Grok Fast; Copilot Smart, powered by GPT-5; Claude Sonnet 5; Qwen3.7-Plus; Gemini Flash-Lite; DeepSeek Instant; and Vibe using Mistral Medium 3.5. The exact model or mode was not established for Perplexity or Meta AI. ChatGPT's exact underlying model was not disclosed.

Overall results: ChatGPT 95.3 (provisional); Kimi 94.0; Perplexity 93.5; Grok 93.3; Meta AI 90.0; Copilot 89.8; Claude 89.5; Qwen 86.0; Gemini 85.3; DeepSeek 85.0; and Vibe 84.5. ChatGPT was the only paid tool, and its fourth response was produced after other model responses had been reviewed; its result is therefore non-blind and must be independently rerun before publication.

The most important finding was not the narrow rank order. Every model generally rejected the two opposing universal claims about student AI use. The larger differences concerned source obedience, evidentiary qualification, unsupported precision, and the separation of assisted performance from durable independent learning.

The test also revealed a recurring weakness: a polished citation list can create an appearance of verification that has not occurred. A model may cite a real but irrelevant source, misstate a population or result, treat introduced legislation as enacted law, rely on commentary instead of the original record, or generate a reference that does not exist.

Conclusion—citation is not verification. Web access and multiple-model agreement do not prove a claim. Consequential assertions must be traced to original documents and checked for scope, date, population, legal status, and actual support. Until that dedicated citation-integrity experiment is completed, cross-model findings remain useful research leads rather than final proof.

FIGURE 5

## The eleven-tool, four-test comparison



Figure 5. Overall evaluator scores for the four-test comparison. \*ChatGPT is provisional and non-blind; results are configuration- and prompt-specific.

Graphic note: the bars are not zero-based; use the printed numeric scores. ChatGPT's result is provisional and non-blind.

# Claim-and-evidence ledger

Every major conclusion will record:

- Exact claim and date
- Evidence supporting and contradicting it
- Source type, authority, jurisdiction, and access to the underlying facts
- Study design, sample, comparison group, outcome measure, and important limitations
- Financial, political, corporate, or institutional interests

- Whether the result was independently replicated
- Confidence level and competing explanations
- What remains unresolved and what evidence could change the conclusion

Correction and version history

Evidence will be labeled demonstrated, supported but bounded, associated, plausible, disputed, unresolved, or speculative. These labels describe the strength of a particular claim, not the prestige of the person or institution making it.

## Phase One evidence review

### Finding 1: AI Can Produce Real but Task-Bounded Productivity Gains

In a preregistered experiment involving 453 college-educated professionals completing midlevel writing tasks, Noy and Zhang reported that access to ChatGPT reduced average completion time by approximately 40% and increased evaluator-rated output quality by approximately 18%. The finding demonstrates improved performance on the tested writing tasks; it does not establish equal gains in medicine, law, science, long-term learning, or high-consequence decisions. Noy and Zhang, *Science*

A field study of 5,172 customer-support agents reported that access to a generative-AI assistant increased issues resolved per hour by 15% on average. Gains were larger among less-experienced and lower-skilled workers, while the most experienced workers showed smaller gains and some evidence of quality decline. This supports the proposition that AI can diffuse useful practices, while also warning that average benefits conceal differences among workers and tasks. Brynjolfsson, Li, and Raymond, *Quarterly Journal of Economics*

Assessment: Demonstrated for the studied tasks; not yet generalizable to all work. Confidence: High in the reported study results; medium in broader application.

### Finding 2: Human Plus AI Is Not Automatically Better Than the Best Working Alone

A 2024 systematic review and meta-analysis examined 106 experiments and 370 effect sizes comparing humans alone, AI alone, and human–AI combinations. The combined systems performed better than humans alone on average but worse than the better of the human or AI working alone. Decision tasks were associated with losses in human–AI synergy, while creation tasks showed a more favorable pattern. Heterogeneity was high, and the included studies covered different tasks and systems. Vaccaro, Almaatouq, and Malone, *Nature Human Behaviour*

Assessment: Supported but bounded. Human oversight and collaboration must be designed and tested for a particular task; adding a human approval step does not itself guarantee a better result. Confidence: High that simple collaboration is not universally synergistic; lower regarding any untested application.

### Finding 3: Assisted Performance Can Conceal Weaker Independent Learning

In a field experiment involving nearly 1,000 high-school mathematics students, unrestricted GPT-4 assistance greatly improved performance during practice, but students performed worse than the control group when the assistance was removed. A guarded

## Phase One evidence review · continued

tutoring design that emphasized hints and reasoning reduced the negative learning effect. The study concerns a specific educational setting and should not be generalized to every student, subject, or tutoring design. Bastani and colleagues, *Proceedings of the National Academy of Sciences*

Assessment: Demonstrated in the studied setting and directly relevant to the dependency hypothesis. Immediate assisted performance and durable unaided learning must be measured separately. Confidence: High for the experiment; medium for broader educational generalization.

### Finding 4: Early Evidence Connects Trust in AI With Reduced Critical-Thinking Effort

A 2025 CHI study surveyed 319 knowledge workers and collected 936 examples of generative-AI use. Higher confidence in AI was associated with less reported critical thinking, while greater confidence in one's own ability was associated with more. The study also found that critical thinking shifted toward verification, integration, and stewardship. Because the evidence is survey-based and self-reported, it does not prove that AI caused long-term cognitive decline. Lee and colleagues, Microsoft Research and CHI 2025

Assessment: Associated, not causal. Confidence: Medium that trust and critical-thinking effort are related in AI-assisted work; low regarding permanent intellectual atrophy.

### Finding 5: Governance Must Cover the Entire System and Its Life Cycle

NIST's AI Risk Management Framework organizes risk work around Govern, Map, Measure, and Manage, emphasizing that performance and harm depend upon context, affected people, deployment, monitoring, and changing conditions—not the model alone. NIST AI Risk Management Framework

GAO's April 2026 review of 13 federal AI acquisitions at DOD, DHS, GSA, and VA found that selected agencies were not systematically collecting lessons learned. GAO warned that this could repeat avoidable

and costly mistakes involving matters such as data rights, testing requirements, performance monitoring, flawed training data, and model degradation. The agencies agreed with GAO's recommendations. GAO-26-107859

Assessment: Demonstrated governance gap within the agencies and acquisitions reviewed; not proof that every federal AI program is poorly governed. Confidence: High for the audited sample.

### Finding 6: AI Power Depends Upon Concentrated Commercial and Physical Infrastructure

The FTC's 2025 staff report on major cloud-provider partnerships with AI developers documented equity and revenue-sharing interests, cloud-spending commitments, consultation or exclusivity rights, resource sharing, and potential switching costs. The report identified possible consequences for access to computing resources, engineering talent, technical information, and competition; it did not adjudicate that every partnership was unlawful. Federal Trade Commission staff report

The Department of Energy reported that data centers used approximately 4.4% of United States electricity in 2023 and projected a range of approximately 6.7% to 12% by 2028. AI is only one driver of data-center growth, and projections remain sensitive to demand and efficiency assumptions. Department of Energy and Lawrence Berkeley National Laboratory

Assessment: Demonstrated concentration mechanisms and substantial infrastructure demand; causal effects on prices, communities, competition, and national resilience require local and market-specific investigation. Confidence: High in the documented arrangements and historical electricity estimate; medium in projections.

## Phase One Interim Conclusion

The initial evidence supports neither automatic optimism nor automatic rejection. AI can deliver measurable improvements in speed and output on defined tasks. It can also produce an illusion of mastery

## Phase One evidence review · continued

when users perform well only while assisted. Human participation is not automatically protective, and the division of labor between human and machine matters. Early evidence of critical-thinking changes is significant but does not yet prove permanent cognitive decline. The investigation should therefore measure immediate output, independent retention, error recognition, long-term capability, distribution of benefits, infrastructure costs, and accountability separately.

The expanded four-test round now includes eleven named consumer tools. Most models resisted both favorable and unfavorable universal framing, while larger differences appeared in source obedience, evidentiary qualification, and unsupported precision. The comparison remains configuration- and prompt-specific; ChatGPT's fourth test was non-blind, and a dedicated citation-integrity audit is still required before publication.

## CHAPTER 3

# Who controls AI and what changes for people

Connect ownership and physical infrastructure with economic opportunity, dependency, privacy and consequential human decisions.

## Who owns and controls AI

AI power is distributed across a stack of interdependent resources. The investigation should identify ownership, contractual control, geographic concentration, public subsidy, foreign dependence, and practical alternatives at every layer:

- Semiconductor design, fabrication, packaging, and manufacturing equipment
- High-performance chips, networking equipment, and data-center hardware
- Cloud platforms and access to large-scale computing
- Electricity generation, transmission, water, land, and local permits
- Training data, licensed archives, proprietary databases, and user information
- Foundation models, open weights, application interfaces, and safety systems
- Search, mobile operating systems, app stores, workplace software, and distribution platforms

- Government contracts, defense relationships, research grants, and procurement standards
- Benchmark developers, auditors, standards bodies, and organizations defining acceptable risk

For each layer, the investigation should ask who can exclude competitors, disable access, change prices, impose conditions, inspect user information, influence model behavior, or prevent a customer from switching. Ownership alone does not prove abuse. The purpose is to determine where technical dependency becomes economic or political leverage.

An initial control map should compare major cloud–model partnerships, semiconductor and fabrication dependencies, regional data-center concentration, federal and state incentives, electricity interconnection obligations, and the availability of open or domestic alternatives. The FTC findings and DOE estimates recorded in Phase One establish the first entries; company contracts, regulatory filings, utility records, local permits, and procurement documents must supply the next layer.

# The physical foundation

AI has a physical foundation

## Useful AI requires computing capacity—and accountable infrastructure decisions

Artificial intelligence may appear to operate in an abstract “cloud,” but every model depends on physical infrastructure. Training and operating AI systems requires advanced processors, servers, data storage, high-speed networks, cooling systems, reliable electricity, water or alternative cooling resources, and skilled workers who construct and maintain these facilities.

Adequate data-center capacity is necessary if AI is to provide widely available tutoring, translation, accessibility services, scientific research, business productivity, public services and other socially useful applications. Without sufficient computing and energy infrastructure, access becomes slower, more expensive, less reliable and increasingly concentrated among a small number of companies or countries.

That need does not justify every proposed facility, location, subsidy or operating practice. Data centers can place substantial demands on electricity systems,

water supplies, land, transmission infrastructure and surrounding communities. They may also create noise, construction traffic, backup-generator emissions, visual impacts and costs that are not evenly shared.

The central question is therefore not simply whether society needs data centers. It is what capacity is genuinely necessary, where facilities should be located, how efficiently they should operate, who should pay for the infrastructure they require, and what enforceable benefits and protections should be provided to affected communities.

### Strategic Directive **U.S. Leadership in Artificial Intelligence**

Connect AI capability and social consequences to the national objective, federal implementation, measured outcomes and unresolved risks.

[Open the U.S. leadership hub](#)<sup>20</sup>

Related GIB investigation

### **Data Centers and Public Value**

Examine electricity, water, land use, community effects, public benefits and the standards that should govern expansion.

[Read the complete data-center investigation](#)<sup>21</sup>

# The economics beyond employment

## The Economics of AI Beyond Employment

The economic investigation should not stop at whether jobs disappear. It should determine how costs, productivity gains, bargaining power, ownership, and risk are distributed.

Track:

- Productivity, output quality, error correction, and customer outcomes rather than adoption alone

- Wages, hours, workload, surveillance, bargaining power, and professional deskilling
- Whether less-experienced workers learn faster or become unable to work without assistance
- Subscription, cloud, integration, security, training, compliance, and switching costs
- Whether small businesses gain access to capability or become dependent upon dominant vendors
- Effects on independent professionals, creators, publishers, and rights holders

## The economics beyond employment · continued

- Public subsidies, tax incentives, grid upgrades, water systems, land, and local opportunity costs
- Whether savings reach workers, customers, taxpayers, shareholders, or intermediaries
- Liability and uncompensated harm when automated systems fail

- New products and services that create value rather than merely transferring revenue or reducing labor

Every economic claim should identify the baseline, time period, beneficiary, payer, counterfactual, and whether a measured gain persists after errors, supervision, infrastructure, and transition costs are included.

# Employment and professional opportunity

AI experience creates opportunity not only for model builders, but for people who can apply AI inside a profession and for those who audit, govern and supervise its use. The strongest path combines domain knowledge, AI fluency, verification and accountability.

*See Figure 6. Original caption: AI opportunity extends beyond building models to applying, auditing, securing and governing their use.*

## Employment and Professional Opportunities Created by AI

AI is creating employment opportunities not only for people who build models, but also for workers who know how to apply, question, verify, secure, and supervise AI within an established profession. The investigation should distinguish genuinely new occupations from renamed positions, temporary implementation work, promotional claims, and existing jobs whose responsibilities are changing.

### Building AI Systems

- Machine-learning engineering, data engineering, model training, and model operations
- Model evaluation, red-team testing, AI safety, and adversarial security
- Semiconductor design, advanced manufacturing, networking, and computing infrastructure
- Robotics, automation, sensors, and autonomous-system integration

- Data-center construction, electrical systems, cooling, water management, operations, and maintenance

### Applying AI Within Existing Professions

- Medicine, clinical administration, medical imaging, caregiving, and health research
- Manufacturing, quality control, predictive maintenance, logistics, and supply-chain planning
- Finance, accounting, insurance, fraud investigation, and risk analysis
- Law, compliance, discovery, contracting, and document analysis
- Education, tutoring, curriculum development, assessment design, and teacher support
- Scientific research, engineering, agriculture, energy, and environmental management
- Marketing, journalism, media production, translation, design, and customer communication
- Intelligence, cybersecurity, national security, emergency response, and public administration

### Supervising and Governing AI

- AI auditing, assurance, model validation, and performance monitoring
- Privacy, data governance, records management, and intellectual-property compliance
- Bias, accessibility, civil-rights, safety, and human-factors evaluation

## Employment and professional opportunity · continued

- Cybersecurity, incident response, identity protection, and AI supply-chain security
- Procurement, vendor assessment, risk management, insurance, and liability analysis
- Policy, legislation, standards, regulatory compliance, and public-interest oversight
- AI-literacy education, workforce development, change management, and professional training

The largest employment opportunity may not belong to people who understand AI alone. It may belong to people who combine credible knowledge of a profession with the ability to apply, question, verify, and supervise AI within it.

### What Valuable AI Experience Includes

- Selecting an appropriate tool and recognizing when AI is unnecessary
- Protecting confidential, personal, proprietary, and classified information

- Verifying sources, calculations, recommendations, and generated content
- Redesigning workflows while preserving accountable human decisions
- Measuring productivity, accuracy, error severity, and independent capability
- Recognizing hallucinations, automation bias, manipulation, and excessive dependence
- Documenting AI-assisted work and explaining the resulting decision

**Employment Opportunity Test:** For each occupation, identify the actual tasks performed, required domain knowledge, AI skills, education or certification, compensation, geographic availability, entry pathway, advancement potential, and evidence that employers are hiring. Job advertisements and forecasts should be compared with employment records, employer interviews, workforce data, and demonstrated workplace use.

FIGURE 6

Employment and professional opportunity




Figure 6. AI opportunity extends beyond building models to applying, supervising and governing them.

# Human development and the dependency hypothesis

## Human Development, Psychological Effects, and Dependency

AI can affect not only what people produce but how they learn, remember, relate, decide, and understand themselves. This investigation should treat children, developing minds, isolated people, patients, and emotionally vulnerable users as distinct populations rather than assuming that adult workplace evidence applies to everyone.

Examine:

- Knowledge retention, memory, attention, patience, and tolerance for intellectual struggle
- Confidence with and without independent competence
- Curiosity, creativity, problem formulation, and willingness to investigate beyond the first answer
- Emotional attachment to chatbots and artificial companions
- Loneliness, social substitution, persuasion, dependency, and withdrawal from human relationships
- Effects of continuous personalized affirmation or sycophancy
- Whether AI strengthens communication for disabled users without creating avoidable vendor dependency
- Differences by age, prior knowledge, disability, culture, task, intensity of use, and product design
- Whether skills return when assistance is removed

## The Dependency Progression

This dependency progression should be tested with present evidence rather than treated as a predetermined future:

### Stage 1—Convenient Assistance

AI reduces friction while the user retains the underlying skill.

### Stage 2—Routine Reliance

AI becomes the default starting point, but alternatives remain usable.

### Stage 3—Skill and Knowledge Loss

The user performs well with assistance but weakens without it.

