

Practical AI Fundamentals

Concepts, Tools, Prompts, Guardrails,
and Near-Term Opportunities

Presented to: IACC | June 3, 2026

John Fullmer

Managing Partner, Blue Field Ventures

john@gobluefield.com

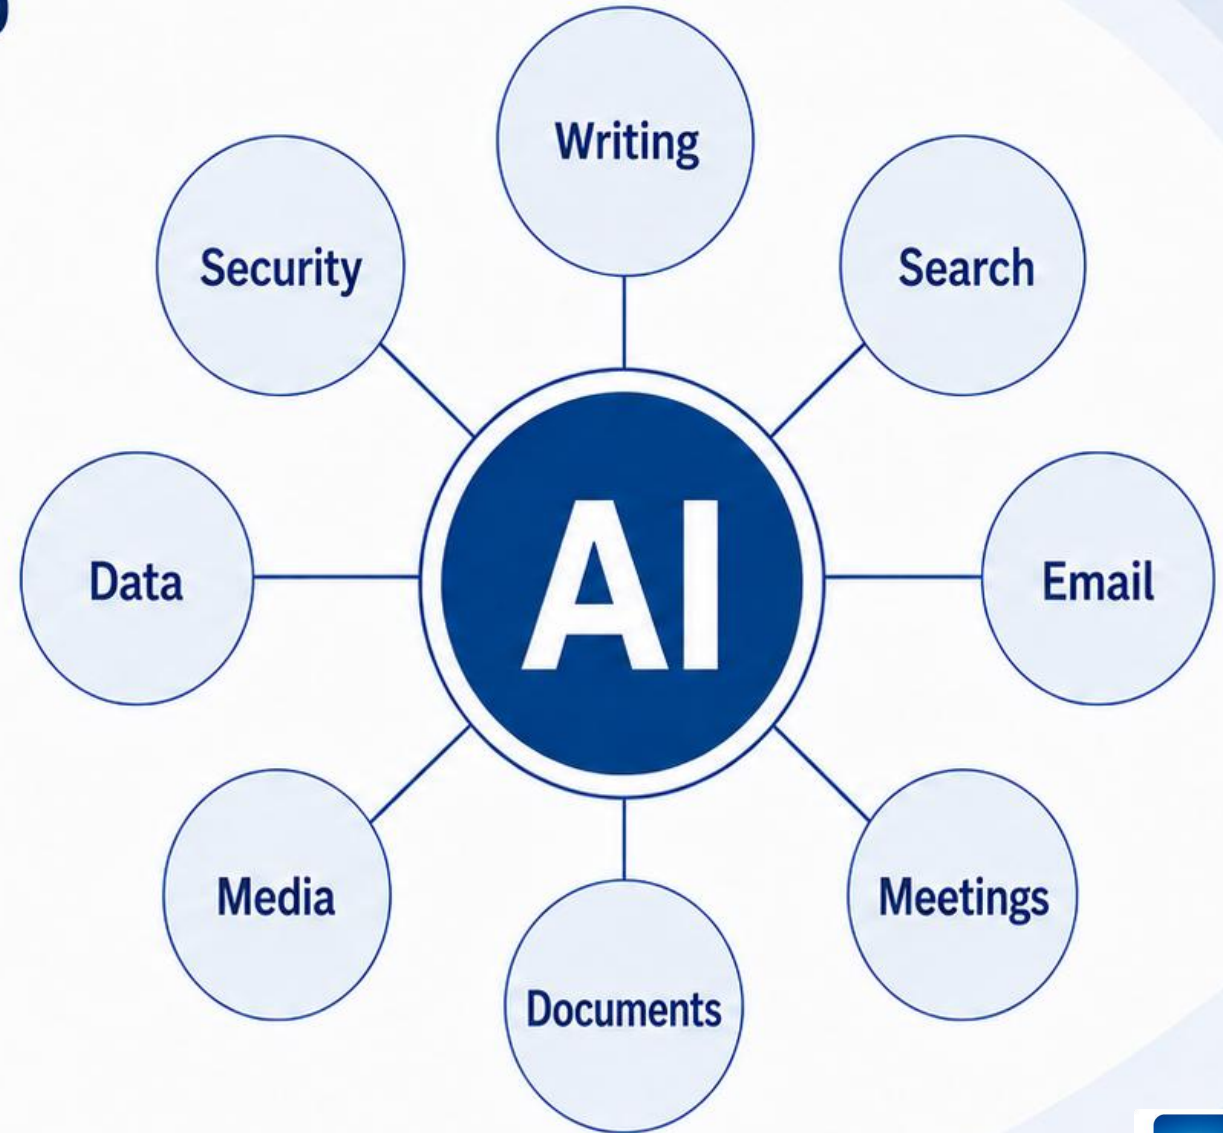


What is AI?

