About This Glossary

These definitions aren't dictionary perfect, and that's intentional. This document and the live page exists to keep these concepts clear for me, someone without a tech background, without all the complicated jargon.

If it makes things clearer for you too, even better. Share it with someone else who could benefit from it too.

I'll be adding to this over time, so come back to The Futureproof Directive to check for updates occasionally.

Here's what I've got so far:

Current Count: 101 (Version 1)

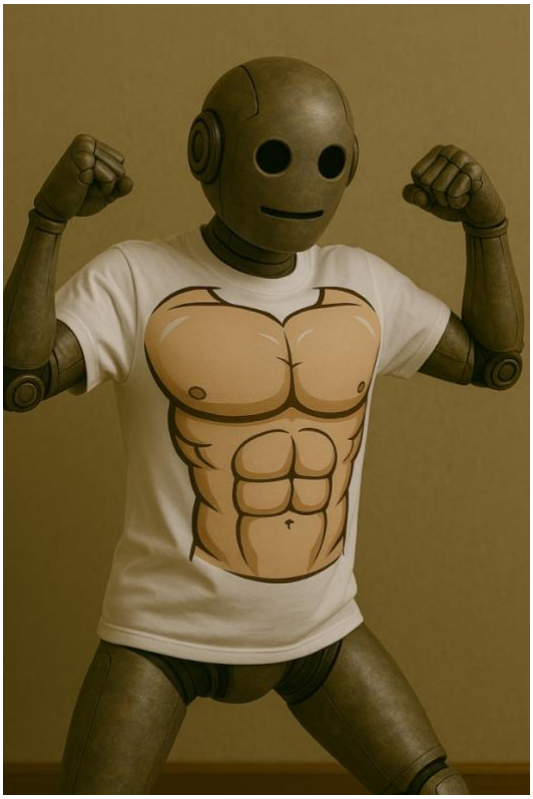

Because understanding AI terms is easier than building these abs.

Table of Contents

Active-Prompting
AGI (Artificial General Intelligence)
Alignment
Algorithm
Application
Application Programming Interface (API)
Artificial Intelligence (AI)
Artificial Neural Networks
Attention Mechanism
Auto-encoder
Automatic Prompt Engineer (APE)
Automatic Reasoning and Tool-Use (ART)

B.

Backpropagation
BERT (Bidirectional Encoder Representations from Transformers)
Bias
Big Data
Build a Neural Network

C.

Chain-of-Thought
Chain-of-Verification (CoVe)
Chatbot
Cognitive Computing
Convolutional Neural Network (CNN)
Context Window
CREATE Prompting

D.

Data Mining
Decision Traceability
Deep Learning
Directional Stimulus Prompting (DSP)

E.

Embeddings
Embeddings Search

F.

Few-Shot Chain-of-Thought Prompting
Few-Shot Learning
Few-Shot Prompting
Fine-Tuning

G.

GAN (Generative Adversarial Network)
Generated Knowledge Prompting

Generative AI
Generative Pretrained Transformer (GPT)
Gradient Descent
Graphics Processing Unit (GPU)

H.

Hard Prompting
Hallucinations
Human in the Loop
Hyperparameter Tuning

I.

Image Diffusion Models
Inference

J.

K.

L.

Large Language Model (LLM)
Learning from Mistakes
Least-to-Most Prompting

M.

Machine Learning (ML)
Massive Datasets
Meta-Prompting
Model
Multimodal AI

N.

Natural Language Processing (NLP)
Natural Language Generation (NLG)
Natural Language Understanding (NLU)
Neural Network
Notebook

O.

One-Shot Prompting
Outcome
Overfitting

P.

Predictive AI
Program-Aided Language Models (PAL)
Prompt
Prompt Chaining
Prompt Engineering

Q.

R.

RAG (Retrieval-Augmented Generation)
ReAct Prompting
Reasoning Engines
Recurrent Neural Network (RNN)
Reflexion
Reinforcement Learning
Responsible AI

S.

