

FREE GUIDE · CONSTRUCTION EDITION

# How AI works, in plain English.

A short, visual guide for business owners and leadership teams. Part 1 covers what AI is and how businesses use it. Part 2 explains how it actually works, so you can use it well and know when not to trust it.



PART 1

# The big picture


What AI is, what it's good at, and how a business actually puts it to work.

- 
- 01 What AI is
  - 02 What it's good and bad at
  - 03 Three ways businesses use it
  - 04 Start with a workflow
  - 05 What it costs
  - 06 A 30-day pilot
  - 07 Protecting your data
-

# 01

## "AI" is a family of tools, not one thing.

Most of what people mean by AI today is a language model: software that has learned the patterns of written language well enough to read, write and reason in it. It sits inside a much older field.



## 02

# It's very good at some work and unreliable at other work.

The pattern is simple. Language models are strong with words and patterns, and weak wherever exactness or facts they weren't given matter.

---

### STRONG

---

Drafting and rewriting

---

Summarizing long documents

---

Pulling details out of messy files

---

Comparing documents side by side

---

Answering from sources you provide

---

### CHECK CAREFULLY

---

Exact math across large data

---

Facts it wasn't given

---

Recent events

---

Judgment calls about people

---

Anything you can't verify

# 03

## Businesses use it in three ways, each with more independence.

An assistant answers when you ask. An automation runs the same step every time. An agent works toward a goal and decides its own next steps. More independence means more value and more need for a person to sign off.

### ASSISTANT



### AUTOMATION




### AGENT



# 04

## Start with a workflow, not a tool.

Map one job your team does often, step by step, and mark the steps where AI could take on the reading, sorting or first draft. People keep the decisions. Here's bid leveling on a construction job.



PART 1 · THE BIG PICTURE

# 05

## The software is only part of the cost.

There are three kinds of cost to plan for. The third is the one most budgets leave out.

### Setup and your team's time

Mapping the work, connecting your systems, training people and checking output

### Usage

Charged per token when you build with AI directly, so the bill grows with how much you run

### Seats

A monthly subscription per person for tools like ChatGPT, Claude or Copilot

# 06

## Prove it with a 30-day pilot.

Pick one workflow, measure how it runs today, and give a small group a month with the new version. Training sticks when it's built around the work people already do.



# 07

## Know where your data goes before you paste it in.

A few habits cover most of the risk, and none of them require a technical team.



### **Use business plans**

Business and enterprise plans usually keep your data out of model training. Free and personal plans may not.



### **Decide what never goes in**

Write down the data your team shouldn't paste into any AI tool, such as client financials or personal records.



### **Keep a list of your tools**

Know which AI tools your team uses and what each one can see. A one-page policy is enough to start.

---

## PART 2

# How AI actually works

You don't need to be technical to use AI well. You do need a working picture of what's happening inside, because it explains both what AI is good at and where it goes wrong.

- 
- 08 Tokens
  - 09 Predicting the next word
  - 10 How a model is trained
  - 11 The context window
  - 12 Why it makes things up
  - 13 Answering from your documents
  - 14 Tools and agents
  - 15 Giving good instructions
  - 16 Checking the work
-

# 08

## It reads and writes in small pieces called tokens.

A model doesn't see words the way you do. It breaks text into tokens, which can be a whole short word, part of a longer word, or a punctuation mark. Tokens are also how usage is measured and billed.

TYPE SOMETHING



11 tokens in this illustration (· marks a space)

An illustration. Every model splits text its own way, so real token counts differ.

# 09

## It writes by predicting the next token, over and over.

At each step the model weighs every possible next token, picks a likely one, adds it, and repeats. Fluent writing comes out of that loop. So does a confident sentence that isn't true.

## The estimator checked the



An illustration. The bars show relative likelihood, not real figures.

# 10

## It learned from a vast amount of text, then was taught to be helpful.

First the model reads an enormous amount of writing and learns how language fits together. Then people show it good answers and rate its attempts, which shapes it into an assistant that follows instructions.

