REFERENCE GUIDE

A Basic Guide to *Generative AI*

A reference for common AI terms, concerns, and models



How AI actually works (and why it matters)

SARAH GIBSON

AReady | The Council of Independent Colleges

INTRODUCTION

What This Guide Covers

1. **How AI Works.** How generative AI uses learned patterns, and why that differs from human understanding.
2. **Core AI Concepts.** A fast-reference glossary of key terms.
3. **How AI Thinks.** Tokens, context, and why AI sometimes gets simple things wrong.
4. **AI Concerns Explained.** Bias, privacy, academic integrity, and more.
5. **AI Models & Tools.** A landscape of today's leading models and platforms.

How Generative AI Works

Generative AI systems are trained on large collections of data to recognize patterns, such as how language, images, or information are typically structured.

When you ask a question, the AI does not search for "the right answer" or think like a human; it predicts a response based on patterns it has learned and the context you provide.

This is why AI can be helpful for drafting and organizing ideas, but still requires human judgment to check accuracy and meaning.

What AI Is (and Is Not)

AI can perform sophisticated reasoning tasks, but it does not have human understanding, intentions, or reliable knowledge of truth.

Core AI Concepts, Part 1

TERM	DEFINITION
Artificial Intelligence (AI)	Computer systems designed to perform tasks that normally require human intelligence, such as analyzing information, writing text, or recognizing patterns.
Generative AI	A category of AI that creates new content — text, images, audio, video, or code — rather than retrieving existing material.
Large Language Model (LLM)	An AI model trained on massive amounts of text to predict and generate language. Think of it like autocomplete on steroids: it predicts likely next words from learned patterns, which lets it handle complex tasks without human understanding. Examples: the GPT, Claude, and Gemini models behind assistants such as ChatGPT.
Prompt	The instruction, question, or task given to an AI system that guides its output. Prompt template: Role, Task, Context, Voice, Format.
Responsible AI	An approach to AI adoption that emphasizes ethics, transparency, privacy, equity, and human oversight.
Bot	An AI configured to perform a specific role or task using instructions, goals, and boundaries — given a specialized set of instructions and knowledge to perform the role.
Agent	An AI system that can perform complex tasks, interact with webpages, and operate semi-autonomously toward a goal within defined constraints. Agents do not act independently of human goals; they operate within instructions and constraints set by people.



Bot vs. Agent: Two Ways AI Can Operate

Bot (Reactive Assistant)

Example: You ask a chatbot to rewrite an email. It generates a response based only on your prompt and the immediate context, then stops.

What's happening:

- One request → one response
- May remember the conversation, but does not pursue a goal on its own
- No follow-up unless you ask again

Think of it as: A helpful responder waiting for instructions.

Agent (Goal-Oriented Actor)

Example: You tell an AI agent to "plan a conference." It drafts an agenda, checks dates, revises the plan when conflicts appear, and adjusts recommendations based on constraints you approve.

What's happening:

- Works toward a goal over multiple steps
- Can make intermediate decisions
- May loop, revise, or pause for human input

Think of it as: A task manager that acts within boundaries you set.

Core AI Concepts, Part 2

TERM	DEFINITION
Agentic Workflow	A structured process where one or more AI agents complete tasks in stages (e.g., analyze → draft → revise → review).
AI Wrapper	A software layer that embeds an AI model inside a tool or workflow, shaping how users interact with it through buttons, menus, and constraints. Wrappers reduce friction but also obscure decision-making, making human oversight especially important.
Augmentation	Using AI to support and enhance human work rather than replace human judgment or expertise.
AGI (Artificial General Intelligence)	A hypothetical form of AI that could match or exceed human-level cognitive abilities across a wide range of tasks. Current AI systems are not AGI; they are specialized tools designed for specific purposes.
Turing Test	A test proposed by Alan Turing to evaluate whether a machine can imitate human conversation well enough to be indistinguishable from a human. Passing it indicates convincing language use, not understanding, reasoning, or intelligence.
Guardrails	Rules, constraints, or safeguards placed on AI systems to reduce risk and guide appropriate behavior.
Human-in-the-Loop	A design approach where humans guide, review, or approve AI outputs before final use.
BoodleBox	A collaborative AI workspace that provides access to multiple leading AI models in one environment. Users can switch between models, compare responses, build custom bots, connect shared knowledge, and collaborate with other people and AI in the same workspace.