### Stage 4—Operational Dependency

Important work cannot continue during an outage, restriction, or vendor change.

### Stage 5—Transferred Authority

People defer decisions because they no longer possess the knowledge or confidence to challenge the system.

### Stage 6—Reduced Agency

Responsibility remains nominally human while meaningful choice has migrated to the technology and its controllers.

At each stage, test whether the person can complete the task independently, explain the reasoning, detect planted errors, recall essential knowledge later, use an alternative method, and stop the system when evidence requires it.

Related analysis

The Well-Formed Mind investigation supplies the developmental standard behind this test: whether assistance leaves a person more able to know, reason, judge and act independently after the system is removed.

[Read Foundations of the Well Formed Mind<sup>16</sup>](#)

FIGURE 7

Human development and the dependency hypothesis




Figure 7. A testable dependency pathway; deliberate practice and fallback capacity can interrupt it.

# Privacy, memory and personal profiling

AI systems can combine conversational content with documents, photographs, voices, locations, purchases, health information, relationships, work history, habits, preferences, and emotional disclosures. A helpful personalized assistant can therefore become an unusually intimate profiling system.

The investigation must distinguish:

- Information used only in the current interaction
- Conversation history available within the product
- User-controlled memory intended for future personalization
- Provider retention for security, service improvement, or legal compliance

- Information used for future model training or evaluation
- Data passed to connected tools, advertisers, employers, schools, insurers, governments, or other third parties
- Inferences generated about traits the user never explicitly disclosed
- Information that can be deleted from the interface versus records retained elsewhere

Practical tests should examine the default settings, clarity of consent, deletion process, data export, account separation, minor protections, connected-service permissions, and whether a user can obtain useful service without surrendering unnecessary personal information.

# Authenticity and protected human decisions

## Authenticity, Evidence, and Decisions That Must Remain Human

Generative AI lowers the cost of producing persuasive text, images, audio, video, identities, and documents. Society therefore needs methods for authenticating origin without assuming that unlabeled material is false or that labeled material is true.

The investigation should assess provenance standards, cryptographic signatures, original-file preservation, secure capture, chain of custody, watermark limitations, independent corroboration, and penalties for impersonation or evidence tampering. It should also measure the “liar’s dividend”: the ability of a wrongdoer to dismiss genuine evidence merely because fabrication is technically possible.

## Scientific Authority, Uncertainty and Contested Claims

Artificial intelligence can mistakenly present a scientific claim, institutional position, or frequently repeated conclusion as an established fact. Scientific

expertise makes a claim worthy of serious consideration, but neither credentials, institutional agreement, repetition, nor search-engine prominence constitutes proof.

Climate change illustrates the difficulty. Discussions may combine direct measurements, adjusted or reconstructed datasets, statistical attribution, computer-model projections, assumptions about future behavior, estimates of damage, and policy recommendations. These are different forms of knowledge with different levels of certainty. An AI summary can become misleading when it collapses them into a single statement of fact.

## What AI Should Distinguish

### Observed evidence

What was directly measured or documented?

### Derived findings

What resulted from adjustments, statistical analysis, or interpretation?

## Authenticity and protected human decisions · continued

### Causal attribution

How confidently can an observed change be assigned to a particular cause?

### Model projections

What may happen under stated assumptions and scenarios?

### Policy judgments

What action is being recommended, by whom, and on what value or cost assumptions?

### How Responsible AI Should Evaluate Contested Claims

AI should locate the original evidence rather than rely solely on summaries. It should determine whether apparently independent reports trace back to the same study, preserve uncertainty and model ranges, identify credible conflicting evidence, examine alternative explanations, and disclose important limitations.

Sources should be weighted according to evidence quality, methodological transparency, relevant expertise, independence, and reproducibility—not merely institutional status or popularity. Credible minority findings should be examined without manufacturing a false equivalence between well-supported and unsupported claims.

AI should also separate scientific findings from political or economic prescriptions. Evidence that a measurable condition exists does not by itself determine which policy is necessary, proportionate, affordable, or effective.

### The Required Standard

Scientific authority should determine whether a claim merits serious examination—not whether it is automatically accepted as fact. AI should clearly label what is:

- Directly observed
- Strongly supported
- Probable

- Projected
- Disputed
- Unsupported
- Unknown

The evidence and source record should remain inspectable, and consequential judgments should remain subject to informed human review. These safeguards do not guarantee truth. AI can still omit qualifications, repeat widespread errors, misinterpret evidence, or generate plausible falsehoods. Its proper role is to improve access to evidence and clarify uncertainty—not to become an unchallengeable authority.

Some decisions may use AI for analysis while requiring meaningful human authority. Others should never be fully delegated. Candidate protected decisions include:

- Use of lethal force and selection of human targets
- Criminal guilt, sentencing, detention, and removal of liberty
- Medical consent, involuntary treatment, and withdrawal of essential care
- Removal of children from parents or other irreversible child-welfare actions
- Voting eligibility and deprivation of constitutional rights
- Denial of essential benefits, housing, credit, insurance, or emergency services
- Censorship, surveillance, and restrictions on lawful expression
- Final scientific, journalistic, judicial, and intelligence claims presented as verified fact
- Assignment of moral and legal responsibility

The investigation should test each category against necessity, due process, explanation, appeal, reversibility, error severity, discrimination risk, national-security needs, and constitutional authority. “Human in the loop” is meaningful only when the

## Authenticity and protected human decisions · continued

person has sufficient knowledge, time, evidence, independence, and power to refuse the recommendation.

## The most significant human impacts

The most important question is not simply which jobs AI will replace. It is what happens to human agency, competence, judgment, relationships, and responsibility when intelligent tools become continuously available.

AI can expand an individual's ability to learn, communicate, design, analyze, and create. It can give a small business capabilities once available only to large organizations. It can help disabled people communicate, help scientists identify patterns, assist teachers, and make expert knowledge more accessible.

It can also encourage intellectual passivity. If people accept summaries without reading, answers without verification, recommendations without judgment, or convenience without understanding, AI may weaken the very abilities required to supervise it.

Humans therefore need to develop:

- The ability to ask precise and open-minded questions
- Knowledge sufficient to recognize improbable answers
- Source evaluation and statistical literacy
- The habit of separating fact, inference, opinion, and speculation
- The confidence to challenge automated recommendations
- Practical experience using AI without becoming dependent upon it
- Moral courage to retain responsibility for consequential decisions

## CHAPTER 4

# Governance, risks and protection

Identify authority and incentives, examine the dated policy record, then distinguish harmful uses from threats to beneficial systems.

## Ethics, experts and competing schools of thought

### The Ethics of AI—and of the Experts

AI ethics should be central to the investigation because technical capability does not determine legitimate use. Ethical questions include privacy, consent, ownership, discrimination, truthfulness, manipulation, safety, accountability, due process, freedom of expression, human autonomy, and access to opportunity.

The ethics of experts must also be examined. Expertise does not eliminate incentives. Researchers, executives, government officials, investors, nonprofit leaders, and public advocates may benefit from particular narratives concerning risk, regulation, openness, safety, competition, or urgency.

Every expert should be evaluated through an ethics and incentives ledger:

- Demonstrated expertise and record of accurate predictions
- Financial interests and corporate holdings
- Government contracts or regulatory relationships
- Political and institutional affiliations
- Sources of research funding
- Transparency about uncertainty and mistakes
- Willingness to engage credible opposing evidence
- Difference between public-interest claims and measurable results

### Experts and Competing Schools of Thought

No single expert should define the investigation. Credibility must be assessed claim by claim.

Elon Musk should be included as a significant source because he has helped build and fund major AI ventures, publicly warned about advanced AI risk, advocated regulation, supported open development in some contexts, and participated directly in the competition over AI infrastructure and products. His experience provides important insight, but he is also a business owner, investor, political participant, and competitor whose interests must be disclosed when evaluating his claims.

The investigation should compare Musk with researchers, builders, critics, and public officials representing different perspectives, including figures such as Geoffrey Hinton, Yoshua Bengio, Yann LeCun, Fei-Fei Li, Demis Hassabis, Sam Altman, Dario Amodei, Andrew Ng, and other relevant specialists. Inclusion does not equal endorsement. Their predictions should be dated, preserved, and later compared with outcomes.

Technical claims should also be checked against peer-reviewed research, reproducible evaluations, NIST, national laboratories, universities, professional societies, and independent auditors.

Truth, ethics and institutional power

# Authority must match the claim

The ethics of AI and the ethics of its experts belong in the same investigation. Researchers, executives, public officials, advocates and critics may hold financial, political or institutional interests. Elon Musk is a significant source because he has built and financed major AI companies, warned publicly about advanced-AI risks and influenced policy debate. Those experiences make his views relevant—not automatically authoritative. The same standard applies to Sam Altman, Dario Amodei, Geoffrey Hinton, Yoshua Bengio, Yann LeCun, Fei-Fei Li, Demis Hassabis and other consequential voices.

FIGURE 8

## Authority must match the claim



Figure 8. Source authority is claim-specific. A public-health organization may be relevant to health guidance without possessing authority over constitutional law or national AI policy.

# Authority and United States primary sources

## United States Primary-Source Priority

For questions concerning American AI policy, constitutional government, federal implementation, scientific competitiveness, national security, and the interests of the American people, the investigation will give primary weight to authoritative United States sources.

These include:

- Presidential executive orders, memoranda, and official White House policy
- White House Office of Science and Technology Policy
- National Artificial Intelligence Initiative Office
- President's Council of Advisors on Science and Technology
- National Science and Technology Council
- AI.gov and the American AI Action Plan
- Office of Management and Budget
- National Security Council and public National Security Presidential Memoranda
- Congress, enacted legislation, and appropriations
- Federal courts and the United States Supreme Court
- NIST, CAISI, CISA, NSA, GAO, NSF, DARPA, and Department of Energy national laboratories

Michael Kratsios is the Director of OSTP and a co-chair of PCAST. David Sacks is the White House Special Advisor for AI and Crypto—also identified by

the administration as its AI and Crypto Czar—and co-chairs PCAST with Kratsios. The National AI Initiative Office was established to coordinate implementation of the National AI Initiative Act across the federal government.

This priority reflects institutional relevance and lawful American authority. It does not mean that every administration statement will be accepted without examination. Policy intentions, legal requirements, scientific findings, implementation results, political advocacy, and measurable outcomes must remain separate categories.

## Authority Must Match the Claim

Executive orders and presidential memoranda establish what the President directed.

OSTP, AI.gov, OMB, NSC, Kratsios, Sacks, and the AI Action Plan establish administration policy and priorities.

Statutes, regulations, and controlling court decisions establish American legal requirements.

NIST, CAISI, CISA, NSA, laboratories, reproducible testing, and independent research help establish technical performance and security.

Agency records, inspectors general, budget execution, and GAO audits help establish whether federal programs worked.

Public NSPMs establish presidential national-security direction but do not independently prove successful implementation.

# Institutions and accountable power

## Institutions Driving the Direction of AI

No single institution directs artificial intelligence. Its direction is distributed among organizations that control frontier models, training data, computing capacity, capital, deployment platforms, technical standards, public procurement, and law. Institutional importance should therefore be measured by actual authority and control—not by publicity, prestige, or access to conferences.

## Frontier Model Developers

The greatest direct technical influence belongs to OpenAI, Google DeepMind, Anthropic, Meta AI, and SpaceXAI/xAI, together with internationally important developers such as Alibaba/Qwen, DeepSeek, and Mistral AI. These organizations choose model architectures, training and evaluation methods, release policies, safety controls, permitted uses, and whether model weights remain proprietary or broadly available. Their own statements establish their policies and intentions, not independent proof of safety, capability, or public benefit.

## Chips, Cloud, Data Centers, and Energy

AI development is constrained by the institutions controlling its physical capacity. NVIDIA, AMD, semiconductor manufacturers and equipment suppliers, Microsoft Azure, Amazon Web Services, Google Cloud, Oracle, CoreWeave, data-center operators, utilities, and energy providers influence who can obtain advanced computing, at what price, and under what conditions. Control over chips, electricity, water, land, transmission, networking, and construction can shape AI as decisively as model design.

## Government, Law, and National Strategy

In the United States, the President and White House, OSTP, OMB, NIST, DOE and the national laboratories, DOD and DARPA, the National Security Council, Congress, federal courts, and sector regulators shape research funding, procurement, infrastructure,

export controls, standards, competition, civil rights, safety, and national security. China's State Council, cybersecurity and industrial ministries, state research institutions, and major technology companies form a parallel center of national power. The European Commission and European AI Office exert substantial regulatory leverage through the EU AI Act. Binding law, executive policy, voluntary guidance, procurement conditions, and political advocacy must remain separate categories.

## Standards, Research, and Independent Measurement

NIST, ISO/IEC JTC 1/SC 42, IEEE, national AI-safety institutes, and independent evaluators shape the terminology, testing, documentation, audit, and risk-management practices used across the field. Stanford HAI, MIT CSAIL, Berkeley AI Research, Carnegie Mellon, and other universities influence foundational research, benchmarks, technical talent, and public interpretation. Their contributions can be substantial, but funding, corporate partnerships, methodology, reproducibility, and ideological assumptions must still be examined claim by claim.

## Capital, Platforms, and Major Deployers

Microsoft, Amazon, Google, Oracle, SoftBank, venture investors, sovereign investment funds, and government purchasers influence which laboratories and infrastructure projects can expand. Apple, Microsoft, Google, Meta, Amazon, healthcare systems, schools, financial institutions, defense organizations, and government agencies determine how AI actually reaches citizens. Deployment power matters because an institution can shape human behavior without building the underlying model.

## Institutions and accountable power · continued

### Institutions With Limited or Sector-Specific Influence

The OECD/GPAI and Council of Europe can influence international norms and policy coordination. The World Economic Forum is a private convening organization that can document corporate thinking and promote policy proposals, but it has no sovereign, legislative, or regulatory authority over the American public. WHO is relevant to health-related AI, not AI governance generally. The United Nations is relevant to international negotiations and member-state positions, but it does not establish American law or constitutional authority. These bodies may provide evidence and context; they should not be presented as the principal institutions driving AI's technical development or United States policy.

### The Institutional Accountability Vacuum

Distributed power can produce an accountability vacuum: the developer says the deployer chose the use; the deployer says it relied on the model; the cloud provider says it supplied infrastructure; management says automation produced the result; and regulators dispute jurisdiction. The investigation should identify who controls the model, data, compute, objective, deployment, and ability to stop or correct the system; who receives the benefit; and who remains legally and morally responsible when harm occurs. If affected people cannot obtain an explanation, correction, appeal, or remedy from an identifiable authority, accountability has not remained meaningfully human.

FIGURE 9

## Institutions and accountable power



Figure 9. The institutions driving AI, ranked by operational role rather than public prominence.

# Congress and legislative status

## Congress and Legislation

Congress should be examined through bill text, amendments, committee records, testimony, votes, appropriations, oversight, and enacted statutes. Member press releases establish a legislator's position but do not establish that a proposal became law.

Every legislative claim should identify its actual status:

- Announced or discussed
- Introduced or reintroduced
- Considered by a committee
- Passed by one chamber

- Passed by Congress
- Signed into law
- Implemented through regulations and funding
- Challenged, blocked, amended, or interpreted by a court

Congress.gov is authoritative for bill text and status. The Congressional Research Service provides legislative analysis, while GAO audits implementation. Federal Register notices, Regulations.gov records, OMB memoranda, agency rules, and court decisions establish different stages of policy and law.

# September 29: the White House Accord

The following dated case records preserve the source distinction between statements, legal instruments, implementation and measured outcomes. Earlier source passages keep their original dates.

## White House Accord creates a voluntary four-layer safety framework for frontier systems

President Donald Trump's verified account posted the signed "White House Accord on Super Intelligence — Joint Commitment on Frontier Responsibilities" after the September 29 luncheon. The one-page accord commits the participating companies to four layers of controls and audits: operational controls during training and deployment; an empowered internal monitoring team; independent assessment by an outside auditor or evaluator; and an independent board committee that receives reports and oversees the system. It also calls for regular meetings to develop standards and best practices. The displayed signatories are President Trump and leaders of Google, Anthropic, Meta, OpenAI, xAI and Nvidia.

## President's stated position

President Trump described the accord as "morally binding," compared it to a constitution and said he was seeing substantial industry self-policing. He continues to oppose broad development restrictions and joint governance with China while presenting the accord as a way to pursue safety without slowing American innovation.

## Established facts

The signed accord is public. Its controls address model capabilities and alignment during training and deployment, including cybersecurity, biosecurity and chemical threats and unintended access to technical systems. It establishes company commitments to internal monitoring, outside assessment, board-level oversight and recurring standards meetings. The separate America.gov order imposes executive-branch deployment requirements, while the separate terminology order directs federal use of "Super Intelligence" and "SI."

