AI FOR IMPACT

boulevardlabs | JMU

UNDERSTANDING & USING AI WITH CONFIDENCE.

AI Fluency Field Guide

A practical foundation for understanding, evaluating, and working effectively with modern AI.

UNDERSTAND • PRACTICE • COLLABORATE • ASSESS

AI for Impact | BoulevardLabs + James Madison University

CONTENTS

What is inside this field guide

01	The five capabilities behind AI fluency	4
02	What modern AI is actually doing	5
03	AI vocabulary you should actually know	6
04	Where AI tends to help—and where it needs caution	8
05	Give AI a better job to do	10
06	Treat the first output as a starting point	12
07	Confidence is not accuracy	14
08	The human part still matters	15
09	Use AI with boundaries	16
10	Choose the right role for AI	17
11	Build skills that outlast the tool	18
12	AI fluency quick reference	19
NEXT	Move from understanding to practice	20
REF	Further reading	21

HOW TO USE IT

Read front to back for a practical foundation, or jump directly to the section that matches the AI question in front of you.

START HERE

AI fluency is more than knowing how to prompt.

AI tools are getting easier to access. The harder question is whether people understand them well enough to use them effectively, question them appropriately, and adapt as the technology changes.

This field guide is designed to build that foundation. It is not a catalog of today's "best" AI tools, a list of magic prompts, or a technical manual. It focuses on durable concepts and habits that remain useful even when the model, interface, or vendor changes.

THE CORE IDEA

AI fluency means understanding what AI is doing well enough to use it intentionally, evaluate its outputs, recognize its limits, and know where human judgment still matters.

Use this guide to:

- build a practical vocabulary for modern AI;
- understand what generative AI can and cannot reliably do;
- improve the way you frame tasks and provide context;
- recognize when AI output needs verification or human review;
- separate useful capability from impressive demos and hype;
- develop habits that transfer across unfamiliar AI tools.

What this guide is not

Not a governance policy. Your organization may have specific rules for approved tools, data handling, privacy, security, records, or regulated work.

Not a substitute for expertise. AI can assist subject-matter experts; it does not remove the need for expertise in consequential work.

Not a one-time certification. Fluency grows through practice and needs to evolve as AI capabilities change.

USE YOUR ORGANIZATION'S RULES

Before entering sensitive, proprietary, regulated, personal, or client information into any AI system, follow your organization's approved-tool, data-handling, privacy, and security requirements.

01 The five capabilities behind AI fluency

You do not need to become an AI engineer to become AI-fluent. You do need a repeatable way to think about the technology and your role in using it.

CAPABILITY	WHAT IT MEANS	QUESTION TO ASK
UNDERSTAND	Know what AI is, what different systems do, and where common limitations come from.	*"What am I actually working with?"*
EVALUATE	Judge whether an AI tool, approach, or output fits the problem and deserves trust.	*"Is this useful, appropriate, and supported?"*
APPLY	Frame tasks, provide context, set constraints, and use AI where it adds real value.	*"How should I use it here?"*
VERIFY	Check important claims, sources, assumptions, calculations, and omissions before relying on them.	*"How do I know?"*
ADAPT	Transfer your understanding to new models, tools, interfaces, and workflows as they change.	*"Can I work effectively with what comes next?"*

These capabilities reinforce one another.

Understanding without application stays theoretical. Application without verification creates risk. Verification without context can miss what matters. Adaptability keeps the whole skill set useful when the software changes.

A USEFUL DISTINCTION

AI access is having a tool. AI experience is having used it. AI fluency is understanding enough to use it well across changing situations.

02 What modern AI is actually doing

The easiest way to misuse AI is to treat it like a technology it is not. Generative AI, search, databases, automation, and analytics can overlap in a product, but they do not work the same way.

Generative AI in plain English

Generative AI is a class of AI designed to produce new content—such as text, images, audio, video, code, or other digital outputs—based on patterns learned from data. A language model does not simply look up a fixed answer. It generates a response based on the input, context, instructions, and patterns encoded in the model.