Select the Best Algorithm
Self-Ask Prompting
Self-Consistency Decoding
Self-Supervised Learning
Semantic Search
Single-Shot Learning
Skeleton-of-Thought Prompting
Soft Prompting
Stochastic Parrots
Strong AI vs. Weak AI
Supervised Learning

T.

Temperature (Sampling Temperature)
Tensor Processing Unit (TPU)
The Activation Bias
The Four Ds of Automation
The Generative AI Revolution
Tokens
Top-p (Nucleus Sampling)
Training Data
Transformer Architecture
Tree-of-Thoughts Prompting
Two-Shot Prompting

U.

Unsupervised Learning
Using AI Systems

V.

VAE (Variational Autoencoder)

W.

X.

Y.

Z.

Zero-Shot Chain-of-Thought

Zero-Shot Prompting

AI Terms/Definitions

A.

Active-Prompting

A prompting approach where you (or an algorithm) actively select which examples will teach the AI best. Instead of random examples, you pick the ones you think are most useful to guide its output.

Prompt Example: “Here are five customer complaints. Pick the one that will teach the model the most about tone and rewrite it positively.”

AGI (Artificial General Intelligence)

AI that would think like a human across any task, not just one job. Basically, sci-fi level intelligence that doesn’t exist yet.

Alignment

Making sure AI systems do what humans actually want and don’t cause harm. The holy grail of AI ethics and safety research.

Algorithm

A set of step-by-step instructions that tells a computer how to solve a problem or get a task done. Think of it as the recipe behind every AI output.

Application

How an AI tool is actually used in the real world – like customer service chatbots, spam filters, or movie recommendations.

Application Programming Interface (API)

A set of rules that lets two programs talk to each other and share info without drama. The digital world's universal translator.

Artificial Intelligence (AI)

Machines or software that do tasks in a "smart" way, mimicking human decision-making without needing someone to micromanage every move. Fancy math turned into stochastic parrots.

Artificial Neural Networks

AI models built to work a bit like human brains, with layers of "neurons" that process information to recognize patterns, translate languages, or generate images.

Attention Mechanism

A trick inside AI models that helps them focus on the most relevant parts of input data. Like skimming an article and only reading the bolded lines to get the gist.

Auto-encoder

An AI model that learns to compress data into a smaller form and then rebuild it back again. Handy for things like reducing image file sizes without losing important details.

Automatic Prompt Engineer (APE)

Using AI to generate and test multiple prompts to find the most effective instruction for a task. The AI writes its own prompts so you don't have to guess.

Prompt Example: "Generate five different prompts to best extract action items from meeting notes."

Automatic Reasoning and Tool-Use (ART)

An AI framework where the model decides when to use tools, like calculators or web searches, while thinking through a problem. Basically, AI learning to Google mid-thought.

Prompt Example: "Calculate the square root of 547 using Wolfram Alpha, then explain how you got the result."

Back to Top

B.

Backpropagation

How neural networks learn – adjusting weights backwards from output to input to reduce errors. The secret sauce that actually makes deep learning work.

BERT (Bidirectional Encoder Representations from Transformers)

A language model that reads text in both directions to understand context better. Great for understanding what words actually mean in a sentence.

Bias

When an AI model's outputs are unfair or skewed, usually because it learned from biased data. Garbage in, biased garbage out.

Big Data

Data sets so massive and messy that normal tools can't handle them. AI thrives on them to find patterns humans would miss.

Build a Neural Network

Creating a brain-like AI model by designing its layers and nodes to tackle tasks like predicting handwritten numbers or classifying photos.

Prompt Example: "Build a simple neural network in Python to classify handwritten digits."

Back to Top

C.

Chain-of-Thought

A prompting technique where you tell the AI to show its reasoning step by step before giving a final answer, making it more accurate on tough problems.

Prompt Example: "Let's think through this step by step: What is the square root of 144?"

Chain-of-Verification (CoVe)

Prompting AI to draft an answer, generate questions to check itself, answer those, then produce a corrected final. AI double-checking its own homework.