Core Concepts: Context, Modality, and Knowledge

TERM	DEFINITION
Hallucination	When an AI generates information that is incorrect, fabricated, or unsupported but presented confidently.
Context	The information provided to the AI to shape its response — for example, a prompt, documents, links, or chat history.
Context Window	The maximum amount of information an AI model can consider at one time when generating a response. Context-window sizes vary widely by model and change frequently; a larger window does not automatically mean better understanding or performance.
Multi-Modal AI	AI systems that can work across multiple formats — text, images, audio, video, or data — within the same task.
Knowledge / Knowledge Base	The curated collection of documents an LLM can reference, from its training or from materials you explicitly provide. AI does not automatically retain institutional knowledge unless it is intentionally provided or connected through approved systems.
Tokens	The fundamental building blocks an AI language model uses — small units of text that can be a whole word, part of a word, or punctuation. The model breaks text into tokens, represents each numerically, and predicts what should come next based on training and context. This sequential prediction is the core mechanism behind its seemingly intelligent output.

The "Strawberry" Example

The Question

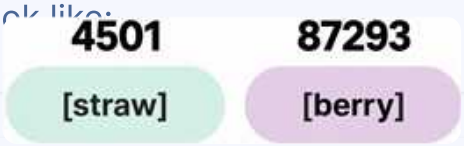
Why has AI sometimes struggled to count the r's in "strawberry"? A human naturally sees individual letters, leading to the correct answer: 3.



How AI Might Process It

An AI model might not process the word letter by letter. Depending on its tokenizer, "strawberry" could be represented as one or more token chunks — which helps explain historical "confident mistakes" on tasks requiring precise character counting, exact spelling, or letter-by-letter inspection.

A simplified token view might look like:



Illustrative example: exact token splits and token IDs vary by model and tokenizer.

The Underlying Lesson

For exact arithmetic or data work, instruct the AI to use a calculator, code interpreter, or another verification tool. While many newer models have improved and can answer this particular example correctly, the underlying lesson remains: fluent prediction is not the same as exact inspection.

A Better Approach

To get accurate character counts or similar exact data, make the task explicit: "Spell strawberry one letter at a time, then count the r's."

WORKING WITH AI

Tips to Work Better with AI

To maximize the effectiveness of AI and ensure outputs align with your objectives, integrating strategic interaction practices is crucial. By understanding how AI processes information, you can guide it more precisely, minimize errors, and achieve more reliable results.

Provide Clear Context

Always furnish the AI with comprehensive details regarding your audience, specific goals, any constraints (e.g., word count, tone), essential source material, and what constitutes a successful output. Supplying only relevant context prevents ambiguity and directs the AI towards a focused, pertinent response.

Structure Complex Tasks

For complex projects, decompose them into a series of smaller, sequential stages. This could involve steps like "analyze this data," "draft a summary," "critique the draft," and "revise based on feedback." This modular approach makes the AI's workflow transparent and simplifies the identification and correction of any missteps.

Leverage External Tools

While language models excel at synthesizing information, they are not always precise with factual recall or exact calculations. For tasks requiring exact figures, character counts, current events, or verified citations, instruct the AI to use integrated tools (like a calculator or search function) or prompt it to perform web searches, and always cross-verify its outputs independently.

Manage Context

Be mindful of the context window, which can become cluttered over time with irrelevant information or outdated assumptions in long conversations. When the task shifts significantly, or if the AI seems to be stuck or misinterpreting the direction, initiate a fresh chat to provide a clean slate and refocus its attention on the current objective.

Pro Tip: Make AI Help You Improve the Prompt

"Ask me questions one at a time until you have enough information to complete this well."

"How could I have asked for this better?"

"Before you do the task, tell me what information would make your answer substantially better."