IMPORTANT CONSEQUENCE

A response can be fluent, detailed, and convincing without being factually correct. The quality of the writing is not proof of the quality of the underlying information.

Model vs. tool

The model is the underlying AI system. The tool is the product or application built around one or more models. A tool can add interfaces, data connections, retrieval, workflow steps, permissions, memory features, safety controls, or organization-specific functions. Two products can feel very different even when they rely on similar underlying models.

Generated vs. retrieved

When a model generates a response from learned patterns, that is different from retrieving a specific fact from a known source. Some AI products combine both approaches. This distinction matters whenever evidence, traceability, or current information matters.

KEEP THIS MENTAL MODEL

AI is not a single thing. “The AI” you interact with may be a model plus retrieval, search, tools, memory, data connectors, workflows, permissions, and other software around it.

03 AI vocabulary you should actually know

You do not need a dictionary full of technical terms. These concepts are worth knowing because they change how you use, evaluate, and talk about AI at work.

Hallucination

PLAIN ENGLISH An AI-generated claim, reference, detail, or answer that is incorrect, unsupported, or fabricated even though it may sound plausible.

WHY IT MATTERS Hallucinations are a reminder that confidence and accuracy are different. High-stakes facts and citations may need independent verification.

KEEP IN MIND A hallucination is not intentional deception. It is a failure mode of generative systems.

Context window

PLAIN ENGLISH The amount of information a model can work with during a given interaction or processing window.

WHY IT MATTERS Uploading a large document does not guarantee that every detail will be interpreted, prioritized, or used the way you expect.

KEEP IN MIND A context window is not the same thing as long-term memory.

Training data

PLAIN ENGLISH Information used during model development to help the model learn statistical patterns, relationships, language, and behavior.

WHY IT MATTERS An answer generated from learned patterns is not automatically traceable to one source document.

KEEP IN MIND Generated output should not be treated as retrieved evidence unless the system actually provides and supports the source.

RAG

PLAIN ENGLISH Retrieval-Augmented Generation: a method that retrieves relevant information from selected sources and supplies it to a model when generating an answer.

WHY IT MATTERS RAG can ground AI output in approved organizational documents or other knowledge sources without retraining the model.

KEEP IN MIND RAG can improve relevance, but it does not guarantee accuracy. Retrieval quality and source quality still matter.

Multimodal

PLAIN ENGLISH AI that can work across more than one kind of information, such as text, images, audio, video, or documents.

WHY IT MATTERS AI use is no longer limited to typing into a chat box. A single workflow may combine several forms of information.

KEEP IN MIND Capabilities vary. Not every multimodal model supports every media type equally well.

AI agent

PLAIN ENGLISH An AI-enabled system designed to pursue a goal through multiple steps, often using tools, information, and actions along the way.

WHY IT MATTERS The more a system can act—not just answer—the more important permissions, boundaries, oversight, and accountability become.

KEEP IN MIND Agents vary widely in autonomy. Many operate inside human-defined workflows and controls.

Fine-tuning

PLAIN ENGLISH Additional training used to adjust an existing model toward particular tasks, behaviors, styles, or domains.

WHY IT MATTERS Fine-tuning is one way to specialize AI, but it is different from providing instructions, adding context, or retrieving documents.

KEEP IN MIND Customizing an AI experience does not necessarily mean the underlying model was fine-tuned.

Prompt

PLAIN ENGLISH The instructions, questions, context, examples, constraints, or other input you give an AI system.

WHY IT MATTERS Prompt quality influences output, but prompting is only one part of AI fluency. Task fit, context, review, verification, and judgment matter too.

KEEP IN MIND There is rarely one “perfect prompt.” Effective use is often iterative.

VOCABULARY TAKEAWAY

Knowing the term is useful. Knowing what changes because of the term is fluency.

04 Where AI tends to help—and where it needs caution

AI is most useful when the task matches the capability. The goal is not to use AI everywhere. The goal is to recognize where it can contribute without quietly taking over work that still requires people.