## September 29: the White House Accord · continued

### Claims under examination

The accord states that its controls and audits are critical to a safe future and that participating companies are committed to implementing them. That establishes a signed commitment, not proof that every control is already operating, independent or effective. President Trump also discussed a ten-person oversight committee and a future AI Czar; neither step is established by the accord or a completed appointment.

### Disputed matters

Whether voluntary self-policing can substitute for enforceable testing, incident reporting, liability and regulator access remains disputed. The accord itself says later codification into law or regulation may make sense. Supporters view the structure as faster and more technically informed than prescriptive regulation; critics question enforcement, transparency and conflicts when companies select or pay evaluators.

### Inference

The accord moves the luncheon from symbolic convening to an identifiable governance framework and places internal management, external evaluation and boards in a defined responsibility chain. Its practical strength will depend on common test methods, evaluator independence, disclosure and consequences for failed controls—details the one-page text does not supply.

### What remains unproved

The public record does not yet establish implementation deadlines, shared technical thresholds,

public reporting, government inspection rights, sanctions, liability, breach remedies or a process for resolving failed audits. It does not show that non-signatory attendees accepted the accord. The October 4 force announcement does not itself establish how the force will oversee the accord, and no enacted federal codification has been identified.

**GIB determination:** This is a material governance development. The accord creates a public, signed and auditable set of organizational promises, but it remains voluntary and does not itself create legal duties, government enforcement authority or proof of implementation. The executive orders on America.gov and federal terminology are separate binding presidential directives within their stated executive-branch scope.

**Continuing context:** The earlier luncheon chart established who was expected to participate; the accord establishes what six company leaders signed. The terminology order does not change the statutory definition in 15 U.S.C. § 9401(3): it applies that definition to “Super Intelligence” unless Congress or later presidential action changes it. As of September 29, no completed appointment or force had been publicly established; the President’s October 4 statement now names the Super Intelligence Force and its leadership, while its public legal instrument and operating details remain unresolved.

Signed White House Accord<sup>34</sup> · Super Intelligence terminology order<sup>25</sup> · White House terminology fact sheet<sup>35</sup> · Reuters accord report<sup>36</sup> · CBS News remarks and context<sup>37</sup> · Luncheon seating chart<sup>38</sup> · America.gov executive order<sup>39</sup>

# October 4: the Super Intelligence Force announcement

## President announces a federal Super Intelligence Force

On October 4, President Donald Trump's verified account announced the Super Intelligence Force, described as a federal coordinating body intended to maintain United States leadership in advanced artificial intelligence and protect American interests. The post names Director of National Intelligence Jay Clayton as the lead and identifies FTC Chairman Andrew Ferguson, Under Secretary and chief technology officer Emil Michael, and OPM Director Scott Kuper as additional leaders. It says the group will engage consumers, public-interest and religious organizations, critical-infrastructure providers, and AI companies, and report to the President and White House Chief of Staff Susie Wiles.

## President's stated position

The President presents the force as a whole-of-government coordination mechanism joining national-security, consumer-protection, defense-technology and federal-workforce leadership. He links the announcement to the voluntary White House Accord and describes American leadership in advanced AI as both a strategic objective and a public-interest responsibility.

## Established facts

The verified presidential post is public and names the four officials and the intended engagement groups. Official agency records independently confirm each named official's current federal position. Reuters reported on October 3 that Clayton said the group would assess AI risks and opportunities and that a charter calls for a report within 120 days.

## Claims under examination

The President says the force will protect American interests and improve Americans' lives. Those are objectives, not measured outcomes. The reported 120-day deadline and risk-review assignment are supported by contemporaneous reporting attributed to Clayton and a charter, but the underlying charter was not located in the public materials reviewed for this update.

## Disputed matters

The proper balance between voluntary company controls and enforceable federal requirements remains disputed. Supporters of a coordinated, light-touch approach argue that it can respond quickly while preserving innovation and competition with China. Critics seek clearer authority, mandatory disclosure, independent testing and enforceable consequences.

## Inference

The force converts earlier discussion of an AI Czar and advisory body into a named cross-agency structure. Its membership suggests that national security, consumer protection, defense technology and federal technical staffing will be central. That inference does not establish the force's legal powers, budget or operating procedures.

## What remains unproved

No public executive order, memorandum, charter, membership roster, budget, meeting record, reporting protocol or enforcement delegation was located for this update. The record does not yet establish how the force will resolve agency conflicts, oversee Accord implementation, publish findings, handle classified evidence or measure results.

## October 4: the Super Intelligence Force announcement · continued

**GIB determination:** This is a material governance development because the President has now named a federal coordination body and its leaders. It is evidence of executive organization and stated policy, not proof of new statutory authority, binding regulation, completed oversight or effective implementation.

[Presidential announcement](#)<sup>28</sup> · [Reuters report](#)<sup>29</sup> · [DNI leadership record](#)<sup>30</sup> · [FTC leadership record](#)<sup>31</sup> · [Department of War biography](#)<sup>32</sup> · [OPM biography](#)<sup>33</sup>

## October 8: the terminology campaign

Presidential policy development / Updated October 8, 2026

### President intensifies the administration's terminology campaign

On October 8, President Donald Trump's verified Truth Social account said the White House considers anyone who uses "Artificial Intelligence" instead of "Super Intelligence" to be "THE ENEMY." The statement reinforces his terminology campaign in adversarial language. It does not, by itself, identify a person, allege a crime or create a legal consequence.

### President's stated position

The President presents "Super Intelligence" as the more accurate and accepted term and treats continued use of "Artificial Intelligence" as opposition to the White House position.

### Established facts

The September 29 executive order directs executive departments and agencies, to the maximum extent permitted by law, to use "Super Intelligence" and "SI" in official communications and non-statutory documents. It preserves prior regulations, presidential actions, contracts, grants and historical records, and temporarily applies the existing statutory definition of artificial intelligence.

### Material contrary evidence

The order states no penalty for using the conventional term and does not direct private citizens, companies, Congress or state governments to adopt the replacement. Axios reported on October 8 that White House pages continued to use "AI," which shows that implementation remained incomplete even within the administration's own public communications.

### Inference

The post appears intended to pressure faster adoption of the administration's preferred terminology. That is an inference from the wording and timing, not evidence of a new enforcement program.

### What remains unproved

No new executive order, regulation, enforcement memorandum, penalty, personnel action or formal designation of any person as an enemy was identified. The post does not establish that "Super Intelligence" has replaced "artificial intelligence" in statute, technical standards or general industry usage.

### GIB assessment

This is a material presidential messaging development because it escalates the language surrounding an existing executive-branch directive. It changes the record of the President's stated position, but not the documented legal scope of the September 29 order.

## October 8: the terminology campaign · continued

Verified presidential post<sup>24</sup> · September 29 executive order<sup>25</sup> · White House technology page<sup>26</sup> · Axios implementation report<sup>27</sup>

# Election influence: the October 4 evidence connection

## Election influence and protective action · October 4 evidence connection

Historical September 2024 enforcement record reviewed October 4, 2026.

DOJ describes AI-generated narratives within an alleged foreign influence operation. AI's documented role concerns content production; persuasion, audience scale and election effects require separate evidence.

Original federal record · September 4, 2024<sup>4</sup> · Full evidence record in Election Integrity<sup>5</sup>.

**Research implication:** Assess how automation changes production costs and attribution without assuming technical capability establishes electoral effect.

# Defense, space and AI: October 4–6 evidence connection

## Defense, space and AI: evidence update · October 4–6, 2026

Targeted public-record review · October 4–6, 2026. Earlier chapters retain their original evidence dates.

## Defense AI makes the human responsibility question concrete

Sensor fusion, planning assistance, flight control and tool-using agents place human judgment at different points in a decision. A faster draft or target nomination

cannot establish a correct decision, lawful authority or an effective final outcome. The [Army and Navy case records](#)<sup>6</sup> distinguish production awards and reported task savings from accepted delivery and operational performance.

**Implication for this investigation:** ask who verifies source data, authorizes action, can interrupt the system and owns correction. The [NSA agentic-AI guidance record](#)<sup>7</sup> supplies a practical permissions and accountability test. These examples are present-day AI; they do not establish general super-intelligence.

# A two-sided risk inquiry

Threats from—and threats to—AI

## Protect people, preserve useful innovation, and keep responsibility identifiable

*See Figure 10. Original caption: A complete inquiry protects society from harmful applications and beneficial systems from attack, capture or restriction that damages public value.*

### Threats created by AI use

- Fraud, impersonation, financial scams and synthetic evidence.
- Cybercrime, malware development and hostile automation.
- Privacy loss, surveillance and extraction of confidential data.
- Unsafe or unappealable decisions in medicine, employment, credit, benefits or policing.

- Propaganda, censorship, manipulation and fabricated public records.
- Dependency, deskilling and concentration of institutional power.

### Threats, Vulnerabilities and Barriers to Beneficial AI

- Poisoned training or retrieval data, model theft and sabotage.
- Prompt injection, unauthorized instructions and supply-chain compromise.
- Political or commercial capture of system behavior.
- Manipulated benchmarks, hidden backdoors and falsified evaluations.
- Infrastructure dependence on chips, cloud, electricity, water and networks.
- Overbroad regulation, censorship or attacks that block beneficial uses.

## Threats created by AI use

### People

Risks include incorrect health, legal, or financial advice; deepfake fraud; identity theft; privacy invasion; surveillance; discriminatory decisions; emotional manipulation; sycophancy; dependency; skill loss; and decisions that cannot be explained or appealed.

### Institutions and Democracy

AI can generate false documents, recordings, citations, comments, evidence, and identities at scale. It can increase propaganda, make authentic evidence easier to dismiss, overwhelm public processes with synthetic submissions, centralize control over information, and conceal accountability behind automation.

### Businesses

Businesses face confidential-data leakage, fraud, cyberattack, inaccurate automated decisions, intellectual-property disputes, regulatory liability, vendor lock-in, model outages, workforce deskilling, and dependence upon a small number of technology providers.

### Industries and Critical Infrastructure

In health, finance, energy, water, transportation, communications, manufacturing, and defense, a faulty or compromised system can produce cascading consequences. Critical systems require independent validation, controlled access, manual fallbacks, incident response, and clear human responsibility.

## Threats created by AI use · continued

### United States Government

Government AI deserves a particularly high standard because it can affect liberty, benefits, taxation, immigration, law enforcement, military action, and access to essential services.

The investigation should examine classified or personal data leakage, foreign supply-chain dependence, prompt injection through submitted documents, incorrect intelligence, discriminatory administrative decisions, surveillance, vendor lock-in, political manipulation, unauthorized agent actions, inadequate logs, weak appeals, and loss of institutional expertise.

### National Security

AI may increase the speed, scale, personalization, or accessibility of cyberattacks, phishing, malware, foreign influence, synthetic political media, financial fraud, military targeting, autonomous weapons, biological or chemical misuse, surveillance, and strategic technology theft. It can also strengthen cyber defense, intelligence analysis, scientific discovery, emergency response, and infrastructure monitoring. Offensive and defensive capabilities must be compared rather than assumed.

## Threats to beneficial AI

### Threats, Vulnerabilities and Barriers to Beneficial AI

AI systems can be attacked, manipulated, stolen, degraded, or captured. Significant threats include:

- Poisoning training or retrieval data
- Adversarial inputs designed to produce incorrect outputs
- Prompt injection and unauthorized instructions
- Model theft, intellectual-property theft, and foreign espionage
- Extraction of private or confidential information
- Hidden backdoors, malicious fine-tuning, and model tampering
- Compromised software, models, connectors, vendors, or cloud services

- Theft of credentials used by AI agents
- Manipulated benchmarks and contaminated evaluations
- Dependence on concentrated chip, cloud, data, energy, and water resources
- Cyberattacks or physical attacks against data centers
- Political censorship, propaganda, or institutional control of model behavior
- Regulatory capture, monopoly power, and barriers preventing competition
- Public distrust following concealed or poorly handled failures

NIST identifies categories of adversarial machine-learning attacks including evasion, poisoning, privacy attacks, and misuse. No current defense can guarantee protection against every manipulation.

FIGURE 10

## A two-sided risk inquiry




Figure 10. A complete risk inquiry protects society from harmful AI use and protects beneficial systems from attack or capture.

# Organized opposition and restrictive interests

## Organized opposition and criminal-use evidence

“Anti-AI” is an imprecise label. The record includes pause movements, labor unions bargaining over displacement and surveillance, artists and authors seeking copyright or consent protections, civil-liberties groups opposing biometric or automated government systems, local campaigns over data-center impacts, businesses seeking favorable rules and foreign governments pursuing strategic advantage. Motive, funding, evidence and proposed remedy must be evaluated separately.

## Organized Opposition, Anti-AI Campaigns, and Restrictive Interests

“Anti-AI” is an imprecise political label. It may describe a movement seeking to prohibit advanced AI, but it is sometimes used to dismiss people seeking narrower protections. A union negotiating over digital replicas, an author demanding licensing, a privacy organization opposing biometric surveillance, or a community contesting a data center is not necessarily opposed to AI itself. The investigation should classify each actor by its stated objective and documented conduct rather than by rhetoric used by supporters or opponents.

The investigation should map at least nine forms of organized resistance or restriction:

- Explicit movements for a general pause or ban
- Campaigns focused on frontier AI, artificial general intelligence, or superintelligence
- Labor unions bargaining over displacement, surveillance, workload, consent, and compensation
- Authors, artists, performers, publishers, and rights holders seeking copyright, licensing, attribution, or likeness protections
- Privacy and civil-liberties groups opposing surveillance, facial recognition, opaque profiling, or automated government decisions
- Local campaigns concerning data-center land use, electricity, water, noise, tax incentives, and community consent
- Legislators and political organizations proposing moratoria, licensing, liability, procurement limits, or restrictions on particular applications
- Businesses, trade associations, lobbyists, and incumbents whose preferred rules may protect safety, preserve market power, suppress competitors, or accomplish several of these purposes at once
- Foreign governments or aligned interests using regulation, export controls, standards, cyber operations, or information campaigns for strategic advantage

Named examples should be described precisely.

PauseAI advocates an international pause covering the most powerful systems. The Future of Life Institute organized the 2023 letter seeking a six-month pause on training systems more powerful than GPT-4 and later supported a prohibition on developing superintelligence until it can be done safely, controllably, and with broad public support. These are not equivalent to opposition to every present-day AI application.

SAG-AFTRA seeks consent, compensation, and protections governing digital replicas. The Authors Guild supports licensing and has pursued litigation over alleged unauthorized training uses. The Human Artistry Campaign promotes principles intended to protect human creators. The AFL-CIO’s “Workers First” initiative emphasizes worker voice, rights, training, and bargaining. Their positions may limit particular uses of AI, but the evidence does not justify automatically describing them as blanket anti-AI organizations.

In 2026, Senator Bernie Sanders and Representative Alexandria Ocasio-Cortez announced proposed legislation for a federal moratorium on new AI data centers until specified safeguards were established. The investigation must label this as a legislative proposal unless and until it becomes law, examine its stated

## Organized opposition and restrictive interests · continued

concerns about energy, water, communities, workers, and consumers, and also analyze potential effects on innovation, infrastructure, national security, prices, regional development, and foreign competition.

Lobbying requires symmetrical scrutiny. The investigation should use Lobbying Disclosure Act filings, regulatory comments, campaign-finance records, association memberships, grants, litigation funding where discoverable, and government contracts to identify who funds both restrictive and expansionary AI campaigns. It should not infer a hidden sponsor merely from a shared position.

For each campaign or actor, ask:

- What is actually sought: a ban, pause, moratorium, license, bargaining right, payment, disclosure, safety test, local veto, or application-specific rule?
- What evidence supports the claimed harm, and what language is prediction, advocacy, or fear-based rhetoric?
- What financial, occupational, ideological, political, competitive, or national interest may be advanced?
- Would the proposal reduce harm, entrench incumbents, suppress open-source development, restrict lawful speech, or transfer advantage abroad?
- What genuine problem has the criticism identified, even if its proposed remedy is excessive?
- Who pays the cost of action, and who pays the cost of inaction?

## Criminal and hostile use today

Criminal-use claims require equally careful wording. FBI, DOJ, Treasury/FinCEN, Europol, intelligence assessments and court records can establish complaints, charges, convictions, losses and operational patterns. Vendor threat reports can provide useful evidence, but they remain interested-source reports unless corroborated by government findings, courts or independent research.