How Models Change and Learn

TERM	DEFINITION
Fine-Tuning	Additional training applied to a model to improve performance for a specific task or context.
Model Drift	Changes in AI performance over time due to shifts in data, usage patterns, or context.
Training on Your Data	This refers to whether information entered into an AI tool may be stored and later used to improve the system over time. AI models do not learn from individual prompts in real time like human students, but some tools may incorporate user data into larger datasets that influence future versions of the model, depending on their policies.
Fast / Low-Reasoning Responses	Many current models can respond quickly using little or no additional reasoning computation. This is useful for straightforward questions, drafting, summarizing, classification, and routine work.
Reasoning / Thinking Modes	Many current models can devote additional computation to difficult problems such as analysis, planning, coding, mathematics, or multi-step decision support. Some models determine automatically when more reasoning is useful; others allow the user or system to select a reasoning or effort level.
GPT (Generative Pre-trained Transformer)	A training approach where language models are pre-trained on large text datasets using the Transformer architecture to generate human-like language.

System 1 and 2 Thinking

These terms are often used as an analogy for how an LLM produces responses. System 1 thinking is fast, automatic, and intuitive. It takes less processing power. System 2 thinking is slower, deliberate, and analytical. It is used when we reason carefully, evaluate evidence, solve complex problems, or reflect on our assumptions. It takes more processing power and energy.

System 1 (Fast, Intuitive)

What's happening:

- Quick pattern matching
- No explicit step-by-step reasoning
- Efficient but more prone to errors

Think of it as:

Autocomplete on steroids.

System 2 (Slow, Deliberate)

What's happening:

- Explicit reasoning steps
- Uses additional computation for tasks
- Useful for complex analysis and problems

Think of it as:

An assistant that spends more effort working through the problem.

Reasoning Approaches & Grounding

Retrieval-Augmented Generation (RAG)

A method where AI pulls information from approved documents before generating a response. This helps reduce hallucinations by grounding AI responses in approved, institution-specific sources.

Example: An AI answers questions about your organization's policies by pulling directly from your uploaded documents instead of relying on its general training.

What's happening:
- Retrieves relevant source material
- Generates responses grounded in that material
- Reduces hallucinations and improves accuracy

Think of it as: An AI that checks the book before answering.

Chain of Thought (CoT)

A term used to describe step-by-step reasoning used to work through a problem. Modern reasoning models may perform substantial reasoning internally, but the explanation shown to a user is not necessarily a complete or literal record of that internal process.

Example: An AI solves a math problem by explaining each step in sequence from start to finish.

What's happening:
- Linear reasoning
- One path forward
- Clear but limited if the first path is wrong

Think of it as: Walking straight down a trail.

Tree of Thought (ToT)

A reasoning or prompting approach that encourages exploration of multiple possible paths before selecting a solution, similar to branching human thinking.

Example: An AI considers multiple possible strategies for solving a problem, explores each briefly, then chooses the best option.

What's happening:
- Branching reasoning paths
- Comparison between options
- Better for planning and open-ended problems

Think of it as: Standing at a fork and exploring several routes before choosing.

Bias

The concern: AI can reflect biases present in the data it was trained on.

What's really happening: AI does not decide to be biased; it reflects patterns in the data it was trained on. Many large datasets overrepresent Western, English-language, and historically dominant perspectives, which can shape outputs in uneven ways.

Bottom line: bias is a reason for oversight, not avoidance.

WHAT TO DO

- Treat AI as a drafting and support tool, not a final authority
- Review outputs critically, especially in grading, advising, or hiring contexts
- Keep humans in the loop for judgment and equity

Data & Privacy (Why "Free" Tools Can Be Risky)

The concern: information entered into some free AI tools may be stored, reviewed, or used for future training.

What's really happening: not all AI tools follow the same data-handling rules. Privacy depends on the specific service, your account settings, its terms or contract, and the data controls it offers, not simply on price. Many free consumer tools use what you enter to train future models by default, and opting out is not always possible.

Bottom line: the risk isn't AI, it's where your data goes.

WHAT TO DO

- Do not enter grades, student records, health data, or confidential information
- Use institution-approved tools with contractual data protections, and check each tool's data settings
- Assume anything sensitive should stay out of public AI systems

Environmental Impact

The concern: AI uses energy and computing resources, contributing to environmental costs.

What's really happening: large AI models require significant energy because they run on powerful data centers that use electricity for computation and cooling, and water is often used to remove heat from servers.

