XOTAVO · AI POLICY HANDBOOK

What artificial intelligence is, in current usage

Provided by ohmyai.ai

A plain-language handbook for public officials, staff, and organizations. Products are named only in Appendix B, as a dated reference list. Positions are attributed to their sources.

Version 16.1
Last updated: November 3, 2026

IN ONE MINUTE

- In everyday government and vendor language, “AI” most often means generative systems—especially large language models—that predict tokens and produce text, images, audio, code, or video.
- Training builds a model from data; inference is everyday use. Agents add tool use and multi-step actions on top of a model.
- Documented limits include fabricated facts (“hallucinations”), training cut-offs, and measured performance gaps across groups and tasks.
- As of Nov 3, 2026, U.S. rules are mostly state and sector-specific; the EU AI Act is phasing in by risk category; several countries use promotion laws, voluntary codes, or pending bills.
- Public-sector decisions typically turn on data location, training-on-inputs terms, citation/retrieval behavior, records retention, and human sign-off—not on brand names.

WHAT PEOPLE USUALLY MEAN BY “AI”

When officials, vendors, and news coverage say “AI” in 2025–2026, they usually mean generative AI: systems that create new content rather than only classifying existing records. The most common office tools are built on large language models (LLMs). An LLM predicts the next token of text given prior tokens. Stacked at scale, that prediction can draft memos, summarize transcripts, translate documents, or write code. Related systems generate images, audio, or video. Computer vision, speech recognition, and classical machine-learning scoring systems remain in use for inspections, licensing queues, and fraud detection; they are AI under broader definitions, but they are not what most “AI assistant” procurements refer to.

Researchers and companies also use the labels AGI (artificial general intelligence) and ASI (artificial superintelligence). AGI usually means a system that could match or exceed human performance across most cognitive tasks. ASI usually means a system that substantially exceeds human ability across nearly all domains. As of Nov 2026, those terms describe research goals and policy debates, not a settled product category with shared technical criteria. This primer keeps the everyday name “AI” for deployed generative and predictive systems, and treats AGI/ASI as forward-looking vocabulary that appears in hearings and vendor materials.

Operational description used throughout this primer: a model maps inputs to outputs using patterns learned from training data and any tools connected at inference time. Predicting fluent language is not the same as verifying facts.

MODELS, DATA, COMPUTE, AND AGENTS

Model

The trained artifact: billions of parameters (weights) that encode statistical patterns. Foundation models are trained broadly; they can be adapted with fine-tuning or prompting.

Training data

Text, images, code, and other corpora used to set the weights. Composition affects capability, bias measurements, copyright disputes, and privacy risk.

Inference

Running the trained model on a new prompt. Most day-to-day "using AI" is inference, billed by tokens, seats, or API calls.

Compute

GPUs, TPUs, and related accelerators used for training and inference. Training frontier-scale models requires large clusters; inference cost scales with usage.

Open vs closed weights

Open-weight releases let an organization download and run a model under a license. Closed models are reached only as a service. Licenses and support differ from classic open-source software.

Agents

Systems that call tools, browse, write files, or take multi-step actions. They raise permission, logging, and oversight questions beyond a chat box.

CAPABILITIES THAT DRIVE ADOPTION

Documented strengths of current generative systems, as described in vendor model cards and independent benchmarks, include drafting and revision of prose, translation, summarization of long text, code assistance, and question-answering when retrieval is available. Computer-vision systems continue to support document intake, license-plate or form extraction, and inspection workflows that predate generative AI. Speech-to-text systems feed those pipelines. The policy question for a legislature or city is rarely whether any of these tasks can be automated in part; it is which tasks may be automated under what disclosure, records, procurement, and human-review rules.

WHY THE TOPIC REACHED GOVERNING BODIES

Public bodies report three concurrent shifts: model capability rose; consumer interfaces removed the need for programming skill; and cloud pricing made pilots inexpensive to start. That combination moved AI questions from IT departments into procurement, records, workforce, elections, child-safety, and infrastructure dockets. Parallel debates over data-center load, training-data copyright, and synthetic media appear in statehouses and national legislatures at the same time (as of 2024–2026 legislative trackers summarized later in this primer).

XOTAVO · AI POLICY HANDBOOK

How modern systems work

Provided by ohmyai.ai

Training, post-training, inference, retrieval, and agents—without the mathematics, with the failure modes each step creates.

Version 16.1
Last updated: November 3, 2026

FROM DATA TO A DEPLOYED ASSISTANT

1. **Pre-training on large corpora.** Developers train a neural network—today usually a transformer—on large collections of text and, for multimodal models, images or other media. The objective is next-token (or similar) prediction. The model is not given a verified database of facts; facts appear only insofar as they were consistent in the data.

2. **Post-training (including RLHF and related methods).** Developers further adjust the model with supervised examples and preference feedback so it follows instructions, refuses certain requests, and matches a product's tone. Reinforcement learning from human feedback (RLHF) is one common family of methods. Post-training shapes helpfulness and refusal behavior; it does not install a truth oracle.

3. **Inference.** At use time the model receives a system prompt and user prompt, then samples tokens. Settings such as temperature change how varied outputs are. There is no separate reasoning engine that checks a ledger unless the product adds tools.

4. **Optional retrieval and tools (RAG, search, code execution).** Many products first retrieve documents from the web or a private corpus, then condition generation on those passages (retrieval-augmented generation, RAG). Agents may call calendars, browsers, or internal APIs. Quality then depends on retrieval coverage and tool permissions as much as on the base model.

5. **Logging, evaluation, and monitoring.** Enterprises and governments that deploy AI typically keep inventories, run evaluations (benchmarks, red teams, bias tests), and monitor incidents. NIST's AI RMF frames these activities as Govern, Map, Measure, and Manage (as of the Jan 2023 AI RMF 1.0 and the Jul 2024 Generative AI Profile).[25,26]

SCALE, TOKENS, AND COST VOCABULARY

Model size is often discussed in parameters (billions of weights). Training cost is discussed in FLOPs or GPU-hours. Inference cost is usually metered in tokens: pieces of text the model reads and writes. A long PDF stuffed into the prompt consumes context window and budget even before the model answers. Multimodal inputs (images, audio) are likewise converted into tokens or embeddings. Officials comparing vendor quotes are

usually comparing seat licenses, API token rates, minimum commitments, and whether retrieval or agent tool-calls are billed separately.

Post-training alignment methods (including RLHF and related preference optimization) change how models refuse or comply; security teams still test for jailbreaks because those controls are imperfect. When products add browsing or code-execution tools, evaluations must cover tool misuse, not only chat refusals.

OPEN-WEIGHT AND CLOSED DEPLOYMENT

Closed service. The vendor hosts the model. Buyers gain ease of use and frequent updates. Data-processing location, subprocessors, training-on-inputs terms, and export of logs become contract questions.

Open weights on own hardware or a chosen cloud. The organization can keep prompts inside a chosen boundary and inspect configurations. It also accepts patching, access control, evaluation, and incident response as its own duties. “Open weight” is not identical to “open source”; licenses vary.

Hybrid patterns. Some agencies use a closed model for general drafting and a local or virtual-private-cloud model for records that cannot leave a designated environment. Suite-embedded assistants inherit the identity and permission model of the office suite.

Cost drivers. Training frontier models is capital-intensive. Inference cost scales with tokens, context length, tool calls, and concurrency. Self-hosting shifts spend from API invoices to hardware, power, and staff time.

WHAT EACH DESIGN CHOICE IMPLIES FOR GOVERNMENT

DESIGN CHOICE	TYPICAL IMPLICATION FOR A PUBLIC BODY
Base model only (no retrieval)	Answers reflect training patterns and cut-off; harder to cite primary sources.
RAG over approved corpora	Answers can point to policies you control; errors track corpus gaps and retrieval
Web-grounded answer engine	Better for current events if citations are shown; still requires human verification
Agent with write permissions	Can change systems of record; needs least-privilege access, logging, and kill swit
Consumer free tier	Often weaker contractual limits on training-on-inputs and retention than enterp

A recurring audit question: which model version, which tools, which corpus, and which human approved the output that left the agency?

XOTAVO · AI POLICY HANDBOOK

Families of tools and the differences that matter

Provided by ohmyai.ai

Brand names change. Family, data path, and jurisdiction usually decide the governance problem.

