



FREE INGRID AI ASSISTANT GUIDE

Build Your Own AI System

A practical beginner's guide to turning your knowledge, documents, and workflows into an evidence-driven AI assistant.

RAG · Evidence · TruthMode · Confidence · Citations · Human Control

Evidence First. Intelligence Second. Humans Decide.

INTRO. From Chatbot to AI System

Most people start with AI by asking questions. That is useful, but a practical AI system can do much more.

A custom AI system can use **your information, your documents, your workflows, and your rules** to help answer questions, find knowledge, compare evidence, and support decisions.

The goal of this guide

You do not need to train your own model. You can start with an existing AI model and surround it with the right knowledge, retrieval, validation, confidence, citations, permissions, and human oversight.

The seven building blocks

1	AI model	Reasoning + language
2	Knowledge base	Your approved information
3	RAG	Retrieves relevant evidence
4	Truth layer	Checks evidence + conflicts
5	Confidence	Communicates uncertainty
6	Citations	Shows the source
7	Human control	Keeps important decisions with people

1. Start With the Problem

Do not start with the AI. Start with the work that needs to improve.

Do not begin by asking, **“What AI should I use?”** Begin with, **“What problem should AI solve?”**

- Searching through documents or folders
- Answering the same questions repeatedly
- Reviewing large amounts of information
- Finding policies, procedures, products, or services
- Comparing information from multiple sources
- Preparing reports or summaries
- Supporting customer or employee questions

Better problem statement

Instead of “I want an AI chatbot,” try: “My team spends too much time searching company documents to answer customer questions.”

Your first use case

Where is your team losing time?

What question gets answered over and over?

What information is difficult to find?

What decision would improve with better evidence?

2. Identify Your Knowledge

Your organization probably already has the intelligence your AI needs. The first job is to identify where that knowledge lives.

- Documents and PDFs
- Web pages
- Policies and manuals
- Products and services
- FAQs
- Research
- Spreadsheets
- Project notes
- Internal procedures
- Approved customer information

Important principle

Your AI does not need to know everything. It needs access to the **right information**.

Starter exercise: list five sources your AI should know

1. _____
2. _____
3. _____
4. _____
5. _____

3. Build a Knowledge Base

A knowledge base is the organized information your AI can search when someone asks a question. The cleaner and more current that information is, the better the system can perform.

A simple company structure might look like this:

COMPANY KNOWLEDGE
Products · Services · Policies
Customers · Research · Procedures
FAQs · Projects · Internal Knowledge

Good knowledge-base habits

- Use approved sources.
- Remove obvious duplicates.
- Keep dates and versions when they matter.
- Separate current policies from old ones.
- Use clear titles and categories.
- Decide who is allowed to update the source material.

Better information creates better AI.


A powerful model cannot compensate for a knowledge base full of outdated, contradictory, or poorly organized information.

4. Add RAG


Retrieval-Augmented Generation

RAG is the layer that lets the system retrieve relevant information before the AI answers.

Without RAG



With RAG



The shift

RAG moves the system away from “I think the answer is...” and toward “Here is what your information says.”

5. Make the AI Show Its Work

A useful business AI should not simply produce confident-looking answers. Important factual claims should be connected to evidence whenever possible.

- Source document
- Citation or passage
- Relevant link
- Evidence ID
- Date or version
- Data source

Example

Instead of only saying “Employees receive 15 vacation days,” the system can add: **Source: Employee Handbook - Vacation Policy.**

Add a truth layer

Retrieval is not enough. The system should evaluate the evidence it found.

- Is the evidence relevant?
- Is it current?
- Do multiple sources agree?
- Are there contradictions?
- Is important information missing?

PASS	Evidence strongly supports the answer.
REVIEW	Evidence exists, but uncertainty or conflict remains.
FAIL	There is not enough reliable evidence to support the answer.

6. Communicate Confidence - Keep Humans in Control

AI should not sound equally certain about everything. Confidence should reflect the quality and consistency of the evidence.

HIGH	Strong evidence supports the answer.
MEDIUM	Useful evidence exists, but there are limitations.
LOW	Evidence is weak, missing, outdated, or conflicting.

Define what AI may do and what still requires a person.

AI MAY	HUMAN APPROVAL
✓ Find company policies	✓ Send or sign contracts
✓ Summarize documents	✓ Approve spending
✓ Answer common questions	✓ Make employment decisions
✓ Identify conflicting information	✓ Publish sensitive information
	✓ Execute major decisions

Evidence First. Intelligence Second. Humans Decide.

7. Give Your AI Clear Instructions

Your system needs an operating policy. A good policy tells the AI where answers should come from, what to do when evidence conflicts, and when to admit uncertainty.