Prompt Example: “Draft a definition for ‘deep learning,’ then create verification questions about it, answer them, and refine the original definition.”

Chatbot

A program that chats with people like it’s human, handling questions, support, or random small talk. Usually see these on web pages to handle company specific FAQs.

Cognitive Computing

AI designed to think a bit more like people do. Making sense of language, reasoning through problems, and figuring out decisions in messy real-world situations.

Convolutional Neural Network (CNN)

An AI model type especially good at recognizing images. It sees pixels in chunks to identify patterns, like spotting cats on the internet – endlessly.

Context Window

The limit on how much text an AI model can “remember” in a single prompt or conversation. If you go over it, older parts get forgotten like a goldfish.

CREATE Prompting

A structured prompting framework: Character, Request, Examples, Adjustments, Type of Output, Extras. Helps build clear prompts for consistent outputs.

Prompt Example: “Act as a career coach (Character). Review my resume (Request). Here are three bullets (Examples). Be concise and actionable (Adjustments). Format as a checklist (Type of Output). Include one improvement I didn’t mention (Extras).”

Back to Top

D.

Data Mining

Digging through massive piles of data to find hidden patterns or insights you can actually use to make decisions. Another great use of AI.

Decision Traceability

Being able to look under the hood and explain how an AI made its choice. Crucial for trust, accountability, and when bosses inevitably ask, “Why did it say that?”

Deep Learning

A type of machine learning where AI models use tons of layers to process data and spot patterns, like recognizing faces in photos or turning speech into text. Basically, it’s machine learning on steroids.

Directional Stimulus Prompting (DSP)

Adding an extra mini-prompt to steer the AI in a specific direction before the main prompt. Think of it as whispering hints to nudge the model’s style or focus.

Prompt Example: “Include these keywords in your response: ‘efficiency, automation, ROI.’ Now write a LinkedIn post about AI tools for marketers.”

Back to Top

E.

Embeddings

Turning words or items into number vectors so AI models can process them. Think of it as translating language into math coordinates for the model to work with.

Embeddings Search

Using word or sentence embeddings to find items that are similar in meaning. Powers smarter search functions in AI apps.

Prompt Example: “Given these product descriptions, find the one most similar to this new listing.”

Back to Top

F.

Few-Shot Chain-of-Thought Prompting

Combining few-shot examples with chain-of-thought reasoning. Gives AI sample reasoning paths to follow before tackling your new problem.

Prompt Example: “Here are two worked-out math examples showing steps. Now solve this new problem using the same approach.”

Few-Shot Learning

When an AI only needs a handful of examples to figure out what you want it to do, instead of a massive training set.

Prompt Example: “Here are three product descriptions. Write a fourth in the same style.”

Few-Shot Prompting

Using a small number of examples in your prompt (usually two to five) so the AI understands the task or style before responding. Zero-shot (no examples), one-shot (one example), and two-shot (two examples) are all specific prompting styles under this umbrella.

Prompt Example: “Here are four social media captions we’ve used before. Write a new one in the same voice and length.”

Fine-Tuning

Taking a pre-trained AI model and training it further on your specific data so it performs better on your tasks. Like teaching a trained dog your house rules.

Back to Top

G.

GAN (Generative Adversarial Network)

An AI system with two models playing creative tug-of-war – one tries to create realistic outputs, and the other tries to catch the fakes until it cant tell the difference. That’s how we get deepfake videos and AI art.

Generated Knowledge Prompting

Prompting the AI to generate background knowledge about a topic first, then use that to answer your actual question more accurately.

Prompt Example: “First, summarize what blockchain is. Then, explain how it could be used in supply chain management.”

Generative AI

AI that doesn’t just analyze. It makes new stuff. Text, images, music, you name it. Models like my personal favorite, ChatGPT.

Generative Pretrained Transformer (GPT)

ChatGPT's government name. Also, it's the underlying model that powers tools like ChatGPT. It predicts the next word in a sentence so well it feels like it’s thinking. (Spoiler: it’s not.)