Version 16.1
Last updated: November 3, 2026

SIX FAMILIES PUBLIC BODIES ENCOUNTER

General assistants

Open-ended chat over a foundation model. Broad drafting and Q&A; grounding varies by product mode.

Answer engines

Retrieve live sources first, then generate with citations. Built for lookup; citations still need checking.

Suite-embedded assistants

Inside email, documents, and calendars, using existing identity permissions. Touches organizational content by design.

Real-time / social tools

Wired to live feeds. Useful for situational awareness; weak as an authoritative legal or scientific source.

Task agents

Plan and execute multi-step actions with tools. Highest need for permission boundaries and audit logs.

Self-hosted / private instances

Open-weight or dedicated instances inside a chosen boundary. More operational work; stronger data-control options.

FEATURES THAT CHANGE RISK, INDEPENDENT OF BRAND

Where processing and storage occur

Some tools process prompts only on provider infrastructure in specified regions; others can run on-device or in a customer virtual private cloud. Location determines which governments can compel access under local law. On Jan 31, 2025, the Office of the Texas Governor announced a ban on several foreign-developed apps, including an AI assistant, on state-issued devices and networks.[3]

Whether sources are shown

Citation and RAG features make verification possible. Their presence does not guarantee accuracy; their absence means staff cannot trace a claim to a document without separate research.

Whether inputs are used for training

Consumer terms often allow broader use of prompts for model improvement than enterprise or government terms. Contract language for the specific SKU, not marketing pages alone, is what procurement reviews.

Open-weight versus closed service

Open weights enable self-hosting and offline evaluation. Closed services centralize updates and liability allocation in the vendor relationship. Neither pattern is inherently compliant; controls and contracts decide.

PROCUREMENT CHECKPOINT

Data location and compelling authority

A recurring written question from public buyers: in which country are prompts, files, and embeddings stored and processed, and which legal processes can compel the provider to disclose them? Quality and price do not answer that question.

FEATURE	QUESTIONS COMMONLY RECORDED IN PUBLIC PROCUREMENTS
Processing location	Region options? Customer-managed keys? On-device mode?
Citations / RAG	Can staff see retrieved passages? Can the corpus be limited to approved sources?
Training on customer data	Default off for this SKU? Subprocessors? Written residual rights?
Retention & deletion	Prompt/log retention period? Deletion on request and at contract end?
Agent permissions	Which tools can it call? Who approves expansions?

XOTAVO · AI POLICY HANDBOOK

Documented limits, evaluation, and measurement

Provided by ohmyai.ai

Failures that appear repeatedly in research, vendor model cards, and incident reports—and how organizations measure them.

Version 16.1
Last updated: November 3, 2026

RECURRING FAILURE MODES

Hallucination (fabricated content)

Models can produce citations, quotations, case names, or statistics that do not exist, delivered in the same fluent style as accurate text. The mechanism is next-token prediction, not a verified knowledge lookup. NIST's Generative AI Profile treats confabulation and related reliability issues as risks to manage with measurement and oversight.[26]

Training cut-off and stale knowledge

A base model has no built-in knowledge after its training cut-off unless the product adds retrieval or tools. Without those, answers about recent events may be incomplete or invented.

Bias and disparate performance

Evaluations often show different error rates across demographic groups, languages, or dialects. Measurement methods and thresholds are themselves contested; reporting the metric and the test set matters as much as a single summary score.

Brittle reasoning and arithmetic

Models can produce step-by-step text that looks like reasoning while failing simple logic or math. Tool-augmented systems that call calculators or code interpreters change that profile when those tools are actually used.

Sycophancy and premise acceptance

Post-training for helpfulness can increase agreement with a user's premise even when the premise is false. Explicit instructions to challenge assumptions change behavior but are not automatic.

Security-specific failures

Prompt injection, data leakage through prompts, agent over-permissioning, and model-extraction attempts appear in security guidance from agencies such as CISA and in vendor secure-use documentation. Agents amplify impact because they can act, not only speak.

HOW LIMITS ARE EVALUATED

BENCHMARKS	Shared task suites (language, code, math, multimodal). Useful for comparison; limited as predictors of a local workflow. Scores need an as-of date and model version.
RED TEAMING	Adversarial testing for disallowed content, jailbreaks, and misuse. Can be internal, contracted, or community-based.
BIAS AND QUALITY SLICES	Error rates broken out by language, geography, or protected-class proxies where lawful and relevant to the use case.

OPERATIONAL METRICS	Human edit distance, escalation rate, citation click-through, and incident counts after deployment.

NIST AI RMF 1.0 (Jan 2023) and the Generative AI Profile (Jul 2024) provide voluntary U.S. vocabulary for mapping these activities without mandating a single test.[25,26] ISO/IEC 42001:2023 specifies requirements for an AI management system that organizations may certify against.[27]

When legislators hear “the model scored X%,” useful follow-ups include: which benchmark, which model version, which date, whether tools/retrieval were enabled, and whether the test matches the local workflow (for example, answering residents in Spanish about trash collection rules). A high general benchmark score does not by itself answer those questions.

Accountability for a public statement or decision remains with the office that issues it, regardless of whether a model drafted the text—an allocation reflected in many agency AI policies and in OMB memoranda for federal agencies.[19]

XOTAVO · AI POLICY HANDBOOK

Energy, water, and compute infrastructure

Provided by ohmyai.ai

Figures reported by laboratories, agencies, and operators. Facility design varies; totals and per-task efficiency are separate measurements.

Version 16.1
Last updated: November 3, 2026

WHY DATA CENTERS APPEAR ON LOCAL AGENDAS

Training and inference for large models concentrate electricity demand in data centers. Cooling historically used substantial water in many designs. Proposals for new facilities therefore intersect zoning, ratepayer, water, and emergency-curtailment policy.

PUBLISHED DEMAND FIGURES

U.S. LOAD

Lawrence Berkeley National Laboratory estimated U.S. data centers used about 4.4% of U.S. electricity in 2023 and could use about 6.7% to 12% by 2028 (report as of Dec 2024). The same report estimated about 66 billion liters (about 17 billion gallons) of direct water use in 2023, largely for cooling. [4]

GLOBAL LOAD

The International Energy Agency estimated data centers used about 415 TWh in 2024 (about 1.5% of global electricity) and projected about 945 TWh by 2030 (as of Apr 2025). [8]

PUE

Uptime Institute's Global Data Center Survey 2025 reported an average power usage effectiveness (PUE) of 1.54 industry-wide, with newer sites averaging 1.48 (as of 2025). [6]

DESIGN CHANGES OPERATORS DESCRIBE

Liquid cooling

Coolant delivered near chips rather than only room-scale air conditioning.

Closed-loop water

Microsoft stated that its datacenter designs since August 2024 evaporate no water for cooling (company blog, Dec 9, 2024). [5]

Heat reuse

Meta's Odense fact sheet (Feb 2025) describes heat supply serving about 9,000 homes. [7]

Per-query efficiency can improve while aggregate electricity use still rises if total inference and training volume grow faster. Published totals and published per-task figures answer different questions.

QUESTIONS RECORDED WHEN A LARGE LOAD IS PROPOSED

POWER	Source mix, on-site generation, interconnection costs, and ratepayer allocation.
WATER	Annual withdrawal, closed-loop share, and drought contingency.
PUE / DESIGN	Design PUE, cooling method, and heat-reuse plans.
TEXAS LARGE LOADS	SB 6 (2025) addresses interconnection, cost-sharing, and emergency curtailment for large loads of 75 MW or more; the Public Utility Commission adopted 16 TAC §25.370 on Feb 6, 2026.[9,10]
BACKUP / LAND USE	Generator fuel, test schedules, noise, setbacks, lighting, and traffic.

Related background on cooling history: The data center comes full circle.

XOTAVO · AI POLICY HANDBOOK

Federal U.S. and EU rules already in motion

Provided by ohmyai.ai

Neutral status summary. Not legal advice. Confirm details with counsel. Broader state and country tables are in the next briefing.