Bottom line: responsible use is about how often and how thoughtfully, not never.

WHAT TO DO

- Use AI intentionally, not excessively
- Favor tools and workflows that reduce duplicated work
- Remember that efficient use matters more than total avoidance

Academic Integrity

The concern: students may misuse AI to bypass learning.

What's really happening: AI has changed how work is produced, not the core goals of learning. Good pedagogy addresses many academic integrity concerns by emphasizing learning processes, reflection, and application.

Bottom line: clarity and transparency reduce misconduct more than bans.

WHAT TO DO

- Be explicit about when and how AI is allowed
- Design assignments that emphasize process, reflection, and application
- Treat AI literacy as part of academic honesty, not separate from it

Why AI Feels Smarter Than It Is

Generative AI often sounds confident, fluent, and authoritative, even when it's wrong. That confidence can create the illusion of understanding, leading users to trust outputs more than they should.

What feels like "intelligence" is usually strong pattern recognition combined with polished language. The system is optimized to produce responses that sound right, not to verify truth, intent, or impact.

This gap between confidence and competence is why human judgment still matters. When AI outputs are clear, persuasive, or emotionally resonant, it becomes harder to notice errors, omissions, or subtle misalignments with context and values. Understanding this illusion helps users slow down, ask better questions, and decide when verification, not automation, is the responsible next step.

AI is often most convincing at the exact moment it deserves the most scrutiny.

AI Detectors

The concern: AI detectors produce false positives.

What's really happening: reliable detection is increasingly difficult. As LLMs become more sophisticated, their output becomes harder to distinguish statistically from human writing. Turnitin has said its detector is tuned to miss roughly 15% of AI-generated text to keep false positives low, and research shows accuracy drops further when AI text is paraphrased or edited to appear more human. The "tell-tale signs" faculty look for, repetitive phrasing, overly formal tone, are fading as models improve. More troubling, research has found detectors disproportionately flag writing by non-native English speakers, and they can misjudge students with learning differences or anyone whose writing doesn't match the "standard" patterns the detector was trained on, creating false accusations that can devastate a student's academic record.

Bottom line: authentication over detection — move from "Did a bot write this?" to "How did you use tools to create this?" Transparency and process documentation are your new academic integrity tools.

Sources: Turnitin (2023), statements on AI writing detection; Weber-Wulff et al. (2023), "Testing of detection tools for AI-generated text," International Journal for Educational Integrity; Liang et al. (2023), "GPT detectors are biased against non-native English writers," Patterns.

WHAT TO DO

- Focus on authenticating the process, not detecting the product
- When you suspect AI misuse, approach with curiosity: "Can you walk me through how you developed this idea?"
- Redesign assignments to emphasize drafts, reflections, and presentations that show thinking development

AI CONCERNS EXPLAINED

Copyright (How to Avoid Infringement)

The concern: AI might copy or infringe on copyrighted material.

What's really happening: generative AI usually produces new content based on patterns, but models can sometimes reproduce copyrighted text or images closely, even without being asked to.

Legal landscape: courts are still actively determining how copyright law applies to AI, particularly around the use of copyrighted materials in training. Recent cases suggest that outputs are not automatically infringing, but they also reinforce that transformation, market impact, and how the tool is used still matter.

Bottom line: use AI as a tool, not a substitute for authorship or citation.

WHAT TO DO

- Avoid prompting AI to imitate specific living authors or proprietary sources
- Use AI for brainstorming, outlining, and drafting, not final attribution
- Cite sources when AI assists with research or synthesis

Your Next Steps

Start small. Try one AI task that saves you 15 minutes this week.

Stay transparent. Tell others how you're using AI in your work.

Keep learning. AI literacy is now part of professional development.

AI Models and Tools Overview

Not every AI tool uses the same kind of model. Language models process and generate text, while other models specialize in images, video, or audio. AI platforms such as Copilot and BoodleBox can combine various models with search, retrieval, tools, interfaces, and workflows.

How to read this section: These personas describe how models behave in practice, not intelligence or autonomy. Model landscape verified September 2026.