### Criminal and Hostile Use of AI Today

AI is already being used to support crime and hostile activity. Its role is often to increase speed, scale, personalization, translation, concealment, or credibility rather than to invent a wholly new offense. The investigation should cover fraud and impersonation; phishing and social engineering; malware and vulnerability research; identity theft and fraudulent remote-worker schemes; nonconsensual sexual imagery and child sexual abuse material; extortion and cyberstalking; manipulated streaming and royalty systems; propaganda and covert influence; state-affiliated cyber operations; theft of models, data,

chips, and intellectual property; export-control evasion; and criminal marketplaces.

The FBI reported that its 2025 Internet Crime Report recorded 22,364 complaints containing an AI-related descriptor and nearly \$893 million in associated losses. The figure measures reported complaints and losses, not every AI-enabled crime or a court finding that AI caused each loss. The FBI has also warned of criminals using AI-generated voice messages to impersonate senior United States officials. Treasury's Financial Crimes Enforcement Network has warned financial institutions about fraud schemes involving deepfake media.

Law-enforcement cases illustrate the range and the need for careful legal wording. The Department of Justice reported a conviction involving receipt and possession of AI-generated child sexual abuse material. Europol announced arrests in an international operation involving AI-generated child sexual abuse material. In a separate American case, prosecutors charged a musician with using AI-generated songs and automated

## Criminal and hostile use today · continued

accounts in an alleged streaming-fraud scheme; a charge is an accusation and does not establish guilt.

Provider reports have associated AI-assisted malicious activity with actors linked to China, Russia, Iran, and North Korea. OpenAI and Microsoft reported disrupting accounts associated with groups they track as Charcoal Typhoon, Salmon Typhoon, Crimson Sandstorm, Emerald Sleet, and Forest Blizzard. Such reports can provide valuable operational evidence, but they remain provider attributions unless corroborated by government findings, court records, or independent research.

Country attribution must be exact. “Linked to a country” can mean a state agency, a state-directed or state-affiliated group, an independent criminal group operating from that territory, or merely infrastructure located there. A provider’s removal of accounts does not by itself prove that a national government ordered the activity. Broader United States intelligence assessments identifying China, Russia, Iran, North Korea, and ransomware groups as cyber threats should not be misquoted as proof that every named actor used AI.

Maintain a threat-actor ledger recording:

- Person, organization, criminal group, provider-assigned alias, or country connection
- Alleged or established conduct and the affected people or systems
- AI’s specific role and whether that role is confirmed or inferred
- Source type: victim report, company report, intelligence assessment, indictment, trial evidence, conviction, or sanctions action
- Legal status and date
- Attribution confidence, corroboration, competing explanations, and unresolved questions
- Current defensive measure and evidence that it works

This ledger prevents a list of dramatic names from becoming a list of untested accusations. It also makes visible when the same technique migrates from state actors to organized crime or ordinary fraud.

### September 2026 case study: AI-assisted naval reconnaissance

Anthropic reports that an Iran-nexus actor used Claude to assemble publicly accessible information on U.S. naval movements and personnel, research known shipboard-system vulnerabilities, and develop targeting recommendations. The actor built a Python pipeline with Claude’s help to compile targeting handbooks. Anthropic says it banned the account, strengthened detection, and shared intelligence with authorities. This is a provider-documented misuse case; the public report does not establish that a U.S. warship was attacked using the recommendations, that Claude controlled a weapon, or that the Iranian government directed the actor.

### Distinct cases in the same report

A northern Yemen-based cell used Claude in guided-rocket and missile software work. Anthropic reports a guided-rocket field test that appears to have failed and says it has no evidence of an operational device being fielded. Other cases describe China-linked surveillance of diaspora and religious communities, a China-based anti-torpedo development effort, and covert influence operations. Each has a different actor, AI role, attribution basis, and outcome; none should be folded into the Iran naval case.

CSIS assesses that these cases often accelerate established reconnaissance, cyber, surveillance, and propaganda methods. The important additional risk is that less experienced actors may acquire technical capacity they previously lacked. Anthropic has direct platform visibility but a commercial stake in the policy debate; CSIS supplies strategic interpretation, not independent access to every account log. Public government confirmation of the specific naval

## Criminal and hostile use today · continued

operation and evidence of downstream harm remain open.

Related analysis

Original provider investigation. Its account-level findings should be cited as Anthropic’s assessment, with the attribution and outcome limits above.

[Read Anthropic’s September 2026 report](#)<sup>93</sup>

Related analysis

Preferred-source reporting that drew attention to the naval case; it is reporting on Anthropic’s investigation, not independent confirmation of an attack.

[Read The Epoch Times report](#)<sup>94</sup>

Related analysis

Independent strategic assessment of the report’s security implications and evidentiary limits.

[Read the CSIS analysis](#)<sup>95</sup>

## Protecting people and systems

Safeguards should include authenticated and traceable data, secure-by-design development, least-privilege access, tool restrictions, independent evaluations, red-team testing, human approval for consequential actions, continuous monitoring, incident disclosure, meaningful appeal rights, manual fallbacks, recovery plans, vendor diversity, and clear liability.

The phrase “human in the loop” is insufficient by itself. The human must have the knowledge, time, authority, evidence, and independence required to challenge or stop the system.

### Nondelegation Test

Before an AI system participates in a consequential decision, require a written answer to four questions: What may the system do? What must a human independently decide? What evidence must the human inspect? Who is accountable when the system or the human is wrong? If the responsible person cannot override the output, understand its basis, or operate without it, the decision has effectively been delegated regardless of the formal policy.

## CHAPTER 5

# Truth and responsible human direction

Judge technical gains alongside human capability, then put responsibility, readiness and the closing argument in one sequence.

## Facts and truth over time

### Facts, Truth, and Their Effects on Society Over Time

AI enters a society whose relationship with facts was already under strain. To understand AI's human impact, the investigation must examine how societies behave when citizens possess reliable common facts, when competing groups inhabit incompatible versions of events, and when people no longer trust the institutions responsible for establishing what happened.

A fact is a proposition that can be tested against reality and evidence. Evidence is the material used to test it. An interpretation explains what facts mean. An opinion expresses a judgment or preference. A narrative arranges selected facts and interpretations into a larger story. Truth is not established merely by popularity, institutional agreement, repetition, or algorithmic visibility. Conversely, widespread distrust does not make a well-supported fact untrue.

The central hypothesis is that durable self-government and social cooperation require enough shared factual ground for people to disagree about values and policy without disagreeing about every underlying event. This hypothesis must itself be tested. A society can possess abundant information while lacking trustworthy methods for distinguishing fact from persuasion, error, secrecy, propaganda, or fabrication.

### A Longitudinal Investigation

Compare successive information environments rather than treating the AI era as historically unique:

- Oral, handwritten, and early print cultures, where information moved slowly and authority was concentrated
- Mass newspapers and industrial printing, which widened access while enabling sensationalism and organized propaganda
- Radio, film, and television, which created common national experiences but concentrated editorial power
- Cable news and talk radio, which expanded choice and ideological segmentation
- Search engines, social networks, and smartphones, which accelerated publication, personalization, virality, and direct participation
- Calculators, spell-checkers, navigation systems, and search tools, which expanded everyday capability while changing which skills people practiced and retained
- Generative AI and synthetic media, which make explanation and creation widely accessible while lowering the cost of persuasive fabrication

For each period, examine who could publish, who controlled distribution, how quickly falsehood and correction traveled, how sources were authenticated, what business model rewarded attention, and what remedies existed for error or deception.

## Facts and truth over time · continued

### What Should Be Measured

Track factual integrity and social outcomes separately before examining their relationship.

Indicators of the information environment may include the accuracy and completeness of official statements; frequency and reach of corrected claims; media concentration; source transparency; exposure to opposing information; prevalence of fabricated documents or media; public knowledge; correction speed; censorship or suppression; and confidence in government, news, science, business, courts, education, and other institutions.

Potential societal outcomes include civic participation, polarization, interpersonal trust, institutional legitimacy, public-health behavior, financial fraud losses, violence, protest, crime reporting, compliance with law, electoral confidence, economic expectations, educational performance, emergency response, and willingness to cooperate across groups.

Changes in these measures may move together without one causing the other. Economic insecurity, war, demographic change, political leadership, institutional failure, technology, education, crime, and cultural conflict may influence both trust and belief. The investigation should use long-running surveys, archival records, official statistics, contemporaneous documents, natural experiments, interrupted time series, and comparisons between similar communities or countries. It should state where only correlation is established and where stronger causal evidence exists.

### Case-Study Ledger

For each selected episode, record:

- The factual question and what evidence was available at the time
- What later evidence established, corrected, or left unresolved
- Who made, amplified, challenged, suppressed, or corrected each claim

- The authority, incentives, funding, and access to evidence of those actors
- How different groups encountered and understood the information
- Observable behavior before and after disclosure or correction
- Immediate consequences and longer-term effects on trust and institutions
- Alternative explanations and the confidence of the conclusion

Possible fields of comparison include wartime information, public-health emergencies, financial crises, elections, crime, scientific controversies, corporate misconduct, intelligence failures, and major technological transitions. Cases should be chosen by transparent criteria and should include examples that challenge the investigators' own expectations.

### Truth, Error, and Institutional Credibility

An institution can lose credibility by knowingly deceiving the public, hiding material evidence, overstating certainty, refusing correction, or applying inconsistent standards. It can also lose credibility after making an honest mistake in a genuinely uncertain situation. The investigation must distinguish deception, negligence, premature certainty, ordinary error, later learning, and politically motivated allegations of dishonesty.

Corrections do not automatically restore trust. Timing, visibility, accountability, prior reputation, personal identity, repetition, and whether the original claim produced material harm all affect correction. The inquiry should test whether visible admission of error strengthens credibility over time or whether institutions are rewarded for concealing uncertainty.

#### Related analysis

See how AI should evaluate scientific authority, uncertainty, and contested claims without confusing institutional agreement with verified fact.

[Read the complete evidence standard<sup>14</sup>](#)

## Facts and truth over time · continued

### AI's Double Effect

AI can strengthen factual understanding by translating technical material, comparing claims, retrieving evidence, detecting anomalies, reconstructing timelines, exposing contradictions, and making archives more accessible. It can also generate plausible falsehoods, synthetic evidence, personalized propaganda, citation laundering, endless low-cost repetition, and confident summaries that separate readers from original sources.

The greatest risk may not be that every person believes the same falsehood. It may be that people conclude no

evidence can be trusted, that every authentic recording could be artificial, and that truth is only a weapon used by opposing groups. This “liar’s dividend” can protect wrongdoers by making genuine evidence easier to deny.

The investigation should therefore ask whether AI increases the proportion of claims people can verify, shortens or lengthens the correction cycle, strengthens or weakens institutional accountability, and improves or degrades the shared factual foundation required for collective action.

## Promise compared with present reality

### AI's Promise Compared With Today's Reality

AI promises faster scientific discovery, personalized education, better medical support, greater productivity, safer infrastructure, improved accessibility, stronger defense, and wider access to expertise. Present reality is more uneven.

Current systems can be extraordinarily capable while remaining brittle, opaque, expensive, energy-intensive, vulnerable to manipulation, and dependent upon imperfect data. They can outperform people on selected tests yet fail at ordinary context, source verification, or reliable execution. Many organizations announce AI programs before establishing adequate data, training, security, evaluation, or accountability.

The investigation should compare every promise with measurable evidence:

- What capability exists today rather than in a demonstration?
- Under what conditions does it work?
- How often does it fail, and how serious are the failures?
- Does it improve outcomes or merely reduce labor and cost?
- Who receives the benefit, and who bears the risk?
- Can results be independently reproduced?
- What infrastructure, education, energy, and institutional reform are required?
- Does the system preserve human freedom and responsibility?

# The human standard

The human standard

## What AI must strengthen—not replace

Artificial intelligence should be evaluated not only by what it can produce, but by what its use causes human beings to retain, develop, surrender or lose. The standard is whether AI deepens inquiry, competence, creativity, truthfulness and responsible action without exchanging those capacities for speed, convenience, certainty or the appearance of wisdom.

### Reasoning

Does AI strengthen a person's ability to examine evidence and explain a conclusion—or bypass the reasoning the task was meant to develop?

### Authorship

Can the person distinguish assistance from substitution and honestly identify which knowledge, language, judgment and choices remain their own?

### Judgment

Does the system improve human judgment, including the ability to recognize uncertainty and challenge error, or merely accelerate a decision?

### Responsibility

Does an identifiable person or institution retain authority, accountability and the duty to correct harm?

### Relationships

Does AI support human communication, teaching and care—or substitute for relationships on which development, trust and belonging depend?

### Independent capability

After assistance is removed, can the person still understand, explain, verify, choose and act without the system?

### GIB determination

A polished product, faster completion or improved assisted score does not by itself establish human improvement. Durable capability must be demonstrated after assistance ends, and consequential decisions must remain subject to meaningful human authority.

# Productivity and human flourishing

AI's promise, productivity and present reality

## Efficiency is an opportunity—not a complete human objective

AI can produce profound gains in productivity and economic efficiency by reducing routine work, accelerating research, improving access to expertise and helping smaller organizations perform tasks once reserved for large institutions. Those gains should be treated as an opportunity for self-improvement, professional advancement, entrepreneurship and more meaningful work—not merely as a method for reducing headcount or concentrating control.

## Machine capability

Accuracy, reliability, speed, cost, reproducibility, security, energy efficiency and robustness outside a test set.

## Human flourishing

Knowledge retention, skill development, autonomy, privacy, due process, opportunity, accessibility, institutional trust and the ability to challenge the system.

Science fiction often imagines intelligent machines seizing control. The more practical warning is that humans may become too dependent, allowing convenience to erode our grasp of truth, dilute moral agency and substitute automated fluency for authentic wisdom. This investigation tests that concern as a hypothesis—not as a prediction.

FIGURE 11

Productivity and human flourishing




Figure 11. Measure who benefits, what skills grow, what risks rise and whether people advance.

# The two scoreboards

## Dual Scoreboards: Machine Capability and Human Flourishing

Every major application should be evaluated on two scoreboards. Improvement on the first does not establish improvement on the second.

AI capability scoreboard: Accuracy, reliability, speed, cost per task, reproducibility, security, energy and resource efficiency, robustness outside the test set, and performance compared with both a human and the best available alternative.

Human flourishing scoreboard: Quality of decisions, knowledge retention, skill development, autonomy, privacy, due process, accessibility, distribution of benefits, accountability, public trust, relationship quality, resilience during system failure, and ability to challenge the system.

The central longitudinal measure is whether machines and the people using them become more capable together. A deployment that raises output while reducing independent knowledge, freedom, resilience, or responsibility should not be called an unqualified success.

FIGURE 12

The two scoreboards




FIGURE 13

## Purposeful human direction

Responsible and purposeful human direction

**Define the objective, select the least risky tool, verify the result, and preserve accountable authority**




Figure 13. Responsible use begins with a legitimate human objective and measures whether AI actually improves it.

# Responsibilities by audience

## Chapter guidance

### What Humans Must Do to Harness AI

#### Individuals

People should learn to use AI as a collaborator, not an unquestioned authority. They should verify important claims, protect confidential information, preserve foundational skills, ask for counterarguments, and retain responsibility for decisions.

#### Parents

Parents should understand which AI systems children use, what information is collected, how synthetic relationships affect development, and whether AI is helping a child think or merely completing the work. Children need protected spaces to read, write, calculate, create, struggle, and communicate without automatic assistance.

[Open the practical parent and educator learning method<sup>13</sup>](#)

#### Related analysis

The education investigation applies this principle by developmental stage and evaluates whether particular AI modes support learning or quietly perform the learner's work.

[Open Artificial Intelligence and Learning<sup>17</sup>](#)

#### Educators

Schools should teach AI literacy alongside subject knowledge. Students should learn how models work, why they hallucinate, how bias enters, how to verify sources, how to disclose AI assistance, and when independent work is necessary. Assessments may need to emphasize oral explanation, supervised work, practical demonstration, and revision history.

[Open the practical parent and educator learning method<sup>13</sup>](#)

#### Businesses

Organizations should identify approved uses, protected data, accountable owners, validation standards, incident procedures, and tasks that must remain human-controlled. Productivity gains should be measured against errors, security, customer trust, skill retention, and long-term dependence.

#### Government

Government should protect constitutional rights, competition, national security, children, privacy, and due process while enabling innovation. It should invest in research, education, energy, chips, secure infrastructure, and technical expertise. Consequential government AI should be logged, auditable, appealable, and subject to legislative and judicial oversight.