TASK	AI CAN CONTRIBUTE	HUMAN ATTENTION STILL NEEDED
Drafting	First drafts, outlines, alternative phrasings, ideation	Accuracy, tone, audience fit, final standards
Summarizing	Reducing volume, extracting themes, reorganizing information	Omissions, priority, nuance, source fidelity
Transforming	Reformatting, translating structure, adapting for audiences	Meaning drift, lost caveats, inappropriate simplification
Comparing	Surfacing similarities, differences, patterns, tradeoffs	Incomplete criteria, hidden assumptions, context
Classification	Sorting, tagging, categorizing, routing	Edge cases, bias, ambiguous categories, quality control
Pattern finding	Trends, anomalies, relationships across large information sets	Causation, significance, real-world interpretation
Brainstorming	Possibilities, alternatives, questions, scenarios	Originality, feasibility, relevance, decision ownership
Decision support	Organizing options and evidence, scenario exploration	Values, consequences, accountability, missing information

Common warning signs

- **The task has a single exact answer.** Use an authoritative source or deterministic system when precision and traceability matter more than generation.
- **The stakes are high.** Legal, medical, safety, financial, personnel, eligibility, or similarly consequential work may require qualified review and stricter controls.
- **The information is sensitive.** Do not assume an AI tool is approved for confidential, regulated, proprietary, client, or personal data.
- **The output will be acted on automatically.** Errors can compound when AI moves from suggesting to taking actions.
- **Nobody can explain what "good" looks like.** If the team cannot define quality, it will be hard to judge whether AI improved anything.

A SIMPLE RULE

Use AI where it can create leverage. Keep humans where context, standards, accountability, and consequence matter.

05 Give AI a better job to do

Strong AI use often starts before the prompt. The quality of the task definition, context, and success criteria can matter as much as the wording of the request.

The five-part task brief

1	OBJECTIVE	What are you trying to accomplish? State the outcome, not just the action.
2	CONTEXT	What background does the AI need to understand the situation, audience, source material, or constraints?
3	TASK	What specifically should the AI do: summarize, compare, draft, classify, critique, transform, brainstorm, or analyze?
4	CONSTRAINTS	What limits matter: length, tone, format, approved sources, exclusions, deadlines, assumptions, or boundaries?
5	SUCCESS CRITERIA	What would make the output useful? Accuracy, completeness, audience fit, traceability, clarity, originality, or something else?

EXAMPLE

WEAK *"Summarize this report."*

STRONGER "Summarize this report for a senior leader who has five minutes. Prioritize decisions, risks, deadlines, and unresolved issues. Preserve uncertainty and do not invent missing details. Use five bullets and a short 'what needs attention' section."

Why context changes output

AI does not automatically know your organization's history, priorities, internal definitions, audience expectations, or exceptions. If those details matter, you may need to provide them—when appropriate and permitted. Better context can improve relevance, but more context is not automatically better context. Structure and selectivity matter.

BOUNDARY CHECK

Before adding context, ask: Is this information appropriate to enter into this specific tool under our organization's rules?

06 Treat the first output as a starting point

One of the most useful AI habits is learning to diagnose an output instead of simply accepting it or starting over.

IF THE OUTPUT IS…	TRY THIS NEXT
Too generic	Add specific audience, goal, examples, or source material.
Wrong direction	Correct the assumption or restate the objective.
Too long	Set a length constraint and identify what must be preserved.
Missing something	Ask what may be omitted, then compare against the source yourself.
Weak reasoning	Ask for assumptions, tradeoffs, counterarguments, or evidence.
Good but not ready	Ask it to critique the draft against explicit quality criteria.
Conflicting outputs	Compare the reasoning, source support, context, and assumptions—not just the wording.

A useful collaboration loop

1	**ASK**	Frame the task and provide the right context.
2	**REVIEW**	Read the output critically. Do not outsource attention.
3	**REFINE**	Correct direction, add constraints, or improve context.
4	**CHALLENGE**	Ask for weaknesses, assumptions, alternatives, or counterarguments.
5	**VERIFY**	Check claims, sources, calculations, and omissions that matter.
6	**DECIDE**	Apply human judgment and own the final outcome.