Gradient Descent

An optimization method AI uses to adjust itself step by step until errors are minimized. Think of it as finding the lowest point in a hilly landscape by taking careful steps downhill.

Graphics Processing Unit (GPU)

A super-powered computer chip originally built for gaming graphics, now used to crunch AI data fast so your models don’t take all year to train.

Back to Top

H.

Hard Prompting

Giving the AI strict, direct instructions with no wiggle room. Works when you need consistency over creativity.

Prompt Example: “Write exactly three bullet points summarizing this article in a formal business tone.”

Hallucinations

When an AI confidently spits out something that sounds legit but is actually dead wrong. Basically, it's making stuff up without knowing it.

Human in the Loop

Keeping humans involved to oversee, correct, or guide AI outputs. Because letting AI run wild is how we eventually end up with Skynet… or, realistically, very bad customer service chats.

Hyperparameter Tuning

Adjusting the "settings knobs" of an AI model to improve how well it learns. Tedious but critical for top performance.

Back to Top

I.

Image Diffusion Models

AI models that generate images by starting with random noise and refining it step by step until you get a realistic picture. How tools like Midjourney or DALL-E work.

Inference

The moment when an AI model uses what it's learned to actually make a prediction or generate an output. Training is learning; inference is doing.

Back to Top

J.

Coming soon. (Maybe.)

Back to Top

K.

Coming soon. (Maybe.)

Back to Top

L.

Large Language Model (LLM)

A massive AI model trained on tons of text to understand and generate language, like ChatGPT or Claude.

Learning from Mistakes

The core of how AI models get better. By adjusting their calculations every time they mess up until they get it right. Basically, trial and error on super speed.

Least-to-Most Prompting

Starting with easy subtasks before tackling harder parts of a complex problem. Helps AI build up understanding gradually instead of diving in blind.

Prompt Example: “First, define what a server is. Then explain how it interacts with APIs in a web app.”

Back to Top

M.

Machine Learning (ML)

A type of AI where models learn from data instead of being explicitly programmed. They spot patterns and make predictions based on examples.

Massive Datasets

Huge collections of data that AI models train on to learn patterns, meanings, and connections. More data usually equals better performance. Eating the entire internet and asking for seconds.

Meta-Prompting

Prompting the AI to decide which prompting strategy it should use before tackling a task. It's like telling the AI to plan how it will plan.

Prompt Example: "Before answering, choose the best prompting strategy to solve this logic puzzle and explain why you chose it."

Model

The trained version of an AI algorithm. It's what you actually use after all the learning is done to generate outputs or make decisions. ChatGPT, Claude, Copilot, etc.

Multimodal AI

AI that can handle text, images, and audio all together. Like ChatGPT Vision analyzing a diagram and explaining it.

Prompt Example: "Here's a photo of my handwritten notes. Summarize them into bullet points."

Back to Top

N.

Natural Language Processing (NLP)

Teaching computers to understand and work with human language. Basically, helping them read, write, and talk back without sounding like a robot from the 80s.

Natural Language Generation (NLG)

When AI writes words that actually sound human. Blog posts, summaries, LinkedIn captions, it's all just predicting what the next most logical word would be.

Prompt Example: "Write a 150-word summary of this article in a friendly, conversational tone."

Natural Language Understanding (NLU)

AI trying to grasp what you mean, even if you're vague or cryptic. Like understanding "I'm fine" actually means "I'm exhausted, but carry on."

Neural Network

An AI model built with layers of connected “neurons” that process info, inspired by human brains but without the anxiety. We haven’t figured out how to add that yet.

Notebook

An interactive coding workspace where you write code, run it, and see what breaks, all in one place. Gold standard for data science experiments and people interested in making their own agents.

Back to Top

O.

One-Shot Prompting

Providing the AI with a single example before asking it to do the task. Gives it a hint without spelling everything out.

Outcome

The final thing you produce with AI in your session. The reason you opened ChatGPT in the first place.

Overfitting

When an AI model memorizes training data instead of learning general rules. Perfect grades in class, clueless in the real world. We all know someone like that.