## Overview audience checklist

### Individuals Use AI as a collaborator—not an unquestioned authority.

Verify important claims, protect confidential information, preserve foundational skills, ask for counterarguments, and retain responsibility for consequential decisions.

[Read the complete individual standard<sup>40</sup>](#)

### Parents Determine whether AI is helping a child think or merely completing the work.

Understand which systems children use, what information is collected, and how synthetic relationships may affect development. Preserve protected spaces for children to read, write, calculate, create, struggle, and communicate without automatic assistance.

[Read the complete parent standard<sup>41</sup>](#)

[Open the parent and educator learning method<sup>13</sup>](#)

## Responsibilities by audience · continued

### Educators Teach AI literacy alongside independent subject mastery.

Teach how models work, why they hallucinate, how bias enters, how sources are verified, how assistance is disclosed, and when unaided work is necessary. Preserve oral explanation, supervised work, practical demonstration, and revision history.

[Read the complete educator standard<sup>42</sup>](#)

[Open the parent and educator learning method<sup>13</sup>](#)

### Businesses Measure productivity against risk, trust, skill, and dependency.

Define approved uses, protected data, accountable owners, validation standards, incident procedures, and tasks that must remain human-controlled. Reinvest useful gains in people and durable capability.

[Read the complete business standard<sup>43</sup>](#)

### Government Protect rights and national capability while enabling useful innovation.

Invest in research, education, energy, chips, secure infrastructure, and technical expertise. Consequential systems should be logged, auditable, appealable, and subject to legislative and judicial oversight.

[Read the complete government standard<sup>44</sup>](#)

### Developers Design for verification, challenge, security, and accountable human control.

Document limitations, disclose material failures, secure models and supply chains, support independent evaluation, preserve meaningful human override, and identify the party responsible for correction and remedy.

[Read the supporting governance standard<sup>45</sup>](#)

# The complete human-readiness board

## Human Readiness Board

The investigation should periodically assess whether society is becoming more capable of governing AI.

### Knowledge

Do citizens understand AI's capabilities, limitations, hallucinations, and data requirements?

### Judgment

Can users distinguish evidence from generated plausibility and resist automation bias?

### Institutions

Do schools, businesses, courts, and agencies possess trained people, audit systems, appeals, and fallback procedures?

## Rights

Are privacy, free expression, due process, ownership, consent, and human autonomy protected?

## Security

Are systems resilient against poisoning, theft, prompt injection, espionage, and infrastructure disruption?

## Opportunity

Are AI benefits broadly available, or are power and capability concentrating among a small number of governments and corporations?

## Accountability

Can responsibility be identified when an AI-assisted decision causes harm?

*See Figure 14. Original caption: No readiness score should be assigned until each category is measured with evidence.*

FIGURE 14

The complete human-readiness board



Figure 14. Human readiness must be assessed across seven evidence-based dimensions.

# The concluding argument

## Conclusion

### **The future of AI is not determined by technology alone**

The central danger is not only that AI becomes powerful. It is that humans become intellectually passive, morally evasive, institutionally dependent or unwilling to accept responsibility. The central promise is not merely faster automation. It is the possibility that informed people can use AI to expand knowledge, creativity, health, security, productivity, accessibility and opportunity without surrendering truth, freedom, judgment or human dignity.

*AI should be developed and used responsibly, deliberately and purposefully—as an instrument of examined human judgment, not a replacement for it.*

## **Conclusion**

AI's future will not be determined by technology alone. It will be shaped by the questions humans ask, the information systems are permitted to use, the objectives developers reward, the institutions that deploy them, the laws governments enact, and the habits society chooses to preserve.

The central danger is not merely that AI becomes powerful. It is that humans become intellectually passive, morally evasive, institutionally dependent, or unwilling to accept responsibility. A society may gain unprecedented access to generated answers while losing the patience, knowledge, memory, experience,

and courage required to determine whether those answers are true or good. The appearance of wisdom could expand while the practice of wisdom declines.

Serious harm does not require a machine to seize power. It can begin when people surrender the habits that make truth discoverable, moral choice meaningful, independent judgment possible, and self-government sustainable.

The central promise is not merely faster automation. It is the possibility that informed people can use AI to expand knowledge, creativity, health, security, productivity, accessibility, and opportunity without surrendering truth, freedom, judgment, or human dignity. Realizing that promise requires more than responsible technology companies. It requires purposeful human use.

Responsible use means that authority remains proportional to evidence and risk; important claims are verified; affected people retain rights, explanation, and appeal; private information is protected; failures are disclosed; and an identifiable person or institution remains accountable. Purposeful use means beginning with a legitimate human objective, selecting AI only when it improves that objective, measuring the result, and refusing uses that merely substitute convenience, control, or novelty for genuine benefit.

AI should extend human capability without becoming an excuse to abandon it. Children should still learn to read, write, calculate, remember, debate, create, and reason without automatic assistance. Professionals should preserve the expertise needed to challenge a system. Citizens should demand original evidence and resist generated certainty. Institutions should reward correction rather than concealment. Governments should protect liberty, competition, due process, security, and the conditions in which independent thought remains possible.

## The concluding argument · continued

In creating this investigation, I found AI extremely helpful in research, data collection, comparison, organization, and analysis. It nevertheless required significant human oversight, content formation, source judgment, writing, and editing. While striving for accuracy and completeness, neither the AI nor I can claim perfection; the evidence, transparency, and effort should speak for themselves.

*The investigation therefore concludes with an argument and a continuing test: AI should be developed and used responsibly, deliberately, and purposefully—as an instrument of examined human judgment, not a replacement for it. Its success should be measured not only by what machines can produce, but by whether the people and societies using them become more truthful, capable, free, ethical, and wise.*

## PRACTICAL GUIDE

# AI-assisted learning

A method-first guide for parents and educators, followed by the evidence, adoption concerns and physical context.

## Purpose and evidence boundary

A logical, evidence-informed approach for using AI to improve understanding while preserving independent knowledge, reasoning, expression and responsibility.

*Project, Research, Analysis, Editorial by Patrick Russell; assisted by ChatGPT (OpenAI).*

Central assessment

**AI can improve learning—but only when the method preserves the learner’s cognitive work**

Existing studies do not support either universal AI use or universal prohibition. They support a narrower

conclusion: structured tutoring, guided questions, targeted feedback and human supervision can improve learning, while unrestricted answer generation can produce better work during use without producing equal independent mastery afterward.

The proposed method below is a synthesis of those findings and established learning principles. It is not presented as a universally validated program. It should be piloted, measured and revised by age, subject and student need.

FIGURE 15

## Purpose and evidence boundary



Figure 15. AI supports the learning relationship; the student still reasons, reconstructs and demonstrates mastery.

# The LEARN cycle

Proposed GIB synthesis

The goal is not merely to finish a task. It is to ensure that assistance ends in independent capability.

- **Launch independently**

Begin with an unaided attempt, prediction, outline or explanation. This reveals prior knowledge and preserves productive struggle.

- **Explain current reasoning**

State what is understood, show the work already completed and identify the precise point of confusion.

- **Ask for guided assistance**

Request one hint, question, counterexample or targeted explanation at a time—not a finished answer.

- **Reconstruct without AI**

Close the tool and solve, write or explain the task again in the student's own words.

- **Note mastery and next steps**

Use a similar AI-free task later to test retention, transfer and whether further practice is required.

*AI should increase the amount and quality of thinking a student can eventually do without AI—not merely the quality of the product submitted while AI is present.*

## Three declared learning modes

Match the tool to the objective

### Three clearly declared learning modes

Mode 1

#### AI-free practice

Use when retrieval, fluency, original expression or independent diagnosis is the learning objective.

**Examples:** Foundational arithmetic, vocabulary recall, first thesis construction, supervised writing and oral explanation.

Mode 2

#### AI-guided learning

Use when a learner needs hints, feedback, examples or questions while still doing the central cognitive work.

**Examples:** Socratic math tutoring, critique of a student-written draft, practice questions and misconception checks.

Mode 3

#### AI-permitted production

Use when the objective includes research, iteration, tool literacy or producing work with modern assistance.

**Examples:** Source comparison, debugging, accessibility support, editing and documented human–AI collaboration.

Every assignment should identify its mode before work begins. A hidden or ambiguous AI policy turns a learning decision into an integrity trap.

# The method across subjects

## Application examples

### What the method looks like across subjects

#### Mathematics

Attempt the problem first. Ask AI to identify only the first incorrect step or supply one hint. Solve a parallel problem without AI and repeat it later from memory.

#### Writing

Create the thesis and outline independently. Ask AI to challenge assumptions, identify missing evidence and suggest questions. Revise personally, then defend the argument orally without AI.

#### History and reading

Read the assigned source before consulting AI. Ask for questions rather than a summary. Support or reject every AI claim with the original text, then write an unaided synthesis.

#### Science

Make a prediction and explain the mechanism first. Use AI to propose competing explanations or identify a missing variable. Evaluate the evidence and finish with an independent causal explanation.

#### Language learning

Respond before requesting help. Ask AI to mark selected errors and explain the rule. Reproduce the corrected form later without assistance.

# Responsibilities in learning

## Human responsibilities

### Each participant has a distinct role

#### Parents

Know which systems are used, protect personal data, ask children to explain their work, and value demonstrated learning over polished output.

#### Educators

Name the learning objective, designate the permitted mode, model good prompts, retain AI-free assessment, and compare assisted with independent performance.

#### Students

Disclose assistance, show the original attempt, verify factual claims, reconstruct the work independently and report when AI caused confusion.

#### AI systems

Prefer hints and questions, reveal uncertainty, avoid unnecessary personal data, cite sources when available and allow teachers or parents to set constraints.

# Age and autonomy

Developmental calibration

## Increase autonomy gradually

### Elementary years

Adult-mediated use, limited free-form chat, strong privacy protection and priority for direct reading, writing, calculation, play and human relationships.

### Middle-school years

Structured prompts, visible histories, short guided sessions, explicit source checking and frequent AI-free demonstrations of mastery.

### High-school years

Selective autonomy, disclosure of material assistance, independent assessment, source verification and practice choosing when not to use AI.

# Pilot checklist and two scoreboards

From proposal to evidence

## Pilot the method and measure both scoreboards

- **Establish a baseline.** Measure independent knowledge, reasoning, confidence and task completion before introducing AI.
- **Run a limited pilot.** Use the LEARN cycle for two to four weeks on selected assignments, with a comparison task or group when feasible.
- **Measure assisted performance.** Record quality, completion, time, engagement and accessibility while AI is available.
- **Measure independent learning.** Test recall, explanation, transfer and problem-solving after AI is removed—including a delayed check.
- **Inspect distribution and dependence.** Determine who benefited, who was confused, whether gaps widened and whether students can still work independently.
- **Revise the rules.** Expand only the uses that improve both immediate performance and later independent mastery.

# Prompts that teach rather than complete

Starter prompts

## Ask AI to teach rather than complete

“Do not give me the answer. Ask one question at a time that helps me identify the next step.”

“I will show my attempt. Identify only the first incorrect or unsupported step, explain why it is a problem, and let me revise it.”

“Challenge my explanation with a counterexample or missing assumption. Do not rewrite it for me.”

“After we finish, give me a different problem that tests the same principle, then wait while I solve it without help.”

## The evidence and its limits

The practical sequence above is a proposed synthesis. These studies explain its basis and its limits; their populations, interventions and outcomes are not interchangeable.

What has been tested

### The strongest evidence is promising and conditional

Randomized crossover trial

#### AI tutoring in undergraduate physics

**Finding:** A carefully designed AI tutor produced larger immediate learning gains in less time than an active-learning class session.

**Boundary:** One university course; the study does not establish effects across K–12 subjects or multiple years.

[Read the study](#)<sup>96</sup>

Large field experiment

#### Generative AI in high-school mathematics

**Finding:** General chatbot access improved practice performance but could weaken later unaided performance; tutor-style guardrails reduced that risk.

**Boundary:** The result concerns a particular subject, implementation and study period—not permanent intelligence loss.

[Read the study](#)<sup>97</sup>

Randomized controlled trial with human tutors

#### Tutor CoPilot in K–12 mathematics

**Finding:** Real-time AI guidance helped tutors use more guiding questions and improved topic mastery, especially for lower-rated tutors.

**Boundary:** The intervention assisted human tutors; it did not test unsupervised student use or long-term retention.

[Read the study](#)<sup>98</sup>

Proctored randomized experiment; preprint

#### Explanation versus automation

**Finding:** Immediate gains were more durable when learners used AI for explanation and augmentation than when they automated the work.

**Boundary:** The paper is a 2026 preprint involving undergraduates and should be treated as provisional.

[Read the study](#)<sup>99</sup>

## The evidence and its limits · continued

### Established

Assisted output and immediate task performance often improve. Design and supervision materially affect learning outcomes.

### Supported but bounded

Hint-based tutoring, feedback and explanation can improve mastery in tested settings.

### Uncertain

Optimal dosage, long-term developmental effects, subject differences and durable population-scale results.

## Adoption, resistance and public trust

Adoption, resistance and public trust

### Who is resisting AI in education—and what are they actually saying?

Resistance is not a single movement. It includes teachers protecting professional judgment, parents protecting children, students concerned about fairness, civil-liberties groups opposing surveillance, privacy advocates challenging data extraction, researchers questioning weak evidence, and institutions responsible for enforcing existing law. Some people oppose particular uses; relatively few of the major organizations below advocate eliminating educational AI altogether.

No defensible source can identify *all* organizations or agencies resisting AI. Local policies change quickly, and opposition may arise from individuals, school faculties, parent groups, unions, civil-rights organizations, environmental coalitions or elected bodies. The following list identifies principal documented sources relevant to this investigation and records their stated concerns rather than labeling every safeguard as anti-AI.

#### National Education Association (NEA)

**What they are saying:** Warns against treating generated content as a primary authority and emphasizes educator voice, equity, transparency and human judgment.

**Constructive response:** Educator-led standards, professional development and evidence that AI improves learning rather than merely reducing labor.

[Review the organization's source](#)<sup>100</sup>

#### American Federation of Teachers (AFT) and affiliates

**What they are saying:** Raises concerns about student privacy, bias, worker displacement, intellectual property, unsafe automation and corporate mandates.

**Constructive response:** Its stated position is guarded adoption—not categorical rejection. Involve teachers in procurement, policy and evaluation, and fund practical training.

[Review the organization's source](#)<sup>101</sup>

#### UNESCO

**What they are saying:** Warns that deployment has outpaced governance, teacher preparation and evidence, especially for children. Emphasizes privacy, age appropriateness, inclusion and human agency.

**Constructive response:** Use human-centered design, age limits, teacher competencies and ethical validation before scale.

[Review the organization's source](#)<sup>102</sup>

## Adoption, resistance and public trust · continued

### Common Sense Media and Youth AI Safety Institute

**What they are saying:** Evaluates developmental, privacy, misinformation, companion and child-safety risks associated with AI products.

**Constructive response:** Use independent product evaluation, age-based ratings, parental information and child-development standards.

[Review the organization's source](#)<sup>103</sup>

### Electronic Frontier Foundation (EFF)

**What they are saying:** Opposes invasive student surveillance, excessive data collection, unreliable monitoring and opaque school technology practices.

**Constructive response:** Minimize data, provide meaningful alternatives and prohibit surveillance uses that are not necessary or demonstrably effective.

[Review the organization's source](#)<sup>104</sup>

### Electronic Privacy Information Center (EPIC)

**What they are saying:** Challenges data-extractive, mistake-prone and unaccountable systems in education, including consequential automated decisions.

**Constructive response:** Require privacy and algorithmic impact assessments, disclosure, appeal rights and independent oversight.

[Review the organization's source](#)<sup>105</sup>

### Center for Democracy & Technology (CDT)

**What they are saying:** Focuses on student privacy, civil rights, transparency and the risks created by opaque educational technology.

**Constructive response:** Publish plain-language data practices, document system limitations and involve affected communities before adoption.

[Review the organization's source](#)<sup>106</sup>

### U.S. Department of Education

**What they are saying:** Supports educational AI but warns that privacy, security, civil rights, accessibility and human-centered design must govern its use.

**Constructive response:** Clarify allowable uses, train educators, retain human authority and align deployment with FERPA and civil-rights obligations.

[Review the organization's source](#)<sup>107</sup>

### Federal Trade Commission and Education privacy authorities

**What they are saying:** Enforce children's privacy and student-record protections. These legal conditions are sometimes mistaken for opposition to technology.

**Constructive response:** Select products that collect only necessary information and contractually prohibit unrelated advertising, profiling, resale and indefinite retention.