Foundation & General-Purpose AI Models

MODEL FAMILY	PERSONA	KNOWN FOR	OFTEN A GOOD CHOICE FOR
GPT — OpenAI	The Strategist	Broad capability, reasoning, and tool use	Complex professional work, reasoning, coding, synthesis, planning, research, and agentic workflows
Claude — Anthropic	The Thoughtful Collaborator	Writing, analysis, coding, and long-context work	Documents, nuanced writing, deep analysis, complex coding, and extended projects
Gemini — Google	The Multimodal Native	Multimodal capability and Google integration	Text, image, audio, and video work; research; and Google-connected workflows
Llama — Meta	The Open Builder	Open-weight models and customization	Private deployments, experimentation, customization, and locally hosted AI
Grok — xAI	The Wild Card	Real-time information and X integration	Current-information research, reasoning, coding, knowledge work, and agentic tasks

Foundation & General-Purpose AI Models (continued)

MODEL FAMILY	PERSONA	KNOWN FOR	OFTEN A GOOD CHOICE FOR
DeepSeek	The Mathlete	Reasoning, math, and coding	Mathematics, technical problem-solving, coding, research, and complex analytical work
Amazon Nova	The Enterprise Builder	Cost-efficient multimodal enterprise AI	Document processing, business automation, multimodal analysis, and AWS-centered workflows
Mistral	The Flexible Engineer	Open models, efficiency, and flexible deployment	Coding, multilingual work, customization, private deployment, and general-purpose AI
NVIDIA Nemotron	The Systems Analyst	AI systems, reasoning, and agent development	Building agents, AI infrastructure, reasoning systems, customization, and technical workflows
Cohere	The Enterprise Librarian	Enterprise retrieval and private knowledge	RAG, enterprise search, multilingual work, private data, and knowledge-intensive applications

AI Model Families

GPT, Claude, Gemini, and Llama are families of models, not single models. Just as a car manufacturer produces various models for different needs, AI companies offer multiple models within the same family. Some models prioritize speed and low cost, while others devote more computing power to reasoning and complex tasks.

Understanding these distinctions helps you choose the right tool for the job, much like you wouldn't use a moving truck for a quick trip across town.

Choosing the right model: models within the same family are often differentiated by speed, reasoning depth, overall capability, and compute intensity. The goal is not to always choose the most powerful model — it is to choose the smallest model that can do the job well.

Ratings in the tables that follow, such as Highest, Exceptional, and Very Deep, are instructional judgments to help you compare models, not independently benchmarked measurements.

GPT — OpenAI Family

OpenAI's GPT models are a prominent general-purpose AI family, adept at a wide range of tasks including writing, reasoning, coding, analysis, research, multimodal work, and powering AI agents.

MODEL	SPEED	REASONING DEPTH*	OVERALL CAPABILITY	COMPUTE INTENSITY**	QUICK RULE
GPT-Astra	Fast	Exceptional	Highest	Very High	Tackle the hardest, end-to-end work.
GPT-Sol	Moderate–Fast	Very Deep	Very High	High	Think deeply.
GPT-Terra	Fast	Strong	High	Moderate	Handle everyday work.
GPT-Luna	Fastest	Moderate	Moderate–High	Low	Do it quickly.

Claude — Anthropic Family

Anthropic's Claude models are known for strong writing, analysis, coding, long-context work, and agentic workflows.

MODEL	SPEED	REASONING DEPTH*	OVERALL CAPABILITY	COMPUTE INTENSITY**	QUICK RULE
Claude-Fable	Moderate	Exceptional	Highest	Very High	Run long, complex workflows.
Claude-Opus	Moderate	Very Deep	Very High	High	Tackle the hardest thinking.
Claude-Sonnet	Fast	Strong	High	Moderate	Handle everyday professional work.
Claude-Haiku	Fastest	Moderate	Moderate–High	Low	Get lightweight work done quickly.

Gemini — Google Family

Google's Gemini models are especially strong in multimodal work, reasoning, coding, research, and integration across the Google ecosystem.

MODEL	SPEED	REASONING DEPTH*	OVERALL CAPABILITY	COMPUTE INTENSITY**	QUICK RULE
Gemini-Pro	Moderate	Very Deep	Very High	High	Use for difficult reasoning and complex work.
Gemini-Flash	Very Fast	Strong	High	Moderate to Low	Move quickly across text, images, video, and routine workflows.