Back to Top

P.

Predictive AI

AI that looks at what’s happened before to guess what comes next. Sales forecasts, stock predictions, or suggesting what snack you’ll buy at 3am.

Program-Aided Language Models (PAL)

Prompting AI to write code that solves part of the problem, then running that code for the final answer. Offloads logic to programming instead of pure text prediction.

Prompt Example: “Write a short Python function to calculate compound interest over ten years with a $5,000 initial investment at 5% annual return.”

Prompt

Anything you type into an AI’s text window to tell it what you want. A prompt can be a single word, a full page of instructions, or somewhere in between. It’s the question, command, or context you give to kick off the model’s response.

Prompt Example: “Explain what a neural network is in simple terms that an average 10-year-old could understand.”

Prompt Chaining

Using the output of one prompt as the input for the next to build complex workflows. Like stacking LEGO bricks to build a castle one section at a time. Happens naturally as you go back and forth with model in a thread.

Prompt Example: “Generate five blog title ideas for productivity hacks. After listing them, pick the one most likely to get high clicks on LinkedIn, and write an intro paragraph based on that.”

Prompt Engineering

Like I said, everything you type into a model’s text window is a prompt, but designing prompts to get the AI to actually do what you want instead of going rogue is prompt engineering. Honestly, this is the most fun part of AI learning for me.

Prompt Example: “Rewrite this email to sound confident but still approachable.”

Back to Top

Q.

Coming soon. (Maybe.)

Back to Top

R.

RAG (Retrieval-Augmented Generation)

AI that grabs info from elsewhere before writing an answer, so it isn’t just relying on its memory. Basically, giving it receipts to back up what it says.

Prompt Example: “Using the uploaded research PDF, summarize the key findings for a marketing strategy meeting.”

ReAct Prompting

Combining reasoning steps and tool-use actions in AI prompts. Think internal monologue plus doing the work.

Prompt Example: “Think step by step about solving this math problem, and use Wolfram Alpha if calculations are needed.”

Reasoning Engines

AI systems that don’t just spit out answers but actually think through problems step by step. Giving your AI a brain upgrade for logic-heavy tasks.

Recurrent Neural Network (RNN)

An AI model designed to handle sequences, feeding its output back into itself. Good for stuff like predicting the next word in a sentence or time series data.

Reflexion

Prompting AI to reflect on its last answer, critique mistakes, and try again for a better result. An iterative self-improvement loop without retraining.

Prompt Example: “Here is your last answer. Reflect on where it might be unclear or inaccurate, then rewrite it stronger.”

Reinforcement Learning

AI learning by trial and error, getting rewards for good moves and penalties for bad ones. Basically, training a dog with math instead of treats.

Responsible AI

Designing and using AI systems in ways that are ethical, fair, and safe. Making sure we don’t accidentally build biased, harmful, or downright evil robot overlords.

Back to Top

S.

Select the Best Algorithm

Choosing the right AI model for the job so you're not hammering in screws or using a blowtorch to toast bread.

Self-Ask Prompting

Asking AI a complex question and having it break it into smaller questions it answers one by one before giving a final result. Like watching it interview itself to solve a problem.

Prompt Example: "Why are housing prices so high? Break this into smaller questions and answer each before giving your conclusion."

Self-Consistency Decoding

Prompting AI to generate multiple reasoning chains and then pick the most consistent answer. More reliable than a single chain-of-thought.

Prompt Example: "Generate three different solutions to this logic puzzle and choose the most consistent one."

Self-Supervised Learning

AI figuring out patterns in unlabeled data all on its own. Solving a puzzle with no picture on the box and still getting it right.

Semantic Search

Searching using meaning instead of exact keywords. AI finds results based on what you meant, not just what you typed.

Prompt Example: "Find all customer reviews expressing frustration, even if they don't use the word 'frustrated.'"

Single-Shot Learning

When AI learns from just one example instead of thousands. Humans are still better at this for now.