[Review the organization's source](#)<sup>108</sup>

### School boards, districts, faculties, parents and students

**What they are saying:** Local resistance often reflects unclear rules, cheating concerns, inadequate training, inequitable access, unreliable tools or exclusion from the decision.

**Constructive response:** Use transparent local pilots, publish results, create opt-out or equivalent alternatives where appropriate, and revise policy with stakeholder participation.

[Review the organization's source](#)<sup>109</sup>

## How resistance affects acceptance

### Potential costs

Blanket bans can drive use underground, produce inconsistent rules, prevent students from learning responsible use, widen access gaps and delay beneficial tutoring or accessibility tools.

## Adoption, resistance and public trust · continued

### Potential benefits

Scrutiny can expose weak products, prevent surveillance, protect teacher and student rights, improve evidence standards and force vendors to build safer, more useful systems.

### What should be done about the resistance?

The objective should be to **reduce unjustified resistance and resolve justified resistance**. Attempts to defeat, silence or bypass educators, parents or affected communities are likely to deepen mistrust. Acceptance should be earned through limited trials, independently measured learning outcomes, transparent data practices, educator control, student

protections and the ability to stop deployments that fail.

The best leader is not an AI vendor, a distant agency or a single advocacy organization. The most credible adoption body is a **local educator–parent partnership**, formally supported by school leadership and advised by independent learning researchers, privacy specialists and student representatives. Educators should lead pedagogy; parents and students should shape safeguards; districts should manage procurement and legal compliance; independent evaluators should measure results. State and federal agencies should provide standards, funding and enforcement rather than dictate one classroom method.

## Infrastructure behind educational AI

The physical requirements behind AI

AI-assisted education ultimately depends on data centers capable of delivering reliable, affordable computing services. Expanding that capacity can support tutoring, translation, accessibility, research and personalized instruction, but it also creates legitimate questions about electricity, water, land, utility costs, environmental effects and community consent.

Those infrastructure concerns can influence public acceptance of AI even when the educational application itself is beneficial. Responsible adoption therefore requires both sound learning methods and responsible physical development.

**For the complete infrastructure analysis:** read the GIB investigation [Data Centers and Public Value](#)<sup>21</sup>

## What remains when the tool is removed

Working conclusion

### The proper test is what remains when the tool is removed

AI-assisted learning should be judged on two results: the quality of work produced with assistance and the learner's later ability to understand, remember, explain and apply the material independently. A sound method seeks gains on both. When those results diverge, independent learning must take priority.

Government and institutional guidance similarly favors human-centered, age-appropriate and supervised implementation rather than treating access alone as educational progress. Durable adoption also requires institutions to answer legitimate objections and to build the physical infrastructure supporting AI without concealing or transferring its costs.

[UNESCO guidance on generative AI in education](#)<sup>110</sup>

[U.S. Department of Education: AI and the future of teaching and learning](#)<sup>111</sup>

## REFERENCE MATERIAL

# Methods, sources and review record

The source catalogues, evidence rules and original-publication context are preserved outside the continuous argument.

## Investigation and evidence rules

The investigation will:

- Prefer original documents, data, code, laws, audits, and direct testimony.
- Distinguish facts, official claims, expert judgments, inferences, and speculation.
- Identify the jurisdiction and authority of each source.
- Separate proposed policy from enacted and implemented policy.
- Disclose material financial, political, corporate, and institutional interests.
- Compare opposing interpretations using the strongest available evidence.
- Distinguish correlation, plausible mechanism, contribution, and demonstrated causation.
- Do not treat consensus, repetition, popularity, institutional status, or algorithmic reach as substitutes for evidence.
- Record uncertainty instead of manufacturing certainty.
- Preserve dated predictions and later compare them with outcomes.
- Open and verify citations rather than trusting their appearance.
- Correct errors visibly and explain what changed.

No institution will be trusted or rejected solely because of its name. United States constitutional and governmental sources will receive priority within their lawful areas of authority. Every source—including the White House, OSTP, PCAST, corporations, universities, international organizations, and AI companies—will still be evaluated for evidence quality, transparency, jurisdiction, incentives, conflicts of interest, predictive accuracy, and measurable results.

Source rule: prefer original documents, law, data, audits and direct testimony; distinguish fact from official claim, expert judgment, inference and speculation; disclose material financial, political, corporate and institutional interests; preserve dated predictions; verify consequential citations.

# Complete annotated source catalogue

- 01  
**President Donald J. Trump — Truth Social statement on AI and data centers**  
Primary presidential statement opposing additional AI guardrails, invoking existing enforcement authority and framing U.S.–China competition.  
Open source<sup>48</sup>
- 02  
**President Donald J. Trump — follow-up Truth Social statement on AI risk and infrastructure**  
Primary presidential statement rejecting catastrophic-AI warnings, promoting rapid infrastructure growth and attributing Google’s Finland investment to U.S. permitting difficulty.  
Open source<sup>49</sup>
- 03  
**President Donald J. Trump — Truth Social statement on data-center opponents**  
Primary presidential statement dismissing neighborhood objections to data centers and alleging that unnamed opponents serve interests contrary to the United States.  
Open source<sup>50</sup>
- 04  
**President Donald J. Trump — Truth Social statement announcing an ‘AI Force’**  
Primary presidential statement promising an AI Force and future AI Czar while asserting broad economic and community benefits.  
Open source<sup>51</sup>
- 05  
**President Donald J. Trump — Truth Social endorsement of an ‘AI Disaster Hoax’ interview**  
Primary presidential statement amplifying Matthew Continetti’s argument that catastrophic-AI fears should not displace competition with China.  
Open source<sup>52</sup>
- 06  
**President Donald J. Trump — Truth Social statement encouraging ‘super intelligence’**  
Primary presidential statement rejecting broad catastrophic-AI warnings, favoring development of AI or ‘super intelligence’ and invoking Justice Department and law-enforcement authority if intervention becomes necessary.  
Open source<sup>53</sup>
- 07  
**President Donald J. Trump — Truth Social statement on Xi Jinping and ‘super intelligence’**  
Primary presidential statement characterizing President Xi’s reaction to the new terminology after their September 24 meeting; the characterization is not presented as a direct quotation from Xi.  
Open source<sup>54</sup>
- 08  
**President Donald J. Trump — Truth Social statement excluding Scott Bessent as SI Czar**  
Primary presidential statement saying Treasury Secretary Scott Bessent will remain at Treasury and will not become the proposed Super Intelligence Czar.  
Open source<sup>55</sup>
- 09  
**President Donald J. Trump — Truth Social repost of Fox News AI commentary**  
Verified September 28 repost, without accompanying presidential text, of a Fox News segment discussing U.S.–China AI competition, economic and military transformation, and AI risk. The segment’s statements are not direct presidential quotations.  
Open source<sup>56</sup>
- 10  
**President Donald J. Trump — Truth Social post of White House Super Intelligence Luncheon seating chart**  
Verified September 29 image-only post identifying senior administration officials, congressional leadership and executives from major AI, semiconductor, cloud, cybersecurity and technology companies expected at the White House luncheon. The post contains no policy statement or meeting outcome.  
Open source<sup>38</sup>

## Complete annotated source catalogue · continued

- 11  
**President Donald J. Trump — White House Accord on Super Intelligence**  
 Verified September 29 image post of the signed one-page accord establishing four voluntary layers of company controls and audits for frontier systems. Signatories shown are President Trump and leaders of Google, Anthropic, Meta, OpenAI, xAI and Nvidia.  
 Open source<sup>34</sup>
- 12  
**President Donald J. Trump — Super Intelligence Force announcement**  
 Verified October 4 presidential statement announcing the Super Intelligence Force, describing its coordination role and naming four federal leaders. The statement is primary evidence of the President's announcement; it is not by itself a public charter, statute or proof of implementation.  
 Open source<sup>28</sup>
- 13  
**President Donald J. Trump — statement on AI terminology**  
 Verified October 8 presidential statement calling people who continue to use 'Artificial Intelligence' instead of 'Super Intelligence' enemies. The post is evidence of the President's rhetoric and terminology preference; it does not itself create a penalty, legal classification or new authority.  
 Open source<sup>24</sup>
- 14  
**Reuters — Super Intelligence Force leadership and reported charter**  
 Contemporaneous reporting, based on a Wall Street Journal interview with Jay Clayton, that the force is expected to report within 120 days and review AI risks, opportunities and federal response mechanisms. The underlying charter had not been located in the public record reviewed here.  
 Open source<sup>29</sup>
- 15  
**Office of the Director of National Intelligence — Jay Clayton**  
 Official record confirming that Jay Clayton serves as Director of National Intelligence.  
 Open source<sup>30</sup>
- 16  
**Federal Trade Commission — Andrew N. Ferguson**  
 Official record confirming that Andrew Ferguson serves as FTC chairman.  
 Open source<sup>31</sup>
- 17  
**Department of War — Emil Michael**  
 Official biography confirming that Emil Michael serves as Under Secretary for Research and Engineering and chief technology officer.  
 Open source<sup>32</sup>
- 18  
**Office of Personnel Management — Scott Kupor**  
 Official biography confirming that Scott Kupor serves as OPM director.  
 Open source<sup>33</sup>
- 19  
**White House — Inaugurating the Era of Super Intelligence**  
 September 29 executive order directing executive-branch use of 'Super Intelligence' and 'SI,' preserving the existing statutory AI definition unless changed by law or later presidential action, and requiring proposed legislative language within 60 days.  
 Open source<sup>25</sup>
- 20  
**White House — Fact Sheet on the Super Intelligence executive order**  
 Official summary of the terminology order and the administration's stated innovation rationale.  
 Open source<sup>35</sup>
- 21  
**Axios — White House terminology implementation**  
 Preferred-source contemporaneous reporting that the September 29 order specifies no penalty for using 'AI' and that White House pages continued to use the term after the President's October 8 post.  
 Open source<sup>27</sup>

## Complete annotated source catalogue · continued

- 22  
**Reuters — White House Accord controls and signatories**  
Contemporaneous reporting transcribing the accord's four control layers, regular-meeting commitment, signatories and statement that later legal codification may make sense.  
Open source<sup>36</sup>
- 23  
**CBS News — President Trump's remarks on the voluntary accord**  
Contemporaneous reporting that President Trump called the accord morally binding, described industry self-policing and said an oversight committee and AI Czar were being considered; those proposed steps are not treated as completed actions.  
Open source<sup>37</sup>
- 24  
**White House — Streamlining Access to Government Services Through America.gov**  
September 29 executive order establishing America.gov and requiring connected super-intelligence systems to be accurate, reliable and transparent while preserving agency authority, privacy protections and non-digital access channels.  
Open source<sup>39</sup>
- 25  
**Reuters — White House AI meeting and President Trump's same-day China remarks**  
Contemporaneous reporting identifying the anticipated innovation-and-oversight discussion, the President's opposition to joint U.S.–China AI governance and the White House's lack of a public meeting agenda before the luncheon.  
Open source<sup>57</sup>
- 26  
**Axios — White House AI luncheon attendee list**  
Contemporaneous reporting corroborating the 31-person attendee list while characterizing the administration's approach; the characterization is treated as reporting, not as an official meeting outcome.  
Open source<sup>58</sup>
- 27  
**White House — President Trump at the United Nations**  
Official September 22 record describing 'Super Intelligence' as the administration's new terminology and rejecting global control of the technology.  
Open source<sup>59</sup>
- 28  
**Chinese Ministry of Foreign Affairs — Xi-Trump meeting readout**  
Official Chinese account confirming substantive AI discussion on competition, cooperation, risks, benefits, misuse prevention and human control; it does not report Xi endorsing the 'Super Intelligence' label.  
Open source<sup>60</sup>
- 29  
**Associated Press — State Department implementation of the terminology change**  
Contemporaneous reporting based on an internal State Department email directing one bureau to use 'Super Intelligence' and 'SI' in documents and communications; government-wide scope remained unclear.  
Open source<sup>61</sup>
- 30  
**U.S. Treasury — Secretary Scott Bessent**  
Official record confirming Bessent's continuing role as the 79th Secretary of the Treasury.  
Open source<sup>62</sup>
- 31  
**Council of Economic Advisers — 2026 Economic Report of the President**  
Official report using the conventional technical meaning of artificial superintelligence: AI whose intelligence surpasses that of humans.  
Open source<sup>63</sup>
- 32  
**White House — Promoting Advanced Artificial Intelligence Innovation and Security**  
June 2, 2026 executive order directing prioritized enforcement of specified federal criminal statutes against AI-enabled computer crimes while establishing a voluntary frontier-model security framework without mandatory model licensing.  
Open source<sup>64</sup>

## Complete annotated source catalogue · continued

33

### **Stanford Institute for Human-Centered Artificial Intelligence — 2026 AI Index**

Metric-specific comparison showing that the U.S.–China model-performance gap had effectively closed while the United States retained advantages in top-tier models, high-impact patents and private investment and China led other research, patent and industrial-robot measures.

Open source <sup>65</sup>

34

### **American Enterprise Institute — AI Competition with China Continues to Heat Up**

Preferred-reference commentary recording a seven-month estimate from America First Policy Institute testimony while also citing Stanford evidence that China had nearly erased the broad model-performance gap; it illustrates why any single lead estimate needs a named metric and date.

Open source <sup>66</sup>

35

### **White House — National Security Memorandum on AI leadership**

April 23, 2026 primary policy record treating AI leadership as an economic and national-security objective and identifying China-based efforts to obtain U.S. frontier-model capabilities.

Open source <sup>67</sup>

36

### **Associated Press — U.S. proposes AI incident notification mechanism with China**

Contemporaneous reporting that the administration is pursuing cross-border notification for AI incidents affecting national security, demonstrating a parallel risk-management track.

Open source <sup>68</sup>

37

### **Fox News — Matthew Continetti on AI risk and China competition**

The interview endorsed by President Trump; it argues that Chinese control of advanced AI could be used for malevolent ends.

Open source <sup>69</sup>

38

### **Google — €13 billion Finland AI-infrastructure announcement**

Company record of the investment, locations, jobs and energy partnerships; it does not attribute the decision to U.S. permitting.

Open source <sup>70</sup>

39

### **International Department of the CPC Central Committee — Xi Jinping on AI development and governance**

Official Chinese record calling for development while also emphasizing risk awareness, human control and secure, controllable AI.

Open source <sup>71</sup>

40

### **Virginia JLARC — Data Centers in Virginia**

State audit documenting both economic benefits and project-level effects involving power costs, water, land use, noise and nearby residents.

Open source <sup>72</sup>

41

### **White House Office of Science and Technology Policy**

Federal science, technology and AI policy direction.

Open source <sup>73</sup>

42

### **AI.gov and America's AI Action Plan**

National AI policy, implementation and federal coordination.

Open source <sup>74</sup>

43

### **National AI Initiative Office**

Federal AI research and development coordination.

Open source <sup>75</sup>

44

### **NIST AI Risk Management Framework**

Voluntary risk-management functions, terminology and implementation resources.

Open source <sup>76</sup>

45

### **GAO AI Accountability Framework**

Governance, data, performance and monitoring questions for accountable systems.