- 
- **Pretraining**  
The model reads a huge range of text and learns the patterns of language.
  - **Fine-tuning**  
It learns from examples of good questions paired with good answers.
  - **Human feedback**  
People rate its answers, and it learns which kinds of answer they prefer.
  - **The assistant you use**  
Its knowledge stops at a cutoff date, unless a tool such as web search gives it newer information.
-

# 11

## It only knows what's in front of it right now.

Everything a model can use in a conversation sits in its context window: your messages, any files you add, and its own replies. Today's windows are very large. Anthropic's current models hold up to a million tokens, which it puts at roughly 555,000 words, enough for a whole stack of contracts, specs and subcontractor proposals at once.

Anything outside the window, it doesn't see: last week's chat, unless the product saves it, or the oldest part of a long one once it's dropped or summarized.

Even inside, research has found models use the start and end more reliably than the middle. Position matters, not age: it rereads the whole window each time.

OUT OF THE WINDOW



Last week's chat, or the oldest part of a long one. Not seen.

CONTEXT WINDOW



Your messages, your files and its replies, all at once

- ☒ Start and end: used most reliably
- ☐ Middle: easier to miss

# 12

## When it doesn't know, it can still sound sure.

The model is built to produce likely text, so a gap in what it knows gets filled with something that sounds right. This is often called hallucination. It's why you give it sources and check anything that matters.

What was the change-order total on the Lone Pine job?

The Lone Pine job had four approved change orders.

They totaled \$48,200, mostly for added site lighting.

Sounds right. Isn't in any document it was given.

The largest was approved in March.

An invented example, to show the pattern.

# 13

## Give it your documents, and it answers from them.

Many business tools first search your files for the passages that matter, then hand those to the model along with your question. This is called retrieval, or RAG. Answers get more accurate, and they can point to where each fact came from.

---

○ You ask a question

---

● The tool searches your files

Contracts, specs and past bids

---

● It pulls out the passages that answer the question

---

● The model writes an answer using only those passages

---

● You get the answer with links to its sources


So you can check each fact yourself

---

# 14

## With tools, an agent can do the work itself and report back.

On its own, a model can only tell you what to do. Connect it to tools such as search, a spreadsheet or your email, and it can take a step, look at the result and decide what to do next. That loop is what people mean by an agent. The more it can do on its own, the more a person should approve before anything goes out.



15

Better instructions get better work.

Treat it like a capable new hire on day one. Tell it the situation, the job, what good looks like and the shape you want back.

|           |                                                                                      |
|-----------|--------------------------------------------------------------------------------------|
| CONTEXT   | I run a 40-person general contractor. This is a subcontractor's electrical proposal. |
| TASK      | List everything it excludes or doesn't mention compared with our scope sheet.        |
| FORMAT    | A table: item, included or excluded, page number.                                    |
| GUARDRAIL | If you aren't sure, say so. Don't guess.                                             |

# 16

## Build the check into the workflow.

AI output is a first draft. Decide in advance who reviews it, what they compare it against, and who signs off. Having a second AI review the first helps, but it doesn't replace a person for anything that matters.

|                                                                                       |          |                                               |
|---------------------------------------------------------------------------------------|----------|-----------------------------------------------|
|    | AI       | Writes a first draft and cites its sources    |
|    | REVIEWER | Checks the draft against those sources        |
|    | REVIEWER | Corrects it, or sends it back with notes      |
|  | SIGN-OFF | A named person approves it before it goes out |

## What to remember.

---

It predicts likely text, so fluent isn't the same as true.

---

It only knows what's in its context window, so give it the sources.

---

Start from one workflow, and keep people on the decisions.

---

Budget for setup and your team's time, not just seats.

---

Use business plans and decide what data never goes in.

---

Build the review step in before anything goes out.



## Want help applying this to your business?

We'll set up 30 minutes to talk through where  
AI could help and roughly what it would cost.

[simpleraven.com](https://simpleraven.com)

[hello@simpleraven.com](mailto:hello@simpleraven.com)