Skeleton-of-Thought Prompting

Prompting AI to first outline a high-level skeleton before expanding each point. This forces it to plan before rambling.

Prompt Example: "Outline the main sections of a blog post about AI bias, then expand each into detailed paragraphs."

Soft Prompting

Using subtle cues or phrasing in your prompt to guide the AI's tone or approach without explicit instructions. Like hinting instead of commanding.

Prompt Example: "Could you maybe draft a quick summary of this in a friendlier tone."

Stochastic Parrots

A nickname for models like ChatGPT that just predict the next word based on patterns, not real understanding. They're parroting what they've seen, but with fancy math behind it.

Strong AI vs. Weak AI

Strong AI would think and feel like a human (sci-fi dream). Weak AI does one thing well, like Siri playing your Spotify or getting your requests half-wrong.

Supervised Learning

AI learning from data with answers provided upfront. Like giving students the answer key before the test so they can memorize what's needed.

Back to Top

T.

Temperature (Sampling Temperature)

A setting that controls how creative or risky AI outputs are. Low temperature (0.1) means safe and predictable; high (0.9) means wild and creative.

Tensor Processing Unit (TPU)

Special chips Google built to train AI models faster than regular GPUs. GPUs handle gaming graphics fast; TPUs handle AI math even faster.

The Activation Bias

A quirk in AI math where certain outputs are favored just because of how the calculations work, not because they're actually the best answer.

The Four Ds of Automation

Tasks that are Dirty, Dull, Dangerous, or Difficult. If it sucks for humans, chances are we're building AI or robots to do it for us.

The Generative AI Revolution

The current wave where AI isn't just analyzing data but creating new stuff like text, images, and music. It's changing how we work, whether we're ready or not. I keep saying this, but we're at the same point with AI now that we were with the internet in the 90s.

Tokens

Chunks of text AI models process instead of actual words on the screen. Not always whole words; sometimes parts of words. Everything you type gets split into tokens before AI reads it.

Top-p (Nucleus Sampling)

Another setting controlling output randomness by limiting choices to the top percentage of likely words. Adjusting it fine-tunes creativity vs coherence.

Training Data

The information we feed into an AI model to teach it what's what. Good data makes good models. Again, garbage data makes biased garbage models.

Transformer Architecture

The AI model design behind ChatGPT and BERT. Uses attention mechanisms to handle long text inputs efficiently. Basically, it made everything else look outdated overnight.

Tree-of-Thoughts Prompting

Getting AI to explore multiple solution paths before choosing the best one. Like having three experts brainstorm routes to the answer.

Prompt Example: "Imagine three experts each tackle this question step by step. Compare their answers and choose the best."

Two-Shot Prompting

Providing two examples to guide the AI before asking for the final output. More context, better odds of getting what you actually want.

Back to Top

U.

Unsupervised Learning

AI finding patterns in data without any labels or guidance. Like being dumped in a new city with no map and figuring out where everything is on your own, but probably way faster.

Using AI Systems

The practical side of deploying AI models to solve real-world problems. Less about the fancy math, more about making stuff actually work in business.

Back to Top

V.

VAE (Variational Autoencoder)

A type of AI model that compresses data and then recreates it with small variations. Used for things like generating slightly different images from the same concept.

Back to Top

W.

Coming soon. (Maybe.)

Back to Top

X.

Coming soon. (Maybe.)

Back to Top

Y.

Coming soon. (Maybe.)

Back to Top

Z.

Zero-Shot Chain-of-Thought

Prompting the AI to reason step by step without giving it any examples first. Adding “Let’s think step by step” is the classic trigger phrase here.

Prompt Example: “Let’s think step by step: What causes inflation to rise in an economy?”

Zero-Shot Prompting

Giving the AI a task with no examples and expecting it to figure it out based on what it already knows. Risky, but sometimes it nails it. Basically, faking it convincingly the first time around.

Back to Top

Futureproof Directive | https://futureproofdirective.com

Want more practical, plain-English AI insights like this? Visit the blog and join the mailing list to stay *Futureproof*.