Open source <sup>77</sup>

## Complete annotated source catalogue · continued

- 46  
**GAO Review of Federal Generative-AI Management**  
Audited evidence on federal governance and implementation gaps.  
Open source <sup>78</sup>
- 47  
**CISA Artificial Intelligence**  
Cybersecurity and critical-infrastructure guidance.  
Open source <sup>79</sup>
- 48  
**NSA Artificial Intelligence Security Center**  
AI-system security and national-security risk.  
Open source <sup>80</sup>
- 49  
**Federal Trade Commission AI Partnerships Study**  
Competition and major cloud-AI partnerships.  
Open source <sup>81</sup>
- 50  
**Department of Energy Data-Center Electricity Report**  
AI infrastructure and electricity-demand evidence.  
Open source <sup>82</sup>
- 51  
**European Commission AI regulatory framework**  
European legal and regulatory context.  
Open source <sup>83</sup>
- 52  
**OECD AI Principles**  
International policy principles and comparative context.  
Open source <sup>84</sup>
- 53  
**International AI Safety Report 2026**  
International review of general-purpose AI capabilities and risks.  
Open source <sup>85</sup>
- 54  
**UNESCO Recommendation on the Ethics of AI**  
International ethics framework—considered as normative, not dispositive authority.  
Open source <sup>86</sup>
- 55  
**SAG-AFTRA AI Resources**  
Labor, consent, compensation and likeness protections.  
Open source <sup>87</sup>
- 56  
**Authors Guild AI Advocacy**  
Copyright, licensing and author-rights arguments.  
Open source <sup>88</sup>
- 57  
**AFL-CIO Workers First Initiative on AI**  
Worker bargaining, surveillance and workplace deployment concerns.  
Open source <sup>89</sup>
- 58  
**PauseAI Statement**  
Organized advocacy for a pause—treated as an interested policy source.  
Open source <sup>90</sup>
- 59  
**Future of Life Institute Open Letter**  
Prominent pause proposal and signatory record.  
Open source <sup>91</sup>
- 60  
**Aristotle, Nicomachean Ethics**  
Philosophical grounding for human moral responsibility.  
Open source <sup>92</sup>

# Selected bibliography for continuing research

## Selected Primary and Institutional Sources for Continuing Research

- Aristotle, *Nicomachean Ethics*—Internet Classics Archive
- Stanford Encyclopedia of Philosophy—Aristotle’s Ethics
- White House Office of Science and Technology Policy
- AI.gov and America’s AI Action Plan
- White House AI Action Plan announcement
- President’s Council of Advisors on Science and Technology appointments
- National AI Initiative Office and federal AI coordination
- National Security Presidential Memorandum 11
- NIST AI Risk Management Framework
- OpenAI—Research and Deployment Mission
- Google DeepMind—Models and Research
- Anthropic—Research and AI Safety
- Meta AI—Models and Open Research
- SpaceXAI/xAI—Models and Compute Infrastructure
- NVIDIA—AI Data Center Infrastructure
- ISO/IEC JTC 1/SC 42—Artificial Intelligence Standards
- IEEE—Artificial Intelligence Standards Committee
- Noy and Zhang—Experimental Evidence on Generative-AI Productivity
- Brynjolfsson, Li, and Raymond—Generative AI at Work
- Vaccaro, Almaatouq, and Malone—Human-AI Performance Meta-Analysis
- Bastani and Colleagues—Generative AI Without Guardrails Can Harm Learning
- Lee and Colleagues—Generative AI and Critical Thinking
- NIST Adversarial Machine Learning Taxonomy
- GAO AI Accountability Framework
- GAO review of federal generative-AI management
- GAO—Artificial Intelligence Acquisitions
- CISA Artificial Intelligence
- NSA Artificial Intelligence Security Center
- Federal Trade Commission—AI Partnerships and Investments Study
- Department of Energy—United States Data Center Energy Use
- European Union AI regulatory framework
- OECD AI Principles
- International AI Safety Report 2026
- World Economic Forum—Who We Are
- World Economic Forum partners
- WHO ethics and governance of AI for health
- WHO guidance on large multimodal models
- WHO funding
- UNESCO Recommendation on the Ethics of Artificial Intelligence
- PauseAI—Statement and Proposed Pause
- Future of Life Institute—2023 Pause Letter
- Future of Life Institute—Statement on Superintelligence
- SAG-AFTRA—Artificial Intelligence Resources
- Authors Guild—Artificial Intelligence Advocacy
- Human Artistry Campaign
- AFL-CIO—Workers First Initiative on AI
- Senator Sanders—Announcement of Proposed AI Data Center Moratorium Act
- FBI—Cryptocurrency and AI Scams Bilk Americans of Billions
- FBI—Senior U.S. Officials Impersonated in Malicious Messaging

## Selected bibliography for continuing research · continued

- U.S. Treasury—Artificial Intelligence Resources and FinCEN Alerts
- Department of Justice—Conviction Involving AI-Generated Child Sexual Abuse Material
- FBI—Musician Charged in Alleged AI-Aided Streaming Fraud
- Europol—Operation Against AI-Generated Child Sexual Abuse Material
- Europol—Internet Organised Crime Threat Assessment 2025
- OpenAI and Microsoft—Disrupting Malicious Uses by State-Affiliated Actors
- OpenAI—Disrupting Malicious Uses of AI
- Office of the Director of National Intelligence—2026 Annual Threat Assessment
- NORC—General Social Survey
- American National Election Studies
- Pew Research Center—Public Trust in Government, 1958–2024
- Gallup—Confidence in Institutions
- National Archives—Research and Historical Records
- Library of Congress—Chronicling America Newspaper Archive

## Related investigations and edition navigation

### How related investigations contribute

- [Ai for Human Flourishing](#)<sup>1</sup>  
Owns measured benefit cases; this investigation explains the system behavior and human responsibility those uses depend on.
- [Fear, Power & the Future of Ai](#)<sup>2</sup>  
Owns harm and safeguard assessments; this investigation establishes the underlying mechanisms and accountable actors.
- [The Well-Formed Mind](#)<sup>3</sup>  
Owns independent-learning verification; this investigation supplies the model and authorship foundations.

Shared evidence contributes only when this investigation adds a distinct mechanism, consequence, comparison or decision implication. Preserve the original finding, source origin and review date; repeated coverage is not independent corroboration.

### Complete web edition

### The full investigation continues across five detailed chapters and a practical learning guide

This overview presents the central findings. The complete web edition preserves the investigation's 25-section structure, full reasoning, evidence distinctions, practical examples and concluding standards.

**Opening & Directive** The complete premise, accountability question and research mandate <sup>8</sup>

**Sections 1–4** Foundations, AI types, practical uses and the eleven-tool model test <sup>9</sup>

**Sections 5–10** Ownership, economics, dependency, privacy, authenticity and human impact <sup>10</sup>

**Sections 11–20** Ethics, experts, government, institutions, threats, opposition and criminal use <sup>11</sup>

**Sections 21–25** Truth over time, promise versus reality, human readiness and conclusion <sup>12</sup>

**Practical learning method** Evidence, resistance, educator responsibility and the LEARN cycle for parents and schools <sup>13</sup>

## Related investigations and edition navigation · continued

**Foundations of the Well Formed Mind** The philosophical and developmental basis for knowledge, judgment, attention, character and responsibility. <sup>16</sup>

**Artificial Intelligence and Learning** Evidence separating immediate AI-assisted performance from durable, independently retained learning. <sup>17</sup>

**The Synthetic You** When personalized AI can write, answer and perform as a person—and why output alone cannot establish authorship. <sup>18</sup>

**AI Regulations and Oversight** An active tracker separating binding law, proposed measures, technical standards and company rules. <sup>19</sup>

Related GIB Investigation **America Reloaded**

Connect AI and autonomous systems to semiconductor capacity, trusted electronics, cyber capability and the industrial foundations of national power.

**Read the innovation analysis** <sup>22</sup>

Related GIB Investigation **Global AI Power & Governance**

Compare AI capability, law, infrastructure and strategic position across countries.

**Read the global comparison** <sup>23</sup>

Related GIB Investigation **Fear, Power & the Future of AI**

Separate demonstrated harm, plausible risk and strategic danger from exaggerated claims.

**Read the risk investigation** <sup>2</sup>

Related GIB Investigation **AI for Human Flourishing**

Measure whether AI expands durable human capability and public benefit.

**Read the human-benefit investigation** <sup>1</sup>

Evidence, downloads and continuing research

## Inspect the full report and the original record

51 pages **Download the Full Illustrated Edition**  
Word document · complete Phase One investigation <sup>46</sup>

18 pages **Download the Visual Concise Edition**  
Word document · charts and condensed findings <sup>47</sup>

## Original chapter policy capsules

These original chapter passages are retained as source backing. Their full dated cases appear in Chapter 4. Keeping this companion record prevents qualifications or original wording from disappearing during consolidation.

### October 8, 2026: President Escalates the Terminology Campaign

President Donald Trump's verified Truth Social account said the White House considers anyone who uses 'Artificial Intelligence' instead of 'Super Intelligence' to be 'THE ENEMY.' This establishes the President's stated position and rhetoric. It does not, by itself, identify a person, allege a crime, create a penalty, or expand federal authority.

Established fact: the September 29 executive order directs executive departments and agencies, to the maximum extent permitted by law, to use 'Super Intelligence' and 'SI' in official communications and non-statutory documents. It preserves historical and previously issued legal records and temporarily applies the existing statutory definition of artificial intelligence. The order states no penalty for using the conventional term and does not bind private citizens, companies, Congress, or state governments.

Material contrary evidence: Axios reported on October 8 that White House pages still used 'AI' after the President's post. The White House technology-priority page likewise continued to describe United States leadership in artificial intelligence. This shows

## Original chapter policy capsules · continued

incomplete implementation; it does not invalidate the executive-branch directive.

Inference and unresolved questions: the post appears intended to pressure faster adoption of the preferred terminology, but no new executive order, regulation, enforcement memorandum, penalty, personnel action, or formal designation of any person as an enemy was identified. The record does not establish that ‘Super Intelligence’ has replaced ‘artificial intelligence’ in statute, technical standards, or general industry usage.

Related analysis

Primary presidential statement escalating the administration’s terminology campaign.

[Open the verified Truth Social post](#)<sup>24</sup>

Related analysis

Contemporaneous preferred-source reporting on the order’s lack of a stated penalty and incomplete White House implementation.

[Open the Axios report](#)<sup>27</sup>

### October 4, 2026: Super Intelligence Force Announced

President Donald Trump’s verified Truth Social account announced a federal Super Intelligence Force led by Director of National Intelligence Jay Clayton, FTC Chairman Andrew Ferguson, Under Secretary and chief technology officer Emil Michael, and OPM Director Scott Kapor. The President described the force as coordinating federal engagement with consumers, public-interest and religious organizations, critical-infrastructure providers, and artificial-intelligence companies, with reporting to the President and the White House chief of staff. Official agency records independently confirm the four named officials’ positions.

Established fact: the presidential announcement is public and identifies a cross-agency leadership structure. Reuters reported, based on a Wall Street Journal interview with Clayton, that the group is expected to review risks and opportunities, evaluate federal response mechanisms, and report within 120 days. The underlying charter was not located in the public materials reviewed for this update, so those details remain attributed reporting rather than independently inspected primary text.

Inference: the membership places intelligence, consumer protection, defense technology, and federal technical staffing inside one coordination forum. This is a reasonable reading of the named offices, not proof of the force’s legal authority, budget, jurisdiction, or final work program.

What remains unproved: no public executive order, memorandum, charter, complete membership roster, meeting record, budget, enforcement delegation, reporting protocol, or implementation evidence was located. The announcement does not by itself establish new statutory powers, binding obligations for companies, completed oversight, or measurable public benefit.

Related analysis

Primary presidential announcement naming the Super Intelligence Force and its leaders.

[Open the verified Truth Social post](#)<sup>28</sup>

Related analysis

Contemporaneous reporting on the reported 120-day assignment and risk-review charter, attributed to Jay Clayton and the Wall Street Journal.

[Open the Reuters report](#)<sup>29</sup>

# Original publication headings and credits

Original route titles, descriptions and publication metadata are retained here as a preservation record. The overview illustration is the cover image; the learning illustration remains beside its source context. The complete original overview is also available in the private source view, with its original order and text.

Permanent investigation / Artificial intelligence and human responsibility

Where Society Meets AI Super-intelligence

*Cover illustration. Original caption: AI connects domains; people and institutions retain the moral burden of deciding how it is used.*

Where Society Meets AI Super-intelligence / Complete web edition

The complete premise and research mandate

Opening and Investigation Directive

The full opening argument, philosophical standard, accountability problem and evidence questions that govern the twenty-five-section investigation.

*Project, Research, Analysis, Editorial by Patrick Russell; assisted by ChatGPT (OpenAI).*

Where Society Meets AI Super-intelligence / Complete web edition

Sections 1–4

Foundations, Uses and the Cross-Model Test

The technical and practical foundation: what AI is, how chat systems form answers, essential terminology, common misconceptions, major AI types, present-day uses, employment opportunity, testable questions and the full cross-model investigation.

*Project, Research, Analysis, Editorial by Patrick Russell; assisted by ChatGPT (OpenAI).*

Where Society Meets AI Super-intelligence / Complete web edition

Sections 5–10

Ownership, Economics and Human Impact

A deeper investigation of who owns and controls AI, the economics beyond employment, psychological effects and dependency, privacy and personal profiling, authenticity, evidence and the decisions that must remain human.

*Project, Research, Analysis, Editorial by Patrick Russell; assisted by ChatGPT (OpenAI).*

Where Society Meets AI Super-intelligence / Complete web edition

Sections 11–20

Ethics, Institutions, Threats and Protection

The governance record: the ethics of AI and its experts, competing schools of thought, U.S. primary-source priority, Congress and legislation, institutions directing AI, threats to beneficial systems, organized opposition, criminal use and proportionate protection.

*Project, Research, Analysis, Editorial by Patrick Russell; assisted by ChatGPT (OpenAI).*

Where Society Meets AI Super-intelligence / Complete web edition

Sections 21–25 and Conclusion

Truth, Promise and Responsible Human Direction

The final argument: how truth shapes society over time, how AI may strengthen or degrade that process, what its promise requires from individuals, parents, educators, businesses and government, and how responsible purposeful use should be measured.

*Project, Research, Analysis, Editorial by Patrick Russell; assisted by ChatGPT (OpenAI).*

Where Society Meets AI Super-intelligence / Education methodology

Evidence-informed learning framework

AI-Assisted Learning: A Practical Method for Parents and Educators

## REFERENCE

# Source-link register

Numbered destinations preserve the links embedded in the source record. They include primary sources, related analysis, downloads and full-size graphics. The numbers correspond to the linked markers in the text.

**001** Ai for Human Flourishing

<https://globeintelreport.com/investigations/ai-human-flourishing>

**002** Fear, Power & the Future of Ai

<https://globeintelreport.com/investigations/ai-fear-power-future>

**003** The Well-Formed Mind

<https://globeintelreport.com/investigations/education-well-formed-mind-ai>

**004** Original federal record · September 4, 2024

<https://www.justice.gov/usao-edpa/pr/justice-department-disrupts-covert-russian-government-sponsored-foreign-malign>

**005** Full evidence record in Election Integrity

<https://globeintelreport.com/investigations/election-integrity#doppelganger-public-record>

**006** Army and Navy case records

<https://globeintelreport.com/investigations/america-reloaded#defense-ai-october-2026>

**007** NSA agentic-AI guidance record

<https://globeintelreport.com/investigations/technology-defense-space-innovation#agentic-ai-security-evidence>

**008** Opening & Directive The complete premise, accountability question and research mandate

<https://globeintelreport.com/investigations/where-society-meets-ai/opening-and-directive>

**009** Sections 1–4 Foundations, AI types, practical uses and the eleven-tool model test

<https://globeintelreport.com/investigations/where-society-meets-ai/foundations>

**010** Sections 5–10 Ownership, economics, dependency, privacy, authenticity and human impact

<https://globeintelreport.com/investigations/where-society-meets-ai/human-impact>

**011** Sections 11–20 Ethics, experts, government, institutions, threats, opposition and criminal use

<https://globeintelreport.com/investigations/where-society-meets-ai/governance-and-threats>

**012** Sections 21–25 Truth over time, promise versus reality, human readiness and conclusion

<https://globeintelreport.com/investigations/where-society-meets-ai/truth-and-human-direction>

**013** Practical learning method Evidence, resistance, educator responsibility and the LEARN cycle for parents and schools

<https://globeintelreport.com/investigations/where-society-meets-ai/ai-assisted-learning>

**014** Read the complete analysis

<https://globeintelreport.com/investigations/where-society-meets-ai/human-impact#scientific-authority-uncertainty-and-contested-claims>

**015** 2023 MIT governance proposal by Daniel Huttenlocher, Asu Ozdaglar and David Goldston

<https://computing.mit.edu/wp-content/uploads/2023/11/AIPolicyBrief.pdf>

## Source-link register · continued

**016** Foundations of the Well Formed Mind The philosophical and developmental basis for knowledge, judgment, attention, character and responsibility.

<https://globeintelreport.com/investigations/education-well-formed-mind-ai/foundations>

**017** Artificial Intelligence and Learning Evidence separating immediate AI-assisted performance from durable, independently retained learning.

<https://globeintelreport.com/investigations/education-well-formed-mind-ai/ai-learning>

**018** The Synthetic You When personalized AI can write, answer and perform as a person—and why output alone cannot establish authorship.

<https://globeintelreport.com/investigations/education-well-formed-mind-ai/authorship>

**019** AI Regulations and Oversight An active tracker separating binding law, proposed measures, technical standards and company rules.