Version 16.1
Last updated: November 3, 2026

TEXAS (SELECTED)

TRAIGA	HB 149, effective Jan 1, 2026, requires government agencies that offer AI systems interacting with residents to disclose that interaction in plain language, and restricts specified uses such as social scoring and certain discriminatory uses; the Attorney General enforces.[1,11]
AI TRAINING	HB 3512, effective Sep 1, 2025, requires annual DIR-certified AI training for covered local and state officials and employees who use computers for at least 25% of duties.[2,12]
CLOUD	TX-RAMP certification is required before a state agency contracts for covered cloud services.[13]
DEEPFAKES / CHILD SAFETY	Election Code §255.004 addresses deceptive deep fake video near elections; SB 20 (effective Sep 1, 2025) extends obscenity law to certain AI-generated child depictions.[14,15]

FEDERAL LANDSCAPE IN BRIEF

As of Nov 3, 2026, the United States has no single comprehensive federal AI Act analogous to the EU AI Act. Federal activity is a mix of sector enforcement (for example FTC actions and TAKE IT DOWN), procurement and use rules for federal agencies (OMB memoranda), voluntary standards (NIST), and executive orders that direct agencies. States have filled much of the detailed legislative space on frontier transparency, employment AI, election deepfakes, and chatbot disclosures.

FEDERAL (SELECTED INSTRUMENTS)

TAKE IT DOWN	Public Law 119-12, signed May 19, 2025, criminalizes publication of non-consensual intimate images including certain AI “digital forgeries,” and requires covered platforms to remove validly reported images within 48 hours beginning May 19, 2026, with FTC enforcement.[16,17]

EO 14365	Issued Dec 11, 2025, directs federal agencies regarding state AI laws the administration describes as burdensome. An executive order does not by itself repeal state statutes; the order text preserves space for state child-safety and state-government procurement rules.[18]
OMB M-25-21 / M-25-22	Apr 3, 2025 memoranda set federal-agency governance and acquisition expectations for AI, including inventories and contract practices aimed at reducing lock-in. They bind federal agencies; some localities reference them as models.[19,20]

EUROPEAN UNION AI ACT TIMELINE

The EU AI Act entered into force Aug 1, 2024. Prohibitions and AI-literacy duties applied from Feb 2, 2025. GPAI model obligations applied from Aug 2, 2025. Most remaining rules, including Article 50 transparency duties, applied from Aug 2, 2026. The Digital Omnibus on AI adjustments set Annex III high-risk rules for Dec 2, 2027 and Annex I product-embedded high-risk rules for Aug 2, 2028; new prohibitions on AI that generates non-consensual intimate images or CSAM apply from Dec 2, 2026 (European Commission AI Act Service Desk timeline, accessed for this review Nov 3, 2026).[21]

Status figures in this briefing are as of Nov 3, 2026 unless a source date is stated. Laws change; verify before reliance.

XOTAVO · AI POLICY HANDBOOK

What states and other countries are doing

Provided by ohmyai.ai

Selected enacted laws and notable measures, with status and effective dates as reported by primary or legislative trackers. Snapshot as of Nov 3, 2026. Not exhaustive and not legal advice.

Version 16.1
Last updated: November 3, 2026

HOW TO READ THIS SNAPSHOT

U.S. AI rules are mostly state-driven and topic-specific (employment, elections, frontier-model transparency, chatbots, NCII). Comprehensive federal AI legislation has not been enacted as of Nov 3, 2026. Executive Order 14365 (Dec 11, 2025) directs federal agencies concerning state AI laws the administration describes as burdensome, while stating that a future national standard should not override state child-safety rules or state government purchasing of AI.[18] Abroad, approaches range from binding risk-based statutes (EU) to promotion frameworks, voluntary codes, and pending bills. Court challenges have enjoined or narrowed some state election-deepfake statutes; NCSL summarizes those cases alongside the enacted count.[28]

Pending bills continue to move in many legislatures. Examples reported in 2025–2026 trackers include additional election-synthetic-media disclosure bills, chatbot safety proposals, and expansions of frontier-model reporting (for example, New York follow-on proposals building on the RAISE Act). Status changes quickly; the tables below prioritize enacted measures with dated effective dates, then flag representative international instruments.

SELECTED U.S. STATE MEASURES (COMPREHENSIVE / SECTOR)

STATE	LAW / BILL	STATUS (AS OF NOV 3, 2026)	EFFECTIVE /
Texas	HB 149 (TRAIGA)	Enacted	Jan 1, 2026
Texas	HB 3512	Enacted	Sep 1, 2025
California	SB 53 (Transparency in Frontier AI Act)	Enacted; in force	Jan 1, 2026
California	AB 2013 (Training Data Transparency Act)	Enacted; in force	Jan 1, 2026
Colorado	SB 24-205 (original CAIA)	Repealed/replaced before enforcement	Never opera

Colorado	SB 26-189 (ADMT Act)	Enacted; rulemaking underway	Jan 1, 2027
Colorado	HB 26-1263 (Chatbot Safety Act)	Enacted	Jan 1, 2027
Utah	AI Policy Act + SB 226 / HB 452 amendments	In force (Office sunset Jul 1, 2027)	Disclosure r
New York	RAISE Act (as amended Mar 27, 2026)	Enacted	Jan 1, 2027
Illinois	HB 3773 (Human Rights Act AI amendment)	Enacted; in force	Jan 1, 2026
Illinois	SB 315 (AI Safety Measures Act)	Enacted (tracker)	Jan 1, 2027

ELECTION DEEPFAKES AND NCII (MULTI-STATE)

NCSL reports that, as of its Oct 13, 2026 update, 31 states have enacted laws regulating deepfakes in political messaging—mostly disclosure rules, with prohibition models in states including Texas, Minnesota, Maryland, and (as amended after litigation) California. [28] Federal TAKE IT DOWN (above) addresses non-consensual intimate imagery including certain AI forgeries. Many states also have separate intimate-image or digital-replica statutes; those are tracked in state codes and in federal-state comparisons rather than in the election table alone.

STATE (EXAMPLES)	MEASURE	STATUS NOTES (AS OF NOV 3, 2026)
Texas	Election Code §255.004	In force; Class A misdemeanor for sp
California	Election deepfake / AI ad rules (amended after *Kohls*)	Prior version enjoined 2025; legislat
Minnesota	Minn. Stat. §609.771	Enacted; federal challenges pending
Maryland	Md. Election Code §§2-110, 16-905	Enacted 2026; year-round deceptive
Multi-state	Various disclosure statutes	27 states use disclosure approaches
Federal	TAKE IT DOWN Act	In force; platform removal duty sinc

SELECTED INTERNATIONAL FRAMEWORKS

JURISDICTION	INSTRUMENT	STATUS (AS OF NOV
European Union	AI Act + Digital Omnibus adjustments	In force; phased ap
United Kingdom	Sector-regulator / pro-innovation approach; AI Opportunities Action Plan	No comprehensive
China	Interim Measures for Generative AI Services	In effect since Aug
South Korea	AI Basic Act (Framework Act)	In effect Jan 22, 20
Japan	AI Promotion Act	Fully in force Sep 1

Canada	AIDA (Bill C-27 Part 3)	Not enacted; bill di
Brazil	PL 2338/2023	Pending in Chamb
India	AI Governance Guidelines (Nov 2025); DPDP Act/Rules; MeitY advisories	Guidelines non-sta
Singapore	Model AI Governance Framework / AI Verify	Voluntary tools in a
Australia	Existing law + voluntary guidance; 2026 policy announcements	No comprehensive

MULTILATERAL EFFORTS

OECD	AI Principles (2019, updated) remain a non-binding intergovernmental reference for trustworthy AI.[44]
G7 HIROSHIMA	Hiroshima AI Process Code of Conduct (Oct 2023) is a voluntary code for advanced AI organizations, with OECD reporting support continuing into 2026.[45]
COUNCIL OF EUROPE	Framework Convention on AI, Human Rights, Democracy and the Rule of Law (2024) is the first global legally binding AI treaty open to non-members; entry into force depends on ratifications (EU ratified May 14, 2026 per reporting; treaty not yet in force as of mid-2026 summaries).[46]
UNESCO / UN	UNESCO's Recommendation on the Ethics of AI (2021) is a global soft-law instrument. UN processes continue to discuss international AI governance without a single binding UN AI statute as of Nov 2026.

NCSL and national legislatures update databases frequently. Re-check primary bill text before drafting local policy.

XOTAVO · AI POLICY HANDBOOK

Main policy topics officials encounter

Provided by ohmyai.ai

Issue areas that recur in hearings and briefings. Where groups disagree, positions are attributed.