Where AI Is Showing Up

AI tools are increasingly built into familiar workplace systems, including:


- Writing and editing tools
- Search engines
- Email and office software
- Meeting summaries
- Document review
- Image, audio, and video tools
- Customer service systems
- Data analysis tools
- Cybersecurity threats



What is AI Really?

At a practical level, AI is software that can:

- Recognize patterns
- Make predictions
- Generate text, images, audio, video, or code
- Summarize information
- Classify or organize content
- Respond to instructions in ordinary language



Simple version:

AI uses patterns in data to produce useful outputs.



Major Work Shifts Through Time


Era	Old Tool	New Tool
Agrarian Age	 Shovel	 Plow
Industrial Age	 Horse power	 Steam power
Technology Age	 Abacus	 Computer
Information Age	 Books only	 Internet search
AI Age	 Manual drafting, searching, and summarizing	 AI-assisted work

New tools change how work gets done.



What Happens When You Ask AI a Question

How AI turns instructions into outputs.



Piece	What it means
Data	Material used to train or guide the AI
Model	The trained system that finds patterns and generates responses
Prompt	The instruction or question you give the AI
Context	What the AI can currently see, such as the conversation or provided material
Output	What the AI produces: text, summary, image, code, answer, or recommendation
Review	The human step: checking accuracy, judgment, source, and usefulness

The output is not the end of the process.
The output still needs review.