Llama — Meta Family

Meta's Llama models are open-weight models designed for customization, private deployment, experimentation, and building AI systems.

MODEL	SPEED	REASONING DEPTH*	OVERALL CAPABILITY	COMPUTE INTENSITY**	QUICK RULE
Llama-Maverick	Moderate–Fast	Strong	High	Moderate	Use when you want more capability and flexibility.
Llama-Scout	Fast	Moderate–Strong	Moderate–High	Low	Use for efficient, customizable deployments.

*Reasoning depth is a relative comparison within and across these model families, not a precise benchmark.

**Compute intensity is a relative guide, not a measurement of energy use. Actual compute and energy use vary by task, reasoning settings, model configuration, infrastructure, and length of interaction.

Search, Research & Retrieval Tools

TOOL / PLATFORM	PERSONA	SPECIALTY
Perplexity	The AI Search Engine	Web-grounded search, research, synthesis, citations, and agent-assisted information gathering.
Semantic Scholar	The Academic Expert	AI-assisted discovery and exploration of scholarly literature, papers, citations, and related research.
Tavily	The Web Connector	Gives AI systems real-time web search, extraction, crawling, mapping, and research capabilities.
Gemini Notebook (formerly NotebookLM)	The Research Companion	Research, synthesis, organization, and question answering grounded in selected sources with inline citations; can also create study guides, audio, mind maps, and other outputs from those sources.
ResearchRabbit	The Literature Explorer	Discovers related academic papers, traces citation networks, maps connections between studies and authors, and helps researchers explore and organize a field of literature.

AI MODELS & TOOLS

A Simple Rule for Research Tools

Search finds information. Retrieval brings relevant information into the AI's context. Generative AI helps you make sense of it.

MATCH THE TOOL TO THE TASK

Current information → use a web-grounded search tool like Perplexity

Academic literature → use Semantic Scholar or ResearchRabbit

Work from specific sources → use NotebookLM

Research tip: follow citations back to the original source — an AI-generated summary is not a substitute for reading and verifying the evidence

Image Generation Models

MODEL	PERSONA	SPECIALTY
Nano Banana	The Visual Designer	Google's image-generation and editing family for creating, revising, and refining visuals through natural-language instructions.
FLUX	The Precision Artist	Black Forest Labs' current multimodal visual-intelligence family spanning image, video, audio, and related creative generation.
Stable Diffusion	The Open-Weight Tinkerer	Customizable image generation with Large, Large Turbo, and Medium variants suited to self-hosted and specialized creative workflows.
GPT Image	The Generalist Creator	High-quality image generation, image-conditioned creation, instruction following, and precise editing in OpenAI workflows.

Video Generation Models & Tools

MODEL / TOOL	PERSONA	SPECIALTY
Veo	The Cinematographer	Cinematic video generation with native synchronized audio, creative controls, video extension, and image-guided generation.
Runway Gen	The Indie Filmmaker	Text-to-video and image-to-video generation for creative filmmaking, visual experimentation, and production workflows.
Pika	The Social Creator	Short-form and stylized AI video creation, photo animation, creative effects, and video-oriented workflows.

Audio, Voice & Transcription

MODEL / TOOL	PERSONA	SPECIALTY
Eleven Labs	The Voice Actor	Expressive, natural text-to-speech, emotional delivery, multi-speaker dialogue, and multilingual narration.
GPT Transcribe	The Transcriber	Accurate transcription of recorded speech for interviews, lectures, meetings, accessibility, and other spoken-audio workflows.
Wispr Flow	The Voice Typist	Turns natural speech into polished written text inside other apps, automatically handling punctuation, formatting, filler words, and self-corrections. Works across Mac, Windows, iPhone, and Android.
Granola	The Meeting Memory	Captures meeting audio without adding a meeting bot, combines transcripts with the user's own notes, produces structured notes and action items, and lets users search and chat across past meetings.
Bluedot	The Meeting Scribe	Bot-free meeting capture, transcription, summaries, action items, searchable meeting knowledge, and follow-up workflows.