Version 16.1
Last updated: November 3, 2026

PRIVACY	Prompts can contain PII, case notes, or confidential records. Consumer tools may retain logs; enterprise contracts often set retention, residency, and non-training terms. Sector laws (health, education, children's data, criminal justice) still apply when AI is only the interface. Cross-border transfers raise additional questions when a vendor's subprocessors sit outside the buyer's country. Privacy regulators in several jurisdictions have opened inquiries into scraping and model training; outcomes remain fact-specific.
COPYRIGHT	News publishers, authors, and other rightsholders have filed suits alleging unauthorized use of works in training corpora; developers argue fair use, implied license, or other defenses depending on the forum. As of Nov 3, 2026, major U.S. cases remain in litigation rather than settled into a single national statute. Separately, California's AB 2013 requires certain generative-AI developers to document training-data categories (effective Jan 1, 2026).[30] Outputs that closely reproduce protected works raise distinct infringement questions from the training-data dispute. Public bodies that publish AI-assisted text still need clearance practices for third-party materials they paste into prompts.
DEEPFAKES & ELECTIONS	States split between disclosure and prohibition models; some statutes face First Amendment challenges (NCSL, Oct 13, 2026).[28] Provenance standards such as C2PA Content Credentials record origin metadata; they do not by themselves detect unlabeled deepfakes.[47]
EMPLOYMENT	Illinois's 2026 Human Rights Act amendment requires notice and bars AI use that has a discriminatory effect in covered decisions.[34] Labor organizations argue for bargaining over monitoring and automation; some employers argue AI assists productivity if human review remains.
CRITICAL INFRASTRUCTURE	CISA and sector agencies publish guidance on AI-enabled cyber risk and on defending systems that themselves use AI. Failure modes include poisoned models, insecure agent permissions, and over-reliance on automated alerts.
GOVERNMENT PROCUREMENT & USE	OMB M-25-21 and M-25-22 set federal inventory, high-impact review, and acquisition practices.[19,20] Local examples include published AI inventories and staff guidelines (e.g., City of San José).[22,24]

CHILD SAFETY	Texas SB 20 and federal TAKE IT DOWN address AI-generated CSAM/NCII pathways.[15,16] Colorado's Chatbot Safety Act (effective Jan 1, 2027) adds operator duties for conversational systems used by minors.[31] Schools and libraries face separate acceptable-use and student-data questions.
LIABILITY	Product, professional, and governmental liability doctrines are being tested in courts and legislatures. Agency policies commonly state that the human decision-maker remains responsible for official acts that use AI assistance.[19]

XOTAVO · AI POLICY HANDBOOK

Standards bodies and questions for vendors

Provided by ohmyai.ai

Voluntary frameworks and the written questions that commonly appear in public-sector evaluations.

Version 16.1
Last updated: November 3, 2026

STANDARDS AND FRAMEWORKS FREQUENTLY CITED

NIST AI RMF 1.0

Jan 2023 voluntary U.S. framework: Govern, Map, Measure, Manage.[25]

NIST AI 600-1

Jul 2024 Generative AI Profile with GAI-specific risks and suggested actions.[26]

ISO/IEC 42001:2023

Certifiable AI management-system standard.[27]

C2PA / Content Credentials

Provenance standard for media origin and edits.[47]

OECD AI Principles

Intergovernmental soft law on trustworthy AI.[44]

TX-RAMP / FedRAMP analogues

Authorization programs for cloud services used by government.[13]

QUESTIONS OFFICIALS COMMONLY ASK VENDORS (IN WRITING)

DATA LOCATION	Where are prompts, embeddings, and files stored and processed? Which affiliates and subprocessors?
TRAINING	Are customer inputs used to train or improve models for others under this exact SKU?
RETENTION	Default log retention? Deletion on request and at contract end, with written confirmation?
GROUNDING	Does the product cite sources? Can retrieval be limited to an approved corpus?
EVALUATION	What benchmarks, red-team results, and bias tests are available for the deployed version?
ACCESS CONTROL	SSO, role-based access, customer-managed keys, and admin audit logs?

AGENTS	Which tools can the agent call by default? How are permissions changed and logged?
INCIDENTS	Notification timeline for security or safety incidents? Cooperation with public records requests?
EXIT	Export formats for prompts/outputs/knowledge bases? Transition assistance and fees?
COMPLIANCE MAPPING	How does the vendor map features to TRAIGA disclosure, EU AI Act roles, or NIST AI RMF functions if claimed?

CONTRACT TERMS OFTEN REQUESTED

- No training on customer content for the stated plan, in writing.
- Deletion and certification at contract end.
- Security-incident notice within a defined number of days.
- Right to test performance and bias on agreed datasets before and after go-live.
- Exit assistance and data return without punitive lock-in fees (parallel to themes in OMB M-25-22).[20]
- For Texas state agencies: TX-RAMP where applicable.[13]

XOTAVO · AI POLICY HANDBOOK

Use inside a public body

Provided by ohmyai.ai

Documented use cases, records issues, and examples cities have published. Descriptive, not a recommendation.

Version 16.1
Last updated: November 3, 2026

USES AGENCIES REPORT

Drafting
First drafts of routine notices and correspondence for human edit.

Summaries
Condensing long reports or meeting transcripts with originals retained.

Language access
Translation and plain-language aids reviewed by fluent staff for official use.

311 / FAQs
Answers drawn from published information, with escalation to a person.

Research assist
Orientation on a topic using citation-enabled tools, then verification.

Accessibility
Assistance for staff or residents navigating dense documents.

CONSTRAINTS THAT APPEAR IN WRITTEN POLICIES

Confidential resident data, personnel files, and non-public enforcement material are commonly barred from consumer-grade tools. Public-records and retention rules may treat prompts and outputs as records.

Disclosure to residents when they interact with AI is required for Texas government systems under TRAIGA.[1] Many policies require a named human approver before material becomes public.

PUBLISHED LOCAL EXAMPLES

INVENTORY	City of San José publishes an AI inventory and reviews.[22]
LANGUAGE ACCESS	The same city describes AI translation/captions for council meetings and 311.[23]
STAFF GUIDELINES	Published generative-AI guidelines for staff.[24]

STARTER CHECKLIST ITEMS SEEN IN AGENCY POLICIES

- Inventory of tools, owners, and data categories allowed.
- Approved SKUs versus prohibited consumer use for sensitive data.
- Resident disclosure when legally required.
- Retention schedule for prompts and outputs.
- Human review before public release.
- Training (including DIR-certified AI training where Texas law applies) and an incident path.[2]

This briefing is educational and vendor-neutral. It is not legal advice. Confirm obligations under public-records, privacy, procurement, and election law with counsel.

Provided by ohmyai.ai. Questions: help@xotavo.com or 713.980.2000. xotavo.com

XOTAVO · AI POLICY HANDBOOK

Corporate AI policy

Provided by ohmyai.ai

How organizations commonly write internal AI rules, assign governance, review vendors, and map duties under laws already summarized in this handbook. Descriptive and vendor-neutral. Snapshot as of Nov 3, 2026. Not legal advice.

Version 16.1
Last updated: November 3, 2026

IN ONE MINUTE

- Many organizations start with a written acceptable-use policy, an approved-tools list, and rules for confidential data in prompts.
- Governance usually names owners (legal, security, HR, IT, risk) and may add an AI committee for higher-risk uses.
- Vendor contracts commonly address data use, training opt-out, security attestations (for example SOC 2), indemnity, and subprocessors.
- NIST AI RMF and ISO/IEC 42001 are frequently cited voluntary frameworks for risk assessment and management systems.[25,27]
- Businesses may face deployer or provider duties under the EU AI Act and state rules on hiring tools, notices, and discrimination (Colorado, NYC Local Law 144, Illinois, Texas TRAIGA, California)—fact-specific; confirm with counsel.[21,1,48]

WRITING AN INTERNAL AI ACCEPTABLE-USE POLICY

Published corporate policies typically state purpose and scope (employees, contractors, systems), define “AI” for the organization, list permitted and prohibited uses, require human review for external or high-stakes outputs, and point to incident and escalation paths. Policies often cross-reference existing information-security, privacy, records, and code-of-conduct rules rather than replacing them. As of Nov 3, 2026, there is no single U.S. federal statute that mandates a private-sector AI acceptable-use policy; adoption is driven by risk management, customer contracts, and sector or state rules that apply to particular uses.

SCOPE	Which roles, devices, and environments the policy covers (corporate accounts only vs. any AI used for work).
ALLOWED USES	Drafting, summarization, coding assist, research—with named review requirements before customer or public release.

PROHIBITED USES	Commonly: autonomous legal, medical, or credit decisions without required human process; covert monitoring where law forbids it; generating deceptive content about the company or individuals.
EXCEPTIONS	A written path to approve higher-risk pilots with security and legal sign-off.