Starter system instruction


Answer questions using the approved information available to you. Retrieve relevant evidence before answering factual questions. Clearly distinguish facts, assumptions, and recommendations. Identify conflicting or missing information. Cite important sources when possible. If the evidence does not support an answer, say that you do not have enough information. Do not invent facts. Important decisions remain under human control.

Add rules for your specific use case

- Which sources are approved?
- Which topics are out of scope?
- What must always be cited?
- When should the system return REVIEW instead of an answer?
- What actions require human approval?
- Who may access sensitive information?

8. Design the AI Workflow

A simple evidence-driven workflow can look like this:



9. Test the System

Never assume the system works because it produced one good answer. Test behavior deliberately.

KNOWN ANSWER TEST

Ask a question clearly answered by your approved sources.

MISSING INFORMATION TEST

Ask something your knowledge base does not contain. The AI should admit that it does not know.

CONFLICT TEST

Give the system two sources with different information and see whether it detects the conflict.

CITATION TEST

Check that the cited source actually supports the answer.

HALLUCINATION TEST

Ask a question designed to tempt the AI into inventing information.

10. Start Small - Build Your First System

You do not need to build a giant AI platform on day one. Start with one useful system and expand after it works.

ONE PROBLEM · ONE KNOWLEDGE BASE · ONE AI INTERFACE · ONE WORKFLOW · ONE GROUP OF USERS

What problem will the AI solve?

Who will use it?

What information does it need?

What questions should it answer?

What should it never do?

What actions require human approval?

What evidence should it show?

How will you know it is working?

11. Your Starter Architecture

You now have the core pieces of a practical custom AI system:

AI MODEL	Reasoning and language
KNOWLEDGE BASE	Your approved information
RAG	Finds relevant evidence
TRUTH / VALIDATION	Checks relevance, conflicts, and gaps
CONFIDENCE	Communicates uncertainty
CITATIONS	Shows where information came from
HUMAN CONTROL	Keeps important decisions with people

YOUR CUSTOM AI SYSTEM

NEXT. Build Something Useful

Building AI does not always mean building a new model. You can take existing AI technology and surround it with **your knowledge, your evidence, your rules, your workflows, your permissions, and your human oversight.**

Do not begin with “How much AI can we add?” Begin with:

- Where are people losing time?
- Where is information difficult to find?
- Where are questions being answered repeatedly?
- Where would better evidence improve a decision?

Your next move

Choose **one** real problem. Gather the approved information. Define the rules. Build a small retrieval-and-answer workflow. Test it with known answers. Improve it from evidence.

INGRID AI ASSISTANT

Custom AI That Knows Your Business - and Shows Its Work.

RAG · Evidence · TruthMode · Confidence · Citations · Human Control

Evidence First. Intelligence Second. Humans Decide.

Ingrid Verified AI by Harmony Fusion Records

signal@harmonyfusionrecords.com

BONUS. 75 AI Response Prompts

A SIMPLE RESPONSE-MODE LIBRARY FOR CHATGPT + INGRID

Use these modes to tell AI how you want it to think, analyze, write, or respond. Type the mode name first, then your request.

HOW TO USE

TRUTHMODE Is this plan actually working?

REDTEAM Find the weaknesses in this system.

80/20 What should I focus on first?

You can also combine modes when you want both approaches.

Example: TRUTHMODE + REDTEAM Evaluate this business idea.

Core rule

A mode changes the response style and reasoning objective. It does not give the AI permission to invent facts, hide uncertainty, or ignore evidence.

Evidence First. Intelligence Second. Humans Decide.

INGRID AI ASSISTANT

75 response prompts for thinking, building, writing, learning, and decision-making

CORE MODES

INGRID 75-PROMPT RESPONSE LIBRARY

01 TRUTHMODE

Give a direct, evidence-grounded assessment without softening the conclusion.

02 ELI10

Explain the subject in simple, intelligent language.

03 REDTEAM

Look for weaknesses, risks, blind spots, and failure modes.

04 HORMOZI

Analyze value, leverage, execution, offer strength, and business outcomes.

05 /HUMAN

Write naturally, conversationally, and directly.

06 FUTUREYOU

Evaluate the decision from the perspective of your future self.

07 UNLEARN

Challenge assumptions and rebuild the problem from first principles.

08 80/20

Identify the few actions likely to produce most of the results.

09 SOCRATES

Test assumptions with thoughtful questions before concluding.

10 LINDYMODE

Favor durable principles and long-term resilience over novelty.

REASONING & ANALYSIS

INGRID 75-PROMPT RESPONSE LIBRARY

11 FIRSTPRINCIPLES

Reduce the problem to fundamental truths and rebuild from there.

12 ROOTCAUSE

Find the underlying cause instead of treating symptoms.

13 SECONDORDER

Examine what happens after the immediate consequence.

14 SYSTEMSTHINK

Analyze how all parts of a system interact.

15 CONTRARIAN

Build the strongest credible argument against the obvious answer.