Why AI Feels Intelligent



Understand ordinary language



Follow instructions



Use the current conversation as context



Explain complicated topics simply



Summarize long material



Compare ideas



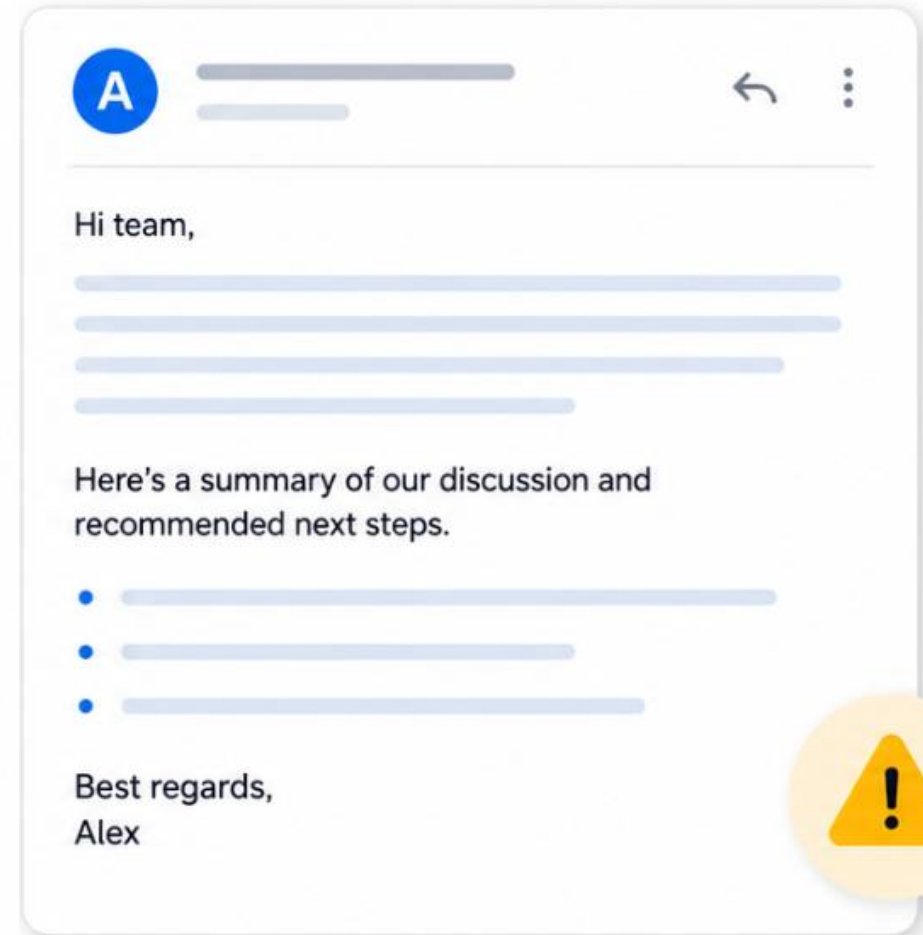
Change tone and format



Revise patiently



Produce polished work quickly



A polished email can still be wrong.

AI can sound clear and professional while still including mistakes.

Subject: Updated Policy Reminder

Hello Team,

This is a reminder that the updated Expense Policy is now in effect. Please review the key points below and ensure your compliance.

All expense reports for April must be submitted by Friday, May 10, 2024. Late submissions will not be accepted.

All expenses over \$75 require manager pre-approval before the expense is incurred.

Personal mileage is reimbursable at \$0.70 per mile effective May 1, 2024.

If you have any questions, please reach out to the Finance team.

Thank you,
Finance Operations



1) Wrong deadline

The email includes a deadline that was never in the original policy.



2) Made-up requirement

It adds a requirement that does not appear in the actual source document.



3) Missing exception

It leaves out an important exception that changes how the rule applies.



4) Assumed detail

It sounds specific, but the source never confirmed that detail.

The output is not the end of the process.
The output still needs review.



Prediction Machine, Not Truth Machine

Prediction Machine

- Pattern-based
- Fast
- Helpful
- Generates likely responses

Truth Machine

- Verified
- Authoritative
- Source-backed
- Requires human review








**AI can help produce the work.
It does not replace checking the work.**



The LLM is a Patient Assistant

It is like a patient assistant that can read, write, summarize, explain, and revise quickly.

It can help you:

-  Draft a first version
-  Rewrite something more clearly
-  Summarize long material
-  Explain a topic at different levels
-  Compare ideas
-  Turn notes into structure
-  Create checklists
-  Suggest next steps
-  Keep revising without getting tired

How the assistant helps



But it still needs clear instructions, source material, boundaries, and review.



AI is patient and useful, It is not accountable.







Where AI Shines and Still Needs Verification

AI is best understood as a tool that can help people work with information.



- Recognize patterns
- Generate useful drafts
- Summarize long material
- Explain ideas in different ways
- Compare options
- Organize messy information
- Revise quickly through conversation

INFORMATION IN



AI PROCESS
Analyzes, understands,
and organizes
information

USEFUL OUTPUTS



AI helps produce useful work, but people still guide and review it.

What AI Is Not

What AI Isn't

AI can be useful, fast, and impressive.
But it is not:

- A search engine by default
- A truth machine
- A database of guaranteed facts
- A private diary
- A human expert
- A final decision-maker
- A replacement for judgment



AI \neq Search Engine



AI \neq Truth Machine



AI \neq Human Judgment

Useful \neq **Verified**



AI Is Not a Search Engine

Search Engine



- Retrieves pages, links, or documents
- Helps you find information to inspect
- Usually points you to sources
- Best for locating information

AI



- Generates likely responses
- Helps draft, explain, summarize, and reformat
- May answer directly without showing full support
- Best for creating and refining outputs

Search helps you find information. AI helps you work with information.

AI is a Probability Engine

AI predicts the most likely next words based on patterns in its training data.



Predicts what is likely useful

It does not “know” facts like a human or a database.



Not an authority engine

It does not determine truth or accuracy.