EMPLOYEE GUIDELINES

Data entry

Do not paste secrets, regulated personal data, non-public financials, or attorney-client material into tools that are not approved for that data class.

Confidential information

Trade secrets, customer lists, unpublished product plans, and credentials stay on approved systems; consumer AI accounts are often barred for work data.

Approved tools

A maintained allow-list (and sometimes a deny-list) tied to contract terms: retention, residency, non-training, logging, and SSO.

Attribution

Some policies require labeling AI-assisted drafts internally until a human owns the final version.

Verification

Staff remain responsible for factual accuracy, citations, and compliance of anything they send externally.

Training

Short, role-based training on approved tools and data classes; Texas DIR-certified training applies to covered public employers, not private firms generally.[2]

AI GOVERNANCE ROLES AND COMMITTEES

Organizations that document AI governance often assign an executive owner, a cross-functional working group, and clear RACI for legal, privacy, security, HR, procurement, and business units. Higher-risk uses (hiring, credit, safety-critical operations, customer-facing agents) are escalated to a committee or risk board. ISO/IEC 42001 describes an AI management system with leadership, planning, support, operation, evaluation, and improvement clauses that some enterprises pursue as a certifiable framework.[27] NIST AI RMF's Govern function similarly emphasizes accountability, culture, and mapped policies.[25]

VENDOR AND CONTRACT TERMS

TOPIC	QUESTIONS BUYERS COMMONLY PUT IN WRITING
Data use	Where prompts, files, and embeddings are stored and processed; subprocessors; retention an
Training opt-out	Whether customer inputs and outputs may be used to train or improve models; how opt-out is
Indemnity	Allocation of IP infringement, privacy, and third-party claim risk; caps and exclusions.
Security	Encryption, access control, logging, vulnerability process; breach notice timelines.

SOC 2 / ISO	Current SOC 2 Type II (Trust Services Criteria) or ISO/IEC 27001 reports; for AI management,
Service levels	Uptime, support, model-change notice, and evaluation rights when the provider swaps under

Attestations reduce but do not eliminate due diligence. Buyers still map data categories to the contract and to sector law.

RISK ASSESSMENT: NIST AI RMF AND ISO/IEC 42001

NIST AI RMF 1.0 (Jan 2023) is a voluntary U.S. framework organized as Govern, Map, Measure, Manage. The Generative AI Profile (NIST AI 600-1, Jul 2024) adds GAI-specific risks and suggested actions.[25,26] Organizations use it to inventory systems, map harms and stakeholders, measure tests and metrics, and manage residual risk with human oversight.

ISO/IEC 42001:2023 is a certifiable AI management-system standard. It is often paired with information-security controls (ISO/IEC 27001) rather than treated as a substitute for them.[27] As of Nov 3, 2026, adoption is voluntary unless a contract or regulator asks for it.

COMPLIANCE OBLIGATIONS UNDER LAWS ALREADY IN THIS HANDBOOK

Whether a business is in scope depends on role (provider vs deployer), location of users or decisions, and use case. The rows below point back to measures summarized earlier; effective dates and status are as of Nov 3, 2026.

INSTRUMENT	BUSINESS-RELEVANT DUTIES (HIGH LEVEL)
EU AI Act	Providers and deployers of in-scope systems face risk-tiered duties (documenta
Colorado ADMT / Chatbot Safety	Automated decision-making in consequential decisions; chatbot operator dutie
NYC Local Law 144	Employers/agencies using covered automated employment decision tools in N
Illinois HB 3773	Employer notice and anti-discrimination limits when AI is used in covered emp
Texas TRAIGA (HB 149)	Among other duties, government AI disclosure to residents; consumer-facing A
California SB 53 / AB 2013	Frontier/transparency and training-data documentation duties for covered dev

HR, HIRING, AND AI

Hiring and promotion tools that score or filter candidates can trigger NYC Local Law 144 bias-audit and notice rules when used in the city, Illinois Human Rights Act AI notice and discrimination limits, and general EEO and state human-rights statutes regardless of “AI” branding.[48,34] Colorado’s ADMT Act (effective Jan 1, 2027) addresses automated decisions in consequential contexts that can include employment.[31] Common corporate controls: vendor bias-audit packages, adverse-impact review, human reconsideration paths, and clear candidate notices. Workforce monitoring and generative drafting of performance reviews raise separate labor and privacy questions that counsel and HR typically handle under existing employment law.

IP AND COPYRIGHT FOR AI OUTPUT

As of Nov 3, 2026, U.S. Copyright Office registration guidance (effective Mar 16, 2023) states that material generated by AI without sufficient human authorship is not copyrightable; applicants must disclose AI-generated content and claim only human-authored portions.[50] Corporate IP policies often require staff to track AI assistance in creative and marketing work, avoid pasting third-party copyrighted material into prompts where licenses forbid it, and clarify ownership of outputs in vendor terms. Training-data litigation against model developers continues in U.S. courts; outcomes remain case-specific rather than a single national statute (see Policy topics).

RECORDKEEPING AND INCIDENT RESPONSE

RECORDS	Retain prompts/outputs where they are business records or evidence of a covered decision; align retention with legal holds and privacy minimization.
INVENTORY	Keep a living register of AI systems, owners, data classes, vendors, and risk tier.
INCIDENTS	Define what counts as an AI incident (data leak via prompt, harmful autonomous action, discriminatory outcome, deepfake brand misuse) and who is on-call.
RESPONSE	Contain access, preserve logs, notify per breach and contract clauses, remediate model or process, and document lessons—folded into existing IR playbooks where possible.

CORPORATE AI POLICY CHECKLIST

- Written acceptable-use policy with scope, allowed/prohibited uses, and human-review rules.
- Approved-tools list mapped to data classes; consumer tools barred for confidential data.
- Named governance owners; committee path for high-risk uses.
- Vendor terms: data use, training opt-out, indemnity, security, SOC 2 / ISO evidence, subprocessors.
- Risk assessments using NIST AI RMF (and ISO/IEC 42001 if pursuing a management system).
- Jurisdiction map: EU AI Act role; Colorado; NYC LL 144; Illinois; Texas TRAIGA; California developer duties.
- HR/hiring controls and notices where AEDTs or automated decisions are used.
- IP guidance for AI-assisted works and vendor output ownership.
- Inventory, retention, and incident-response steps tested at least annually.
- Employee training and an exception/approval path for pilots.

This briefing is educational and vendor-neutral. It is not legal advice. Confirm obligations with counsel for your jurisdictions, sectors, and contracts.

Provided by ohmyai.ai. Questions: help@xotavo.com or 713.980.2000. xotavo.com

XOTAVO · AI POLICY HANDBOOK

Deepfakes, elections, and child safety

Provided by ohmyai.ai

Synthetic media capabilities, enacted rules, and operational practices offices document.

Version 16.1
Last updated: November 3, 2026

DEEPFAKES AND ELECTIONS

A deepfake is synthetic or heavily altered media depicting a real person in words or actions that did not occur. Voice cloning can use short samples. NCSL counts 31 states with political-deepfake laws as of Oct 13, 2026, and notes ongoing constitutional litigation in several jurisdictions.[28]

Texas Election Code §255.004 makes it a Class A misdemeanor to create a deceptive deep fake video and publish it within 30 days of an election with intent to injure a candidate or influence the result.[14]

Practices offices list in readiness plans

- A named channel that confirms or denies a circulating clip.
- Retention of original recordings of public events for comparison.
- Call-back verification before acting on urgent voice requests for funds or data.
- Referral of suspected election-law violations to the appropriate authority.

PROVENANCE TECHNOLOGY

C2PA Content Credentials attach signed provenance metadata. Watermarking may be combined for durability. Content Authenticity Initiative materials state that credentials report origin claims; they are not a universal deepfake detector, and metadata can be stripped.[47]

CHILD SAFETY

Texas SB 20 (effective Sep 1, 2025) addresses obscene AI-generated images that appear to depict children.[15] Federal TAKE IT DOWN requires covered platforms to remove validly reported NCII within 48 hours as of May 19, 2026.[16,17] Colorado's Chatbot Safety Act adds duties for conversational AI operators beginning Jan 1, 2027.[31]

VOICE INTERFACES

Lifelike voice agents raise the same disclosure, records, and fraud-awareness issues as text chat. Texas TRAIGA disclosure duties apply when a government AI system interacts with residents.[1]

Appendix A: terms and acronyms

Provided by ohmyai.ai

Alphabetized plain-language definitions used in AI policy discussions. 109 entries.

