


Introduction to How Generative AI Works


- - Generative AI creates new text, images, and code by learning patterns from data, not by copying or classifying.
- - Traditional AI classifies and predicts; generative AI produces original content from statistical patterns.
- - ChatGPT, Claude, Midjourney, GitHub Copilot, ElevenLabs, and Spotify AI DJ are common 2026 examples.
- - Non-technical pros need a working model to evaluate tools and use AI effectively without a CS degree.
- - You will learn to explain the AI workflow, craft prompts, handle multimodal inputs, and assess risks.

From Classifying to Creating: The Evolution of AI



- Rule-based → ML → deep learning → generative & foundation models
- Traditional AI classifies data; generative AI produces new content
- Compare 2018 autocomplete to 2026 multi-paragraph reasoning & coding
- Mid-2026 models pass PhD-level benchmarks yet fail simple tasks – the 'jagged frontier'
- AI is now a general-purpose creation engine accessible via plain language

The Core Workflow: Training, Pattern Learning, and Generation




- **Training:** Model learns from trillions of words and billions of images, not by memorizing facts.
- **Pattern learning:** Identifies which words, pixels, or sounds tend to appear together.
- **Generation:** Predicts the next most likely element—word, pixel, note—based on learned patterns.
- **‘Autocomplete on steroids’:** Predicts entire paragraphs with vast scale and sophistication.
- **Critical caveat:** Produces statistically plausible, not necessarily true, outputs.


Foundation Models:

The Engines Behind Modern AI

- Large, general-purpose AI trained on broad data, adaptable to thousands of specific tasks
- Pre-training costs millions; adapting a foundation model is cheap and fast
- Major 2026 families: GPT-5, Claude, Gemini, Llama, DeepSeek, and Grok
- Span text, images, audio, video, and code; multimodal models handle several at once
- The model underneath sets your product's cost, continuity, and capability ceiling



How Inputs Become Outputs: Prompts, Tokens, and Context



- Prompts are broken into tokens, which the model reads to predict the next most likely token



- Same architecture creates text, images, or code because it was trained on different data types



- Broad training lets a single model chat, summarize, translate, and brainstorm without retraining



- Context window is the model's working memory—modern models handle 200K+ tokens (~750 pages)



- Output cap matters: even with a huge context window, the model writes a limited amount per turn

Crafting Effective Prompts: The Skill of Instructing AI



Modern house
two floors
big windows
nice look



- Shift from keyword searching to instructing with context, audience, and format



- The five elements:
Role, Context, Task, Format, Constraints



- Context is the highest-leverage element most people skip



- One concrete example beats five paragraphs of description



- Iteration — refining through conversation — is the highest-leverage habit

Advanced Prompting Strategies That Make a Difference



- **Few-shot prompting:** show 1-3 examples instead of describing output.



- **Chain-of-thought:** ask "think step by step" for complex reasoning.



- **Self-critique loop:** have AI evaluate, then improve its own answer.



- **Constraint stacking:** combine audience, format, word count, tone.




- **Skip the incantations:** modern models need clarity, not role-play.




Bringing AI Into the Enterprise: APIs, Copilots, and RAG

- Generative AI in tools you use: APIs, embedded copilots, and chat
- RAG grounds answers in your real documents like policies and manuals
- Retrieve relevant passages first, then generate a sourced answer
- Provides citations, making responses auditable without retraining
- You control document access and review; AI is a reasoning engine



Multimodal AI and the Shift Toward Agents



- Multimodal models: accept text, images, audio in one conversation — show a chart, get analysis
- Six eras from 2022 to 2026: chat → multimodal → search → memory → operations → autonomous agents
- AI agents don't just answer; they take action — searching files, booking meetings, running workflows
- Infrastructure that made agents work: reliable tool calls, huge context windows, shared protocols
- As of Feb 2026, frontier models autonomously complete ~14.5 hours of expert work, doubling every 4 months


Understanding Hallucinations and Why They Happen

- AI generates fluent, confident, but factually wrong content — including invented citations
- The model predicts likely words, it does not know or verify truth
- Hallucination is a fundamental characteristic, not a temporary bug
- Ground outputs in your documents with RAG; always verify AI memory claims
- Trust summaries of provided material; for high-stakes decisions, the human decides



Bias, Privacy, and Ethical Risks Professionals Must Understand

- Training data bias amplifies societal stereotypes in AI outputs
- Bias can enter at any stage: data, training, deployment, or use
- Models can memorize and leak personal data from training sets
- Always check: Does the vendor train on your data? Where is it stored?
- Never input confidential data without consent, safeguards, and a DPIA



The Pattern Loom

Copyright, Accountability, and When Not to Use AI



- Copyright law remains unsettled; training data raises unresolved infringement questions.



- You are accountable for AI outputs you publish, sign, or act upon.



- Automation bias leads users to over-trust fluent but incorrect AI outputs.



- Never use AI for safety-critical decisions without human review or verification.



- Baseline AI literacy is now a core professional obligation.

USE AI?

A dramatic background illustration featuring a stone path that leads to a cliff edge. A glowing blue line traces the path, which is illuminated by small lights. A large, glowing blue sign with the text 'USE AI?' stands at the edge. To the right, a red flag flies on a pole. Further right, two warning signs are visible: one for 'Confidential Data' and another for 'No Human Review'. The scene is set against a dark, stormy sky with mountains in the background.

Confidential Data

No Human Review

Evaluating Generative AI Tools:

A Practical Framework



- Only 23% of districts had AI-specific evaluation processes in 2024; traditional procurement misses critical AI risks.



- Evaluate across six dimensions: educational effectiveness (25%), privacy and governance (25%), and four others.



- Always ask: "Is our data used to train AI models?" Demand written attestation from both vendor and AI provider.



- Require source citation: systems that don't show which passage produced which answer are inadequate for business use.



- Pilot before deploying: 4–6 weeks with 8–12 real users under normal conditions beats any demo.



Privacy Score



Accuracy Gauge



Cost Meter



Designing Your First AI Experiment



- Pick a real task: write an email, summarize a document, or brainstorm an idea you need anyway



- Give clear context, audience, and format instructions; compare output to your own work



- Track time saved, editing needed, and what the AI did better or worse than expected



- Iterate instead of abandoning: tell the AI what was wrong and ask for a revision



- Build a shared library of tested prompt patterns for your team's most common tasks



Communicating AI Capabilities to Colleagues and Students

- Frame AI as a fast, literal assistant needing clear briefs; for students, a study partner who has read everything
- **Good at:** summarizing material, first drafts, brainstorming, format translation, processing large volumes
- **Not good at:** verifying facts from memory, genuine judgment, emotional nuance, accountability for decisions
- AI is here to stay—understanding it wisely is now a core professional competency
- Normalize AI as a tool, not a threat; model critical review openly so others learn to do the same



PersonWise.

PersonWise • AI digital-human course platform



Scan the code or click anywhere to watch this course live, with voice Q&A

personwise.ai/t/e7f5304c/public