Confident does not mean correct

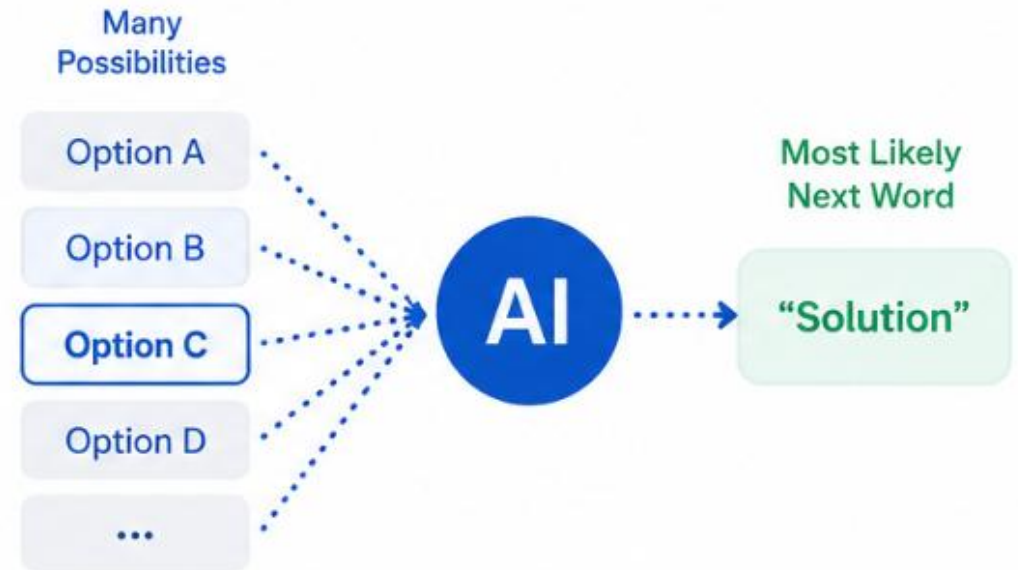
AI can sound certain and still be wrong.



Your role is still essential

You provide goals, context, judgment, and final decisions.

AI looks at many possibilities and chooses the next most likely word.



AI is great at patterns.
You are essential for judgment.



Conversational Does Not Mean Personal

Feels like...

- ✓ Conversational
- ✓ Helpful
- ✓ Patient
- ✓ Fast
- ✓ Polite
- ✓ Confident
- ✓ Always available



Actually is...

- ✗ A tool
- ✗ Not loyal
- ✗ Not private by default
- ✗ Not a trusted advisor by default
- ✗ Not a human expert
- ✗ Not responsible for the outcome

Treat it like a powerful, tireless assistant, but not a friend.