Appendix A
Version 16.1
Last updated: November 3, 2026

Acceptable-use policy (AUP)
A written rule set stating which AI tools and uses are allowed or forbidden for an organization's staff and contractors, usually tied to data classes and human-review requirements.

AEDT (automated employment decision tool)
A term used in New York City Local Law 144 for software that substantially assists or replaces discretionary decision-making in hiring or promotion and meets the City's regulatory definition.

Agent
An AI system that not only answers but takes multi-step actions, such as opening software or filling forms, under defined permissions.

AGI (artificial general intelligence)
A proposed future AI system that could match or exceed human performance across most cognitive tasks; not a description of current commercial products.

AI Act (EU)
The European Union's Regulation (EU) 2024/1689 on artificial intelligence, which phases in obligations by risk category from 2025 through 2028.

AI RMF
NIST's Artificial Intelligence Risk Management Framework (AI 100-1, Jan 2023), a voluntary U.S. framework organized around Govern, Map, Measure, and Manage.

Alignment
The research and engineering goal of making model behavior consistent with stated human or organizational objectives.

API (application programming interface)
A defined way for one software system to request services from another, commonly how applications call an AI model.

ASI (artificial superintelligence)
A hypothetical AI system that substantially exceeds human cognitive ability across nearly all domains; a research and policy concept, not a deployed product class.

Automation bias
The tendency of people to over-trust machine recommendations relative to other evidence.

Benchmark
A standardized test used to compare model performance on tasks such as language understanding, coding, or math.

Bias (statistical / representational)
Systematic differences in model outputs across groups or contexts that reflect imbalances or gaps in training data, objectives, or deployment.

C2PA
Coalition for Content Provenance and Authenticity; an open technical standard for Content Credentials that record origin and edit history of digital media.

Chain of custody (digital)
Documented handling of electronic evidence so its integrity can be shown in audits or proceedings.

Chatbot / assistant
A product interface, usually conversational, built on one or more models that accepts prompts and returns text, images, or other media.

CISA
Cybersecurity and Infrastructure Security Agency, the U.S. federal agency that issues guidance on cyber and critical-infrastructure risks, including AI-related threats.

Closed model
A model offered only as a remote service, without publicly releasing weights that others can run independently.

Closed-loop cooling
A data-center cooling design that recirculates coolant rather than continuously evaporating large volumes of fresh water.

Cloud
Computing capacity provided over the internet by a third party; prompts and files are typically processed on the provider's servers.

Compute
The processing capacity, often measured in FLOPs or GPU/TPU hours, used to train or run a model.

Content Credentials
C2PA-based, cryptographically signed metadata describing how a piece of media was created or edited; not by itself a deepfake detector.

Context window
The maximum amount of text (tokens) a model can consider at once in a single session.

CSAM (child sexual abuse material)
Sexual images or video of children; many statutes also cover AI-generated depictions that appear to show children.

CV (computer vision)
AI methods that analyze images or video, such as object detection or face matching.

Data center
A facility housing servers where training and inference commonly occur.

Data sovereignty
The principle that stored data is subject to the laws of the jurisdiction where it resides.

Deepfake
Synthetic or heavily altered audio, image, or video made to depict a real person saying or doing something they did not.

Deployer (EU AI Act)
Under the EU AI Act, an organization that uses an AI system under its authority, as distinct from the provider that develops or places the system on the market; duties depend on risk tier and role.

Differential privacy
A mathematical approach that limits how much an individual's data can influence published results or models.

DIR (Texas)
Texas Department of Information Resources, which certifies certain state AI and cybersecurity training programs.

Embedding
A numeric representation of text, images, or other data that models use to compare similarity and retrieve related content.

EO (executive order)
A directive from a president or governor to executive agencies; it does not by itself repeal a statute.

ERCOT
Electric Reliability Council of Texas, which operates most of the Texas power grid.

Evaluation / evals
Systematic testing of model behavior on defined tasks, including accuracy, safety, and bias metrics.

Explainability / XAI
Techniques and disclosures intended to make model behavior or decision factors more understandable to humans.

FCC
Federal Communications Commission, the U.S. agency that regulates interstate communications.

Federated learning
Training across devices or organizations so raw data stays local while model updates are shared.

Fine-tuning
Additional training that adapts a base model to a narrower task, domain, or style.

FLOP / FLOPs
Floating-point operations; a unit used to describe the computational work of training or running a model.

Foundation model
A large model trained on broad data that can be adapted to many downstream tasks.

Frontier model
A term used in some U.S. state laws for highly capable foundation models defined by compute, capability, or revenue thresholds.

FTC
Federal Trade Commission, which enforces certain federal consumer-protection and NCII-removal rules.

Generative AI
AI systems that produce new text, images, audio, video, or code rather than only classifying existing inputs.

GPAI (general-purpose AI)
Under the EU AI Act, models that can perform a wide range of distinct tasks and can be integrated into many systems.

GPU (graphics processing unit)
Hardware commonly used to train and run large AI models because it performs many parallel calculations.

Grounding
Connecting model outputs to retrieved documents or tools so answers can cite external sources.

Guardrails
Filters, policies, and technical controls intended to limit disallowed or harmful model outputs.

Hallucination
A fluent model output that presents false or fabricated information as if it were factual.

High-risk AI system
Under the EU AI Act and similar frameworks, an AI use that can seriously affect health, safety, or fundamental rights and therefore faces extra duties.

Human-in-the-loop
A design in which a person must review or approve AI outputs before consequential action.

Incident response (AI)
Procedures for detecting, containing, reporting, and remediating harmful AI behavior or misuse.

Inference
Running a trained model to produce an output for a new input; everyday product use is inference.

ISO/IEC 42001
The 2023 international standard for artificial-intelligence management systems.

Jailbreak
A prompt or technique intended to bypass a model's safety controls.

Latency
The time between submitting a request and receiving a model response.

LLM (large language model)
A neural language model trained on large text corpora to predict tokens and generate or transform language.

Machine learning (ML)
Methods that fit models to data so systems improve at a task without being fully hand-coded for every case.

Model card
A document describing a model's intended use, training data summary, evaluation results, and known limits.

Model inversion / extraction
Attacks that try to recover training data or approximate a proprietary model from its outputs.

Model registry
An inventory of models in use, with owners, versions, and approved purposes.

Multimodal
A model that accepts or produces more than one modality, such as text and images.

NCII (non-consensual intimate imagery)
Intimate images shared or created without the depicted person's consent, including some AI-generated forgeries.

NIST
National Institute of Standards and Technology, a U.S. agency that publishes voluntary AI risk and measurement guidance.

NLP (natural language processing)
The field of methods for analyzing and generating human language.

OECD AI Principles
Non-binding intergovernmental principles on trustworthy AI first adopted by OECD members in 2019 and updated since.

On-device / local processing
Running a model on hardware controlled by the organization or user so data need not leave that environment.

Open source (software) vs open weight
Open-source usually means code under an OSI-style license; open-weight means parameters are released, which may use a different license.

Open-weight model
A model whose trained parameters are released so others can download and run them, subject to the license terms.

OSTP
White House Office of Science and Technology Policy.

Parameters / weights
The learned numeric values inside a neural network that encode patterns from training.

Personalization
Adapting outputs using a user's history or profile; raises distinct privacy and records questions in government use.

PII (personally identifiable information)
Information that can identify a person, alone or combined with other data.

Poisoning (data / model)
Inserting malicious or misleading data into training or retrieval corpora to alter later behavior.

Procurement
The process of buying goods or services, including contract terms for AI tools used by government.

Prompt
The instruction or query submitted to a model.

Prompt injection
An attack in which untrusted text tries to override system instructions given to a model or agent.

Provenance
Information about origin, authorship, and transformations of content or data.

Provider (EU AI Act)
Under the EU AI Act, a person or organization that develops an AI system or general-purpose AI model and places it on the market or puts it into service under its name.

Public records / FOIA analogue
Laws that can make prompts, outputs, and related logs subject to disclosure and retention rules.

PUE (power usage effectiveness)
A data-center efficiency ratio of total facility power to IT equipment power; closer to 1.0 means less overhead.

Quantization
Reducing the numeric precision of model weights to shrink size and cost, often with some accuracy trade-off.