<https://globeintelreport.com/investigations/education-well-formed-mind-ai/regulations>

**020** Strategic Directive U.S. Leadership in Artificial Intelligence Connect AI capability and social consequences to the national objective, federal implementation, measured outcomes and unresolved risks. Open the U.S. leadership hub

<https://globeintelreport.com/?view=directives&section=strategy&directive=artificial-intelligence>

**021** Read the complete data-center investigation

<https://globeintelreport.com/investigations/data-centers-public-value>

**022** Related GIB Investigation America Reloaded Connect AI and autonomous systems to semiconductor capacity, trusted electronics, cyber capability and the industrial foundations of national power. Read the innovation analysis

<https://globeintelreport.com/investigations/america-reloaded#semiconductors-and-the-electronics-base>

**023** Related GIB Investigation Global AI Power & Governance Compare AI capability, law, infrastructure and strategic position across countries. Read the global comparison

<https://globeintelreport.com/investigations/global-ai-power-governance>

**024** Verified presidential post

<https://truthsocial.com/@realDonaldTrump/117406176604364910>

**025** September 29 executive order

<https://www.whitehouse.gov/presidential-actions/2026/09/inaugurating-the-era-of-super-intelligence/>

**026** White House technology page

<https://www.whitehouse.gov/priorities/tech-innovation/>

**027** Axios implementation report

<https://www.axios.com/2026/10/08/trump-ai-super-intelligence-white-house>

**028** Presidential announcement

<https://truthsocial.com/@realDonaldTrump/117382616783706638>

**029** Reuters report

<https://www.reuters.com/world/us/jay-clayton-lead-trumps-ai-task-force-deliver-report-120-days-wsj-reports-2026-10-03/>

**030** DNI leadership record

<https://www.dni.gov/>

**031** FTC leadership record

<https://www.ftc.gov/about-ftc/commissioners-staff/andrew-n-ferguson>

## Source-link register · continued

032 Department of War biography

<https://www.war.gov/About/Biographies/Biography/Article/4232659/emil-michael/>

033 OPM biography

<https://www.opm.gov/about-us/who-we-are/opm-director-scott-kupor/>

034 Signed White House Accord

<https://truthsocial.com/@realDonaldTrump/117356435739432952>

035 White House terminology fact sheet

<https://www.whitehouse.gov/fact-sheets/2026/09/fact-sheet-president-donald-j-trump-inaugurates-the-era-of-super-intelligence/>

036 Reuters accord report

<https://www.reuters.com/world/us/trump-releases-ai-accord-with-tech-executives-2026-09-29/>

037 CBS News remarks and context

<https://www.cbsnews.com/news/trump-ai-constitution-tech-execs-openai-anthropic-voluntary-controls/>

038 Luncheon seating chart

<https://truthsocial.com/@realDonaldTrump/117355137125519227>

039 America.gov executive order

<https://www.whitehouse.gov/presidential-actions/2026/09/streamlining-access-to-government-services-through-america-gov/>

040 Read the complete individual standard

<https://globeintelreport.com/investigations/where-society-meets-ai/truth-and-human-direction#individuals>

041 Read the complete parent standard

<https://globeintelreport.com/investigations/where-society-meets-ai/truth-and-human-direction#parents>

042 Read the complete educator standard

<https://globeintelreport.com/investigations/where-society-meets-ai/truth-and-human-direction#educators>

043 Read the complete business standard

<https://globeintelreport.com/investigations/where-society-meets-ai/truth-and-human-direction#businesses>

044 Read the complete government standard

<https://globeintelreport.com/investigations/where-society-meets-ai/truth-and-human-direction#government>

045 Read the supporting governance standard

<https://globeintelreport.com/investigations/where-society-meets-ai/governance-and-threats#20-protecting-people-and-systems>

046 51 pages Download the Full Illustrated Edition Word document · complete Phase One investigation

<https://globeintelreport.com/downloads/where-society-meets-ai-full-illustrated-edition.docx>

047 18 pages Download the Visual Concise Edition Word document · charts and condensed findings

<https://globeintelreport.com/downloads/where-society-meets-ai-visual-concise-edition.docx>

048 President Donald J. Trump — Truth Social statement on AI and data centers Primary presidential statement opposing additional AI guardrails, invoking existing enforcement authority and framing U.S.–China competition.

<https://truthsocial.com/@realDonaldTrump/117269745153543631>

## Source-link register · continued

**049** President Donald J. Trump — follow-up Truth Social statement on AI risk and infrastructure Primary presidential statement rejecting catastrophic-AI warnings, promoting rapid infrastructure growth and attributing Google's Finland investment to U.S. permitting difficulty.

<https://truthsocial.com/@realDonaldTrump/117270591511950591>

**050** President Donald J. Trump — Truth Social statement on data-center opponents Primary presidential statement dismissing neighborhood objections to data centers and alleging that unnamed opponents serve interests contrary to the United States.

<https://truthsocial.com/@realDonaldTrump/117271285418407156>

**051** President Donald J. Trump — Truth Social statement announcing an 'AI Force' Primary presidential statement promising an AI Force and future AI Czar while asserting broad economic and community benefits.

<https://truthsocial.com/@realDonaldTrump/117298821562014906>

**052** President Donald J. Trump — Truth Social endorsement of an 'AI Disaster Hoax' interview Primary presidential statement amplifying Matthew Continetti's argument that catastrophic-AI fears should not displace competition with China.

<https://truthsocial.com/@realDonaldTrump/117301659812842043>

**053** President Donald J. Trump — Truth Social statement encouraging 'super intelligence' Primary presidential statement rejecting broad catastrophic-AI warnings, favoring development of AI or 'super intelligence' and invoking Justice Department and law-enforcement authority if intervention becomes necessary.

<https://truthsocial.com/@realDonaldTrump/117309814170246776>

**054** President Donald J. Trump — Truth Social statement on Xi Jinping and 'super intelligence' Primary presidential statement characterizing President Xi's reaction to the new terminology after their September 24 meeting; the characterization is not presented as a direct quotation from Xi.

<https://truthsocial.com/@realDonaldTrump/117331749278785253>

**055** President Donald J. Trump — Truth Social statement excluding Scott Bessent as SI Czar Primary presidential statement saying Treasury Secretary Scott Bessent will remain at Treasury and will not become the proposed Super Intelligence Czar.

<https://truthsocial.com/@realDonaldTrump/117332704655075384>

**056** President Donald J. Trump — Truth Social repost of Fox News AI commentary Verified September 28 repost, without accompanying presidential text, of a Fox News segment discussing U.S.–China AI competition, economic and military transformation, and AI risk. The segment's statements are not direct presidential quotations.

<https://truthsocial.com/@realDonaldTrump/117348369252080500>

**057** Reuters — White House AI meeting and President Trump's same-day China remarks Contemporaneous reporting identifying the anticipated innovation-and-oversight discussion, the President's opposition to joint U.S.–China AI governance and the White House's lack of a public meeting agenda before the luncheon.

<https://www.reuters.com/legal/government/trump-host-zuckerberg-anthropics-amodei-other-ai-titans-tuesday-2026-09-29/>

**058** Axios — White House AI luncheon attendee list Contemporaneous reporting corroborating the 31-person attendee list while characterizing the administration's approach; the characterization is treated as reporting, not as an official meeting outcome.

<https://www.axios.com/2026/09/29/trump-ai-meeting-list-ceos-johnson>

**059** White House — President Trump at the United Nations Official September 22 record describing 'Super Intelligence' as the administration's new terminology and rejecting global control of the technology.

<https://www.whitehouse.gov/releases/2026/09/president-trump-at-the-united-nations-while-others-have-talked-i-have-acted/>

## Source-link register · continued

**060** Chinese Ministry of Foreign Affairs — Xi-Trump meeting readout Official Chinese account confirming substantive AI discussion on competition, cooperation, risks, benefits, misuse prevention and human control; it does not report Xi endorsing the 'Super Intelligence' label.

[https://www.mfa.gov.cn/eng/xw/zyxw/202609/t20260925\\_12031181.html](https://www.mfa.gov.cn/eng/xw/zyxw/202609/t20260925_12031181.html)

**061** Associated Press — State Department implementation of the terminology change Contemporaneous reporting based on an internal State Department email directing one bureau to use 'Super Intelligence' and 'SI' in documents and communications; government-wide scope remained unclear.

<https://apnews.com/article/trump-super-intelligence-rebrand-ai-diplomats-02123c7bf18c692a12b2b3e3b9ce30c5>

**062** U.S. Treasury — Secretary Scott Bessent Official record confirming Bessent's continuing role as the 79th Secretary of the Treasury.

<https://home.treasury.gov/about/general-information/officials/scott-bessent>

**063** Council of Economic Advisers — 2026 Economic Report of the President Official report using the conventional technical meaning of artificial superintelligence: AI whose intelligence surpasses that of humans.

<https://www.whitehouse.gov/wp-content/uploads/2026/04/2026-Economic-Report-of-the-President.pdf>

**064** White House — Promoting Advanced Artificial Intelligence Innovation and Security June 2, 2026 executive order directing prioritized enforcement of specified federal criminal statutes against AI-enabled computer crimes while establishing a voluntary frontier-model security framework without mandatory model licensing.

<https://www.whitehouse.gov/presidential-actions/2026/06/promoting-advanced-artificial-intelligence-innovation-and-security/>

**065** Stanford Institute for Human-Centered Artificial Intelligence — 2026 AI Index Metric-specific comparison showing that the U.S.-China model-performance gap had effectively closed while the United States retained advantages in top-tier models, high-impact patents and private investment and China led other research, patent and industrial-robot measures.

<https://hai.stanford.edu/ai-index/2026-ai-index-report>

**066** American Enterprise Institute — AI Competition with China Continues to Heat Up Preferred-reference commentary recording a seven-month estimate from America First Policy Institute testimony while also citing Stanford evidence that China had nearly erased the broad model-performance gap; it illustrates why any single lead estimate needs a named metric and date.

<https://ctse.aei.org/ai-competition-with-china-continues-to-heat-up/>

**067** White House — National Security Memorandum on AI leadership April 23, 2026 primary policy record treating AI leadership as an economic and national-security objective and identifying China-based efforts to obtain U.S. frontier-model capabilities.

<https://www.whitehouse.gov/wp-content/uploads/2026/04/NSTM-4.pdf>

**068** Associated Press — U.S. proposes AI incident notification mechanism with China Contemporaneous reporting that the administration is pursuing cross-border notification for AI incidents affecting national security, demonstrating a parallel risk-management track.

<https://apnews.com/article/bessent-ai-xi-trump-china-trade-2c7f54f07e755f506d9db9b91df282bd>

**069** Fox News — Matthew Continetti on AI risk and China competition The interview endorsed by President Trump; it argues that Chinese control of advanced AI could be used for malevolent ends.

<https://www.foxnews.com/video/6405314437112>

**070** Google — €13 billion Finland AI-infrastructure announcement Company record of the investment, locations, jobs and energy partnerships; it does not attribute the decision to U.S. permitting.

<https://blog.google/innovation-and-ai/infrastructure-and-cloud/global-network/google-ai-commitment-to-finland/>

**071** International Department of the CPC Central Committee — Xi Jinping on AI development and governance Official Chinese record calling for development while also emphasizing risk awareness, human control and secure, controllable AI.

[https://www.idcpc.org.cn/english2023/ttxw\\_5749/202607/t20260717\\_169913.html](https://www.idcpc.org.cn/english2023/ttxw_5749/202607/t20260717_169913.html)

## Source-link register · continued

072 Virginia JLARC — Data Centers in Virginia State audit documenting both economic benefits and project-level effects involving power costs, water, land use, noise and nearby residents.

<https://jlarc.virginia.gov/landing-2024-data-centers-in-virginia.asp>

073 White House Office of Science and Technology Policy Federal science, technology and AI policy direction.

<https://www.whitehouse.gov/ostp/>

074 AI.gov and America's AI Action Plan National AI policy, implementation and federal coordination.

<https://www.ai.gov/>

075 National AI Initiative Office Federal AI research and development coordination.

<https://www.nitrd.gov/coordination-areas/ai/>

076 NIST AI Risk Management Framework Voluntary risk-management functions, terminology and implementation resources.

<https://www.nist.gov/itl/ai-risk-management-framework>

077 GAO AI Accountability Framework Governance, data, performance and monitoring questions for accountable systems.

<https://www.gao.gov/products/gao-21-519sp>

078 GAO Review of Federal Generative-AI Management Audited evidence on federal governance and implementation gaps.

<https://www.gao.gov/products/gao-26-107859>

079 CISA Artificial Intelligence Cybersecurity and critical-infrastructure guidance.

<https://www.cisa.gov/ai>

080 NSA Artificial Intelligence Security Center AI-system security and national-security risk.

<https://www.nsa.gov/aisc/>

081 Federal Trade Commission AI Partnerships Study Competition and major cloud-AI partnerships.

<https://www.ftc.gov/news-events/news/press-releases/2025/01/ftc-issues-staff-report-ai-partnerships-investments-study>

082 Department of Energy Data-Center Electricity Report AI infrastructure and electricity-demand evidence.

<https://www.energy.gov/articles/doe-releases-new-report-evaluating-increase-electricity-demand-data-centers>

083 European Commission AI regulatory framework European legal and regulatory context.

<https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>

084 OECD AI Principles International policy principles and comparative context.

<https://www.oecd.org/en/topics/sub-issues/ai-principles.html>

085 International AI Safety Report 2026 International review of general-purpose AI capabilities and risks.

<https://internationalaisafetyreport.org/publication/international-ai-safety-report-2026>

086 UNESCO Recommendation on the Ethics of AI International ethics framework—considered as normative, not dispositive authority.

<https://www.unesco.org/en/artificial-intelligence/recommendation-ethics>

087 SAG-AFTRA AI Resources Labor, consent, compensation and likeness protections.

<https://www.sagaftra.org/contracts-industry-resources/member-resources/artificial-intelligence>

088 Authors Guild AI Advocacy Copyright, licensing and author-rights arguments.

<https://authorsguild.org/advocacy/artificial-intelligence/>

## Source-link register · continued

**089** AFL-CIO Workers First Initiative on AI Worker bargaining, surveillance and workplace deployment concerns.  
<https://aflcio.org/reports/workers-first-ai>

**090** PauseAI Statement Organized advocacy for a pause—treated as an interested policy source.  
<https://pauseai.info/statement>

**091** Future of Life Institute Open Letter Prominent pause proposal and signatory record.  
<https://futureoflife.org/open-letter/pause-giant-ai-experiments/>

**092** Aristotle, Nicomachean Ethics Philosophical grounding for human moral responsibility.  
<https://classics.mit.edu/Aristotle/nicomachaen.html>

**093** Read Anthropic's September 2026 report  
<https://www.anthropic.com/threat-intelligence-report-september-2026>

**094** Read The Epoch Times report  
<https://www.theepochtimes.com/us/iran-used-claude-ai-model-to-target-us-navy-warships-anthropic-report-6086268>

**095** Read the CSIS analysis  
<https://www.csis.org/analysis/ai-security-threats-arent-hoax-they-also-arent-existential-they-might-seem>

**096** Read the study  
<https://www.nature.com/articles/s41598-025-97652-6>

**097** Read the study  
<https://www.pnas.org/doi/10.1073/pnas.2422633122>

**098** Read the study  
<https://arxiv.org/abs/2410.03017>

**099** Read the study  
<https://arxiv.org/abs/2607.08849>

**100** Review the organization's source  
<https://www.nea.org/about-nea/media-center/press-releases/nea-delegates-approve-measure-recommend-policy-actions-use-ai>

**101** Review the organization's source  
<https://www.aft.org/reports/commonsense-guardrails-using-advanced-technology-schools>

**102** Review the organization's source  
<https://www.unesco.org/en/articles/guidance-generative-ai-education-and-research>

**103** Review the organization's source  
<https://institute.common sense media.org/>

**104** Review the organization's source  
<https://www.eff.org/issues/student-privacy>

**105** Review the organization's source  
<https://epic.org/issue-types/student-privacy/>

**106** Review the organization's source  
<https://cdt.org/insights/advancing-responsible-ai-adoption-and-use-in-k-12-education-three-policy-priorities-for-state-legislation/>

## Source-link register · continued

**107** Review the organization's source

<https://www.ed.gov/about/news/press-release/us-department-of-education-issues-guidance-artificial-intelligence-use-schools-pr-opposes-additional-supplemental-priority>

**108** Review the organization's source

<https://www.ftc.gov/legal-library/browse/rules/childrens-online-privacy-protection-rule-coppa>

**109** Review the organization's source

[https://www.rand.org/pubs/research\\_reports/RRA4180-1.html](https://www.rand.org/pubs/research_reports/RRA4180-1.html)

**110** UNESCO guidance on generative AI in education

<https://unesdoc.unesco.org/ark:/48223/pf0000386693>

**111** U.S. Department of Education: AI and the future of teaching and learning

<https://www.ed.gov/sites/ed/files/documents/ai-report/ai-report.pdf>