16 DEVILSADVOCATE

Deliberately challenge the proposed idea.

17 ASSUMPTIONCHECK

Identify assumptions that may be driving the conclusion.

18 LOGICCHECK

Examine an argument for logical gaps or contradictions.

19 CAUSEEFFECT

Separate correlation from actual causation.

20 TRADEOFFMODE

Show what is gained and lost with each choice.

EVIDENCE & VERIFICATION

INGRID 75-PROMPT RESPONSE LIBRARY

21 **RAGMODE**

Retrieve relevant evidence before answering.

22 **EVIDENCEMODE**

Separate supported claims from unsupported claims.

23 **CONFIDENCEMODE**

State how strong the evidence is behind the conclusion.

24 **SOURCECHECK**

Evaluate the quality and authority of the sources.

25 **CONFLICTCHECK**

Identify conflicting information.

26 **FACTCHECK**

Verify factual claims.

27 **GAPCHECK**

Identify missing evidence or information.

28 **BIASCHECK**

Look for bias in sources, assumptions, or reasoning.

29 **VERIFYFIRST**

Do not conclude until important claims have been checked.

30 **AUDITMODE**

Review the evidence trail and reasoning step by step.

BUSINESS

INGRID 75-PROMPT RESPONSE LIBRARY

31 **CEOMODE**

Think like the person responsible for the whole business.

32 **FOUNDERMODE**

Evaluate the problem from a founder perspective.

33 **INVESTORMODE**

Analyze opportunity, risk, return, and scalability.

34 **CUSTOMERMODE**

View the offer entirely from the customer perspective.

35 **COMPETITORMODE**

Analyze how a strong competitor would respond.

36 **REVENUEMODE**

Focus on how the idea produces revenue.

37 **MARGINMODE**

Focus on profitability instead of just sales.

38 **SCALEMODE**

Determine whether the idea can grow efficiently.

39 **MOATMODE**

Identify defensible competitive advantages.

40 **OFFERMODE**

Strengthen the product, service, price, and value proposition.

STRATEGY & DECISIONS

INGRID 75-PROMPT RESPONSE LIBRARY

41 STRATEGYMODE

Compare possible paths and recommend the strongest.

42 DECISIONMODE

Turn analysis into a clear decision.

43 PRIORITYMODE

Rank what deserves attention first.

44 OPTIONMODE

Generate multiple realistic paths forward.

45 SCENARIOMODE

Analyze best-case, expected, and worst-case outcomes.

46 RISKREWARD

Compare potential upside with downside.

47 OPPORTUNITYCOST

Identify what is sacrificed by choosing one path.

48 BOTTLENECK

Find the constraint preventing progress.

49 NEXTMOVE

Identify the most useful immediate action.

50 WARROOM

Analyze an urgent problem quickly and strategically.

BUILDING & TECHNOLOGY

INGRID 75-PROMPT RESPONSE LIBRARY

51 ARCHITECTMODE

Design the complete system and its components.

52 BUILDERMODE

Move directly from concept to implementation.

53 DEBUGMODE

Diagnose the cause of a technical problem before changing anything.

54 AGENTMODE

Design an AI agent with goals, tools, memory, and controls.

55 AUTOMATIONMODE

Identify repetitive work that can be automated.

56 WORKFLOWMODE

Turn a process into a repeatable workflow.

57 DATAMODE

Examine the information, structure, and data requirements.

58 APIMODE

Think in terms of connections between systems and services.

59 SECURITYMODE

Examine permissions, threats, privacy, and access controls.

60 FAILSAFEMODE

Design how the system should behave when something goes wrong.

WRITING & COMMUNICATION

INGRID 75-PROMPT RESPONSE LIBRARY

61 EDITOR

Improve clarity, structure, grammar, and flow.

62 COPYWRITER

Write persuasive marketing copy.

63 STORYMODE

Communicate the idea through narrative.

64 SCREENWRITER

Write scenes, dialogue, and cinematic storytelling.

65 PITCHMODE

Turn the idea into a concise persuasive pitch.

66 EXECUTIVEBRIEF

Reduce complex material into decision-ready information.

67 EMAILMODE

Write a clear, natural email.

68 SOCIALMODE

Turn an idea into social-media-ready content.

69 HEADLINEMODE

Generate strong titles, hooks, and headlines.

70 TONECHECK

Evaluate whether the message sounds appropriate for its audience.

LEARNING & CREATIVITY

INGRID 75-PROMPT RESPONSE LIBRARY

71 **TEACHERMODE**

Teach the concept step by step.

72 **QUIZMODE**

Test understanding with questions.

73 **ANALOGYMODE**

Explain complicated ideas using comparisons.

74 **BRAINSTORM**

Generate many possible ideas without judging them too early.

75 **INNOVATIONMODE**

Combine ideas in new ways to create something different.