RAG (retrieval-augmented generation)
A pattern in which a system retrieves documents first, then generates an answer conditioned on those documents.

Red teaming
Adversarial testing that deliberately probes a system for failures, misuse paths, or safety gaps.

Responsible scaling policy
A developer-stated plan that ties higher model capability to stronger safety measures before further scaling.

Risk classification
Grouping AI uses by potential harm so controls can scale with impact.

RLHF (reinforcement learning from human feedback)
A post-training method that adjusts a model using human preference ratings so outputs better match desired behavior.

Safety case
A structured argument, with evidence, that a system's residual risks are acceptable for a stated use.

Sandbox (regulatory)
A supervised environment in which firms can test AI systems under regulator oversight with tailored conditions.

Shadow AI
Staff use of consumer AI tools without organizational approval or logging.

SOC 2
An AICPA attestation report on a service organization's controls relevant to security, availability, processing integrity, confidentiality, and/or privacy under the Trust Services Criteria; commonly requested in vendor due diligence.

Social scoring
Rating people based on behavior or personal traits to grant or deny treatment or services; certain forms are restricted under the EU AI Act.

SOTA (state of the art)
The best published result on a given benchmark at a stated time.

Synthetic data
Artificially generated data used for training, testing, or privacy-preserving analytics.

Synthetic media
Audio, images, video, or text produced or substantially altered by AI.

System prompt
Hidden or privileged instructions that steer a model's behavior for a product or session.

Take-down / notice-and-removal
A legal or platform process requiring removal of specified unlawful content within a deadline after valid notice.

Temperature
A sampling setting that affects how deterministic or varied a model's next-token choices are.

Threat model
A structured description of who might misuse a system, how, and with what impact.

Token
A small text unit, often a word piece, that models read and write sequentially.

TPU (tensor processing unit)
A type of accelerator chip designed for neural-network workloads.

TRAIGA
Texas Responsible Artificial Intelligence Governance Act (HB 149), effective Jan 1, 2026.

Training
The process of adjusting model weights on large datasets so the model learns statistical patterns.

Training cut-off
The last date of data included in a model's training corpus, after which the base model has no built-in knowledge.

Training opt-out
A contractual or product setting stating that a customer's prompts, files, or outputs will not be used to train or improve a vendor's models.

Transformer
The neural-network architecture, introduced in 2017 research, that underlies most modern LLMs.

TX-RAMP
Texas Risk and Authorization Management Program, a security authorization required for certain cloud vendors serving Texas state agencies.

Vector database
Storage optimized for similarity search over embeddings, often used in RAG systems.

Watermarking (AI content)
Embedding a signal in generated media intended to help identify AI origin; often used alongside, not instead of, C2PA metadata.

Zero-shot / few-shot
Asking a model to perform a task with no examples (zero-shot) or a small number of examples (few-shot) in the prompt.

Definitions are for a general audience. Where a statute supplies its own definition, the statute controls.

Provided by ohmyai.ai. Questions: help@xotavo.com or 713.980.2000. xotavo.com

Appendix B: named tools, for reference

Provided by ohmyai.ai

Dated snapshot of widely known names grouped by family. Inclusion is not a ranking or endorsement.

Appendix B
Version 16.1
Last updated: November 3, 2026

Names illustrate the families in Briefing 3. Evaluate any product with the written questions in Briefing 9.

FAMILY	EXAMPLES (MAKER)
General assistants	ChatGPT (OpenAI) · Claude (Anthropic) · Copilot (Microsoft) · Gemini (Goog
Answer engines	Perplexity (Perplexity AI). Several general assistants also offer search-grou
Suite-embedded assistants	Copilot in Microsoft 365 (Microsoft) · Gemini in Google Workspace (Google)
Real-time and social tools	Grok on X (xAI)
Task agents	Agent modes offered within several general assistants above
Open-weight models for self-hosting	Gemma (Google) · Llama (Meta) · Mistral models (Mistral AI)

Listed alphabetically within each family. Names as of Nov 2026. Products and features change often.

SOURCES FOR THIS HANDBOOK

Numbers match the bracketed citations in the text. Also online at xotavo.com/sources.

1. **89(R) HB 149, Texas Responsible Artificial Intelligence Governance Act, enrolled version.** Texas Legislature, 2025; effective Jan 1, 2026. https://capitol.texas.gov/tlodocs/89R/billtext/html/HB00149F.htm (accessed Oct 26, 2026).
2. **89(R) HB 3512, enrolled version.** Texas Legislature, 2025; effective Sep 1, 2025. https://capitol.texas.gov/tlodocs/89R/billtext/pdf/HB03512F.pdf (accessed Oct 26, 2026).
3. **"Governor Abbott Announces Ban On Chinese AI, Social Media Apps".** Office of the Texas Governor, Jan 31, 2025. https://gov.texas.gov/news/post/governor-abbott-announces-ban-on-chinese-ai-social-media-apps (accessed Oct 26, 2026).
4. **Shehabi et al., 2024 United States Data Center Energy Usage Report.** Lawrence Berkeley National Laboratory, Dec 2024. https://eta-publications.lbl.gov/sites/default/files/2024-12/lbnl-2024-united-states-data-center-energy-usage-report_1.pdf (accessed Oct 26, 2026).
5. **"Sustainable by design: Next-generation datacenters consume zero water for cooling".** Microsoft Cloud Blog, Dec 9, 2024. https://www.microsoft.com/en-us/microsoft-cloud/blog/2024/12/09/sustainable-by-design-next-generation-datacenters-consume-zero-water-for-cooling/ (accessed Oct 26, 2026).
6. **Uptime Institute Global Data Center Survey 2025.** Uptime Institute, 2025. https://datacenter.uptimeinstitute.com/rs/711-RIA-145/images/2025.Annual.Survey.Report.pdf (accessed Oct 26, 2026).
7. **Meta's Odense Data Center (fact sheet).** Meta, Feb 2025. https://datacenters.atmeta.com/wp-content/uploads/2025/02/Metas-Odense-Data-Center.pdf (accessed Oct 26, 2026).
8. **Energy and AI: Energy demand from AI.** International Energy Agency, Apr 2025. https://www.iea.org/reports/energy-and-ai/energy-demand-from-ai (accessed Oct 26, 2026).
9. **89(R) SB 6, enrolled version.** Texas Legislature, 2025. https://capitol.texas.gov/tlodocs/89R/billtext/pdf/SB00006F.pdf (accessed Oct 26, 2026).
10. **Order adopting new 16 TAC §25.370, ERCOT Large Load Forecasting Criteria (Project 58480).** Public Utility Commission of Texas, Feb 6, 2026. https://ftp.puc.texas.gov/public/puct-info/agency/rulesnlaws/subrules/electric/25.370/58480adt.pdf (accessed Oct 26, 2026).
11. **Consumer AI Rights: Texas Responsible Artificial Intelligence Governance Act.** Office of the Texas Attorney General, undated web page. https://www.texasattorneygeneral.gov/consumer-protection/file-consumer-complaint/consumer-ai-rights (accessed Oct 26, 2026).
12. **FY26–27 Certified AI Awareness Training Programs.** Texas Department of Information Resources, Sep 2026. https://dir.texas.gov/sites/default/files/2026-09/FY26-27%20Certified%20AI%20Awareness%20Training%20Programs%20Final%204.0.pdf (accessed Oct 26, 2026).
13. **TX-RAMP Eligibility and Requirements.** Texas Department of Information Resources, undated web page. https://dir.texas.gov/information-security/texas-risk-and-authorization-management-program-tx-ramp/tx-ramp-eligibility (accessed Oct 26, 2026).
14. **Texas Election Code §255.004.** Texas Legislature, statutes online. https://statutes.capitol.texas.gov/?artSec=255.004&chapter=EL.255&code=EL&tab=1 (accessed Oct 26, 2026).
15. **89(R) SB 20, enrolled bill summary.** Texas Legislature, 2025; effective Sep 1, 2025. https://capitol.texas.gov/billlookup/BillSummary.aspx?Bill=SB20&LegSess=89R (accessed Oct 26, 2026).
16. **S.146, TAKE IT DOWN Act, all actions (Public Law 119-12).** Congress.gov, signed May 19, 2025. https://www.congress.gov/bill/119th-congress/senate-bill/146/all-info (accessed Oct 26, 2026).
17. **"FTC Begins Enforcing the TAKE IT DOWN Act".** Federal Trade Commission, May 2026. https://www.ftc.gov/news-events/news/press-releases/2026/05/ftc-begins-enforcing-take-it-down-act (accessed Oct 26, 2026).
18. **Executive Order 14365, "Ensuring a National Policy Framework for Artificial Intelligence".** The White House, Dec 11, 2025. https://www.whitehouse.gov/presidential-actions/2025/12/eliminating-state-law-obstruction-of-national-artificial-intelligence-policy/ (accessed Oct 26, 2026).
19. **OMB Memorandum M-25-21, Accelerating Federal Use of AI through Innovation, Governance, and Public Trust.** Office of Management and Budget, Apr 3, 2025. https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-21-Accelerating-Federal-Use-of-AI-through-Innovation-Governance-and-Public-Trust.pdf (accessed Oct 26, 2026).
20. **OMB Memorandum M-25-22, Driving Efficient Acquisition of Artificial Intelligence in Government.** Office of Management and Budget, Apr 3, 2025. https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-22-Driving-Efficient-Acquisition-of-Artificial-Intelligence-in-Government.pdf (accessed Oct 26, 2026).
21. **Timeline for the Implementation of the EU AI Act.** European Commission AI Act Service Desk, reflecting the Digital Omnibus on AI in force Jul 27, 2026. https://ai-act-service-desk.ec.europa.eu/en/ai-act/timeline/timeline-implementation-eu-ai-act (accessed Oct 26, 2026).