AI Assistants, Workspaces & Productivity Platforms

TOOL / PLATFORM	PERSONA	SPECIALTY
Microsoft Copilot	The Business Professional	AI assistance across Microsoft workplace applications, documents, communication, research, agents, and workflows. Modern Copilot experiences can use models from multiple providers.
Inflection Pi	The Empathetic Listener	Personal, supportive, emotionally tuned conversation designed around human-centered interaction and everyday reflection.
BoodleBox	The AI Team Room	Access to many leading AI models in one environment, plus model comparison, shared context, custom bots, knowledge, collaboration, and education-oriented controls.
Canva AI	The Design Assistant	Conversational design, editable AI-generated layouts, visual creation, brand-aware work, iterative editing, and creative workflows.
Gamma	The Presentation Builder	Fast creation of presentations, visual documents, websites, social content, and graphics with AI-assisted structure and design.

Where Judgment Lives

AI systems can generate outputs and make intermediate decisions within configured workflows. Human judgment remains essential for deciding whether the system should be used, what authority it receives, and whether its outputs or actions are appropriate. Judgment lives at three levels:

The Model

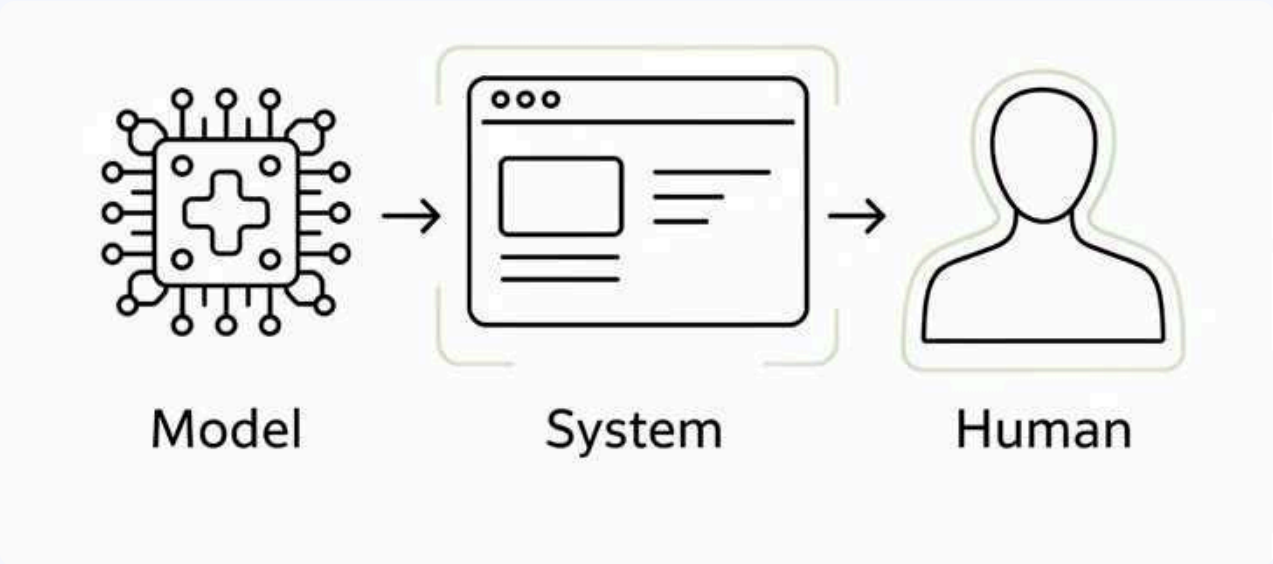
Predicts patterns based on training and context

The System

Embeds the model inside tools, defaults, and workflows

The Human

Decides whether, how, and where the output is used



Most risk does not come from the model alone; it comes from unclear responsibility at the system and human level. Responsible AI use requires knowing where judgment stops and starts.

What to Remember as You Use Generative *AI*

Generative AI is not a single tool, model, or moment. It is a system made up of layered decisions about data, design, defaults, and use. If you can explain what the AI is doing, where it's embedded, and who is accountable for the outcome, you're using it thoughtfully.

© 2026 Sarah Gibson. ProfessorGibson.com.

A Basic Guide to Generative AI is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License (CC BY-NC-SA 4.0). You may share and adapt this work for noncommercial purposes with appropriate attribution. Adaptations must be distributed under the same license.