IMPORTANT

Asking AI to critique or verify itself can be useful, but it is not independent verification. When accuracy matters, use authoritative sources, trusted data, or qualified human review.

07 Confidence is not accuracy

Generative AI is optimized to produce useful-looking outputs. Your job is to know when "looks right" is not enough.

What should be verified?

- facts that influence a decision;
- claims about laws, regulations, policies, contracts, or requirements;
- claims about what a source document says;
- high-consequence recommendations or interpretations.
- names, dates, figures, quotations, citations, and links;
- calculations or quantitative conclusions;
- current information that may have changed;

A simple verification sequence

1	IDENTIFY	Which parts of the output are factual claims versus suggestions, synthesis, or creative content?
2	TRACE	Can the claim be tied to an authoritative source or the supplied material?
3	CHECK	Open the source. Confirm that it exists, is current enough, and actually supports the claim.
4	COMPARE	For consequential claims, compare against another authoritative source or trusted dataset when appropriate.
5	DOCUMENT	When traceability matters, preserve the source and the reasoning behind the final human decision.

SOURCE RULE

A citation only becomes useful evidence after someone checks that the source exists and supports the claim.

08 The human part still matters

AI can generate, summarize, detect patterns, compare options, and predict. People still provide the context, standards, priorities, consequences, and accountability that turn capability into responsible work.

AI CONTRIBUTION	THE HUMAN PART
AI can generate	People judge what is accurate, relevant, appropriate, and worth keeping.
AI can summarize	People decide what deserves attention.
AI can find patterns	People interpret what the pattern means in the real world.
AI can suggest options	People weigh tradeoffs and own the decision.
AI can accelerate	People set direction and define the objective.
AI can process information	People supply context, history, nuance, and exceptions.
AI can predict	People weigh consequences, uncertainty, and risk tolerance.
AI can draft	People define what good looks like and when the work is finished.

HUMAN-IN-THE-LOOP IS NOT A CHECKBOX

Human review only adds value when the reviewer has enough context, authority, time, and expertise to challenge the output. A person clicking “approve” without meaningful review is not effective oversight.

09 Use AI with boundaries

Fluency includes knowing that the best possible output is not always worth the information, permissions, or risk required to produce it.

Before you use AI for a work task, ask:

Is the tool approved? Use organization-approved systems and configurations for organization work.

Is the information appropriate to share? Consider privacy, security, confidentiality, proprietary information, client data, regulated data, and records requirements.

Does the task require disclosure? Some organizations, clients, assignments, or contexts may require transparency about AI assistance.

Who reviews the output? Define who is responsible for checking accuracy, quality, and appropriateness.

What happens if it is wrong? The higher the consequence, the stronger the review and control should be.

Is AI actually the right tool? Do not add AI to a task when a simpler, deterministic, or authoritative method is more appropriate.

THE BOUNDARY PRINCIPLE

What you can enter into an AI system and what you should enter are different questions.

When in doubt

Pause the AI workflow and follow the organization's policy, security guidance, legal/compliance requirements, or subject-matter review process. Fluency includes recognizing when the tool should not be the next step.

10 Choose the right role for AI

Instead of asking "Can AI do this?", ask "What role should AI play in this work?"

SITUATION	AI ROLE	EXAMPLES
Low consequence + easy to review	EXPLORE FREELY	Brainstorming, formatting, first drafts, alternative phrasing
Useful but needs context	ASSIST WITH HUMAN REVIEW	Summaries, research synthesis, comparisons, internal analysis
High consequence or hard to verify	USE CAUTIOUSLY / ESCALATE	Legal interpretations, medical decisions, personnel actions, safety-critical recommendations
Exact / deterministic requirement	PREFER AUTHORITATIVE OR DETERMINISTIC SYSTEMS	Official records, calculations requiring guaranteed logic, source-of-truth lookups
Sensitive information involved	USE ONLY APPROVED ENVIRONMENTS AND RULES	Client data, regulated information, proprietary materials, internal-only content

FIT BEFORE FEATURES

The newest or most capable model is not automatically the best choice. Start with the problem, the workflow, the information, the consequence of error, and the review process.