22. **Artificial Intelligence & Inventory.** City of San José, undated web page. https://www.sanjoseca.gov/your-government/departments-offices/information-technology/digital-privacy/ai-reviews-algorithm-register (accessed Oct 26, 2026).
23. **"San José City Clerk Launches Artificial Intelligence Initiative to Increase Language Access at City Council Meetings".** City of San José, news release. https://www.sanjoseca.gov/Home/Components/News/News/5907/4699 (accessed Oct 26, 2026).
24. **Generative AI Guidelines.** City of San José. https://www.sanjoseca.gov/home/showpublisheddocument/100095/638255600904303329 (accessed Oct 26, 2026).
25. **Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1.** National Institute of Standards and Technology, Jan 2023. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf (accessed Nov 3, 2026).
26. **Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1.** National Institute of Standards and Technology, Jul 2024. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf (accessed Nov 3, 2026).
27. **ISO/IEC 42001:2023, Information technology — Artificial intelligence — Management system.** International Organization for Standardization, Dec 2023. https://www.iso.org/standard/42001 (accessed Nov 3, 2026).
28. **Artificial Intelligence (AI) in Elections and Campaigns.** National Conference of State Legislatures, updated Oct 13, 2026. https://www.ncsl.org/elections-and-campaigns/artificial-intelligence-ai-in-elections-and-campaigns (accessed Nov 3, 2026).
29. **"Governor Newsom signs SB 53…".** Office of the Governor of California, Sep 29, 2025. https://www.gov.ca.gov/2025/09/29/governor-newsom-signs-sb-53-advancing-californias-world-leading-artificial-intelligence-industry/ (accessed Nov 3, 2026).
30. **SB 53 and the TDTA: How California is regulating AI transparency.** IAPP, describing SB 53 and AB 2013 effective Jan 1, 2026. https://iapp.org/news/a/sb-53-and-the-tdta-how-california-is-regulating-ai-transparency (accessed Nov 3, 2026).
31. **Colorado Automated Decision-Making Technology & Chatbot Safety Rulemaking.** Colorado Attorney General, describing SB 26-189 and HB 26-1263 effective Jan 1, 2027. https://coag.gov/ai/ (accessed Nov 3, 2026).
32. **AI FAQ / Artificial Intelligence Policy Act materials.** Utah Department of Commerce, Office of Artificial Intelligence Policy. https://commerce.utah.gov/ai/ai-faq/ (accessed Nov 3, 2026).
33. **NY State Senate Bill 2025-S8828 (RAISE Act chapter amendment), signed Mar 27, 2026.** New York State Senate. https://www.nysenate.gov/legislation/bills/2025/S8828 (accessed Nov 3, 2026).
34. **US AI Law Tracker: Illinois.** Orrick AI Law Center (HB 3773 effective Jan 1, 2026; SB 315 effective Jan 1, 2027). https://ai-law-center.orrick.com/illinois/ (accessed Nov 3, 2026).
35. **AI regulation in the UK (CBP-10003).** House of Commons Library. https://commonslibrary.parliament.uk/research-briefings/cbp-10003/ (accessed Nov 3, 2026).
36. **Interim Measures for the Management of Generative Artificial Intelligence Services.** Cyberspace Administration of China / State Council portal; English translation at China Law Translate. Effective Aug 15, 2023. https://www.chinalawtranslate.com/en/generative-ai-interim/ (accessed Nov 3, 2026).
37. **South Korea: Comprehensive AI Legal Framework Takes Effect.** Library of Congress Global Legal Monitor, Feb 20, 2026 (AI Basic Act effective Jan 22, 2026). https://www.loc.gov/item/global-legal-monitor/2026-02-20/south-korea-comprehensive-ai-legal-framework-takes-effect/ (accessed Nov 3, 2026).
38. **AI Act fully in force — next phase.** Cabinet Office of Japan, noting full enforcement Sep 1, 2025. https://www.cao.go.jp/press/new_wave/20251003 (accessed Nov 3, 2026).
39. **Doing business in Canada: Artificial intelligence and data.** Dentons, Sep 11, 2026 (AIDA not enacted after C-27 prorogation). https://www.dentons.com/en/insights/articles/2026/september/11/doing-business-in-canada-artificial-intelligence-and-data (accessed Nov 3, 2026).
40. **PL 2338/2023 bill page.** Câmara dos Deputados (Brazil). https://www.camara.leg.br/proposicoesWeb/fichadetramitacao/?idProposicao=2487262 (accessed Nov 3, 2026).
41. **India AI Governance Guidelines.** Principal Scientific Adviser / Government of India, Nov 2025. https://psa.gov.in/CMS/web/sites/default/files/publication/India%20AI%20Governance%20Guidelines%205%C2%A0Nov%C2%A02025.pdf (accessed Nov 3, 2026).
42. **Model AI Governance Framework / AI Verify programme materials.** Infocomm Media Development Authority (Singapore), referenced via governance handbooks summarizing voluntary tools (accessed Nov 3, 2026). https://www.imda.gov.sg/how-we-can-help/showcase-products/artificial-intelligence.
43. **Chapter 2 — Regulating the AI Industry in Australia.** Parliament of Australia Senate committee report materials. https://www.aph.gov.au/Parliamentary_Business/Committees/Senate/Adopting_Artificial_Intelligence_AI/AdoptingAI/Report/Chapter_2_-_Regulating_the_AI_Industry_in_Australia (accessed Nov 3, 2026).
44. **OECD AI Principles.** Organisation for Economic Co-operation and Development. https://www.oecd.org/en/topics/sub-issues/ai-principles.html (accessed Nov 3, 2026).
45. **Documents of Achievement | Hiroshima AI Process.** Ministry of Internal Affairs and Communications (Japan) / G7. https://www.soumu.go.jp/hiroshimaaiprocess/en/documents.html (accessed Nov 3, 2026).
46. **Council of Europe Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law.** Council of Europe treaty materials / Fedlex publication. https://www.fedlex.admin.ch/eli/treaty/2024/0301/en (accessed Nov 3, 2026).
47. **How it works — Content Credentials.** Content Authenticity Initiative (C2PA). https://contentauthenticity.org/how-it-works (accessed Nov 3, 2026).

48. **Automated Employment Decision Tools (AEDT) — Local Law 144 of 2021.** NYC Department of Consumer and Worker Protection. https://www.nyc.gov/site/dca/about/automated-employment-decision-tools.page (accessed Nov 3, 2026).
49. **2017 Trust Services Criteria (With Revised Points of Focus – 2022) / SOC suite overview.** AICPA & CIMA. https://www.aicpa-cima.com/resources/landing/system-and-organization-controls-soc-suite-of-services (accessed Nov 3, 2026).
50. **Copyright Registration Guidance: Works Containing Material Generated by Artificial Intelligence.** U.S. Copyright Office / Federal Register, Mar 16, 2023. https://www.federalregister.gov/documents/2023/03/16/2023-05321/copyright-registration-guidance-works-containing-material-generated-by-artificial-intelligence (accessed Nov 3, 2026).