11 Build skills that outlast the tool

AI products will keep changing. Durable fluency comes from learning principles you can transfer to unfamiliar systems.

When you encounter a new AI tool, ask:

1	**WHAT IS IT?**	What model or capability is underneath the product? What can the tool actually do?
2	**WHAT PROBLEM DOES IT SOLVE?**	Is there a real workflow or user need, or only an impressive demo?
3	**WHAT INFORMATION DOES IT USE?**	What data can it access, retain, retrieve, or send elsewhere?
4	**WHAT HAPPENS WHEN IT IS WRONG?**	How visible are errors, and who reviews them?
5	**WHAT CAN IT ACT ON?**	Does it only generate output, or can it take actions through tools and systems?
6	**HOW WILL WE KNOW IT HELPS?**	Define measurable value before scaling adoption.

ADAPTABILITY PRINCIPLE

Do not memorize every AI product. Learn how to interrogate a new one.

Fluent users transfer the same habits across tools:

- define the problem before choosing the product;
- set constraints and success criteria;
- verify when evidence matters;
- keep human accountability clear.
- provide context intentionally;
- review the output instead of deferring to it;
- protect information and follow approved-use rules;

12 AI fluency quick reference

Keep these questions nearby when AI becomes part of the work.

STAGE	QUESTIONS
BEFORE	What is the objective? Is AI a good fit? Is this tool approved? What information is appropriate to provide?
DURING	Did I provide enough context? Did I define constraints? Am I reviewing the output critically?
AFTER	What needs verification? What may be missing? Does the output meet the real standard? Who owns the decision?

Ten habits of AI-fluent professionals

01 They start with the problem, not the tool.

02 They know that polished output can still be wrong.

03 They provide context intentionally rather than dumping information indiscriminately.

04 They distinguish generated content from retrieved evidence.

05 They iterate instead of treating the first answer as final.

06 They ask what assumptions or information may be missing.

07 They verify important claims using authoritative sources.

08 They protect sensitive information and follow approved-use rules.

09 They keep human judgment and accountability visible.

10 They can transfer these habits to a new AI tool without starting from zero.

ONE SENTENCE TO REMEMBER

AI fluency is not knowing every tool. It is knowing how to work intelligently with whatever tool comes next.

NEXT STEP

Move from understanding to practice

This field guide is the foundation. AI capability becomes stronger when people practice these ideas, learn how to collaborate with AI, and understand where their own gaps remain.

AI FOR IMPACT RESOURCE	PURPOSE	WHAT IT ADDS
30-Day AI Skills Challenge	PRACTICE	Short exercises that turn AI concepts into everyday habits.
Human + AI: The Collaboration Lab Workbook	COLLABORATE	Hands-on scenarios for learning how AI contribution and human judgment work together.
AI Capability Self-Assessment	ASSESS	A structured way to identify strengths and gaps across practical AI capabilities.

SUGGESTED PATH

UNDERSTAND → PRACTICE → COLLABORATE → ASSESS

A final note

The goal of AI fluency is not to make every person an AI expert. It is to help people become informed, capable users who understand enough to use AI deliberately, question it when necessary, and keep human responsibility at the center of the work.

SELECTED REFERENCES

Further reading

These resources provide additional background on trustworthy AI, generative AI, and retrieval-augmented systems.

National Institute of Standards and Technology (NIST). Artificial Intelligence Risk Management Framework (AI RMF 1.0). 2023. https://www.nist.gov/itl/ai-risk-management-framework

National Institute of Standards and Technology (NIST). Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (NIST AI 600-1). 2024. https://doi.org/10.6028/NIST.AI.600-1

NIST Computer Security Resource Center. Glossary entry: Generative Artificial Intelligence. https://csrc.nist.gov/glossary/term/generative_artificial_intelligence

IBM. What is Retrieval-Augmented Generation (RAG)? https://www.ibm.com/think/topics/retrieval-augmented-generation

EDITORIAL NOTE

AI capabilities, product features, policies, and terminology evolve quickly. Treat this guide as a practical foundation and confirm tool-specific behavior and organizational requirements at the time of use.