Best Practice

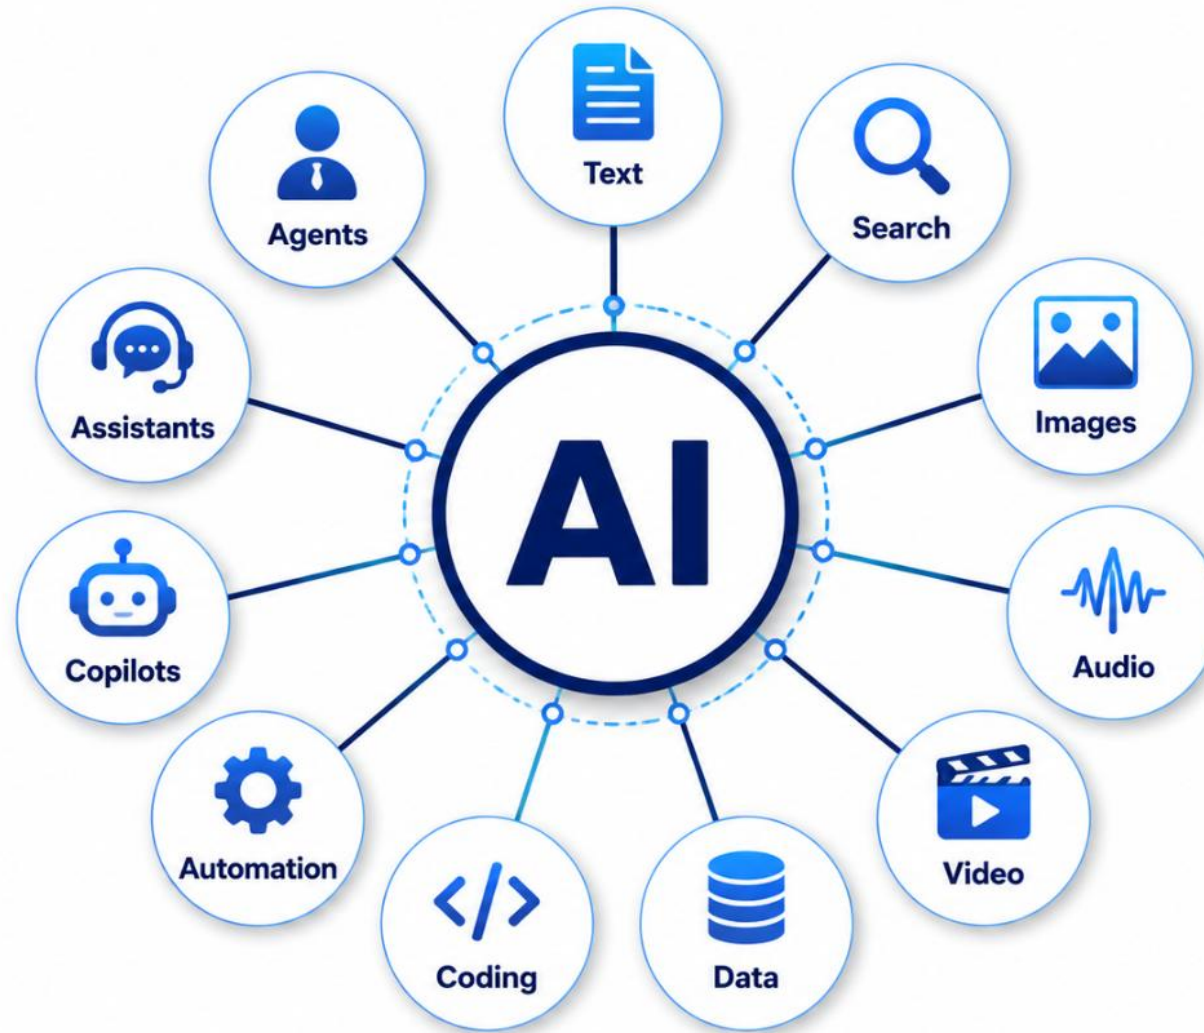


AI can accelerate the work.
People finish the work



What Types of AI Are There?

AI Is Not One Tool



Different AI tools are built for different jobs.



Text Generation Tools

These help create, revise, and organize written content.

Common Examples:

- ChatGPT
- Claude
- Gemini
- Microsoft Copilot

Useful for:

- Drafting
- Rewriting
- Summarizing
- Brainstorming
- Explaining
- Creating outlines
- Turning notes into structure
- Creating checklists
- Comparing ideas

Best for turning rough ideas into usable written drafts



AI Search Tools

These help create, revise, and organize written content.

Common Examples:

- Perplexity
- ChatGPT with search
- Gemini search features
- Microsoft Copilot search feature

Useful for:

- Researching a topic
- Finding sources
- Comparing options
- Summarizing search results
- Getting a quick overview
- Identifying questions to investigate further

Use AI search for discovery. Use sources for verification.



Image, Audio, Video, and Media AI

AI can now generate and edit more than text.

Common Examples:

- Image generation and editing
- Audio and voice generation
- Music generation
- Podcast-style summaries
- Video generation
- Avatar generation

Useful for:

- ChatGPT image tools
- Midjourney
- Gemini image tools
- Suno
- NotebookLM
- Sora
- Synthesia
- HeyGen

Media AI can turn a rough idea into something people can see, hear, and react to



Image, Audio, Video, and Media AI

Media AI helps move from description to demonstration.

Tool	Category
ChatGPT image tools	Image generation and editing
Midjourney	Image generation
Gemini image tools	Image generation and editing
Suno	Music generation
NotebookLM	Podcast-style summaries
Sora	Video generation
Synthesia	Avatar video generation
HeyGen	Avatar and video generation



Copilots and AI Inside Existing Software

AI is increasingly being built into tools people already use.

Best use:

Start by understanding what AI features are already available in the tools you use today.

Ecosystem / Tool	Where AI may show up
1. Microsoft 365 / Copilot	Word, Excel, Outlook, Teams, PowerPoint
2. Google Workspace / Gemini	Gmail, Docs, Sheets, Slides, Meet
3. Adobe tools	Creative editing, document tools, image tools
4. Zoom / Teams	Meeting summaries, transcripts, action items
5. Browsers and search tools	Writing help, search summaries, page summaries
6. Business software	CRM, help desk, document management, reporting

Before buying something new, understand what you already have.



Custom Assistants and GPTs

A custom assistant or GPT is a reusable AI setup built for a repeated task.

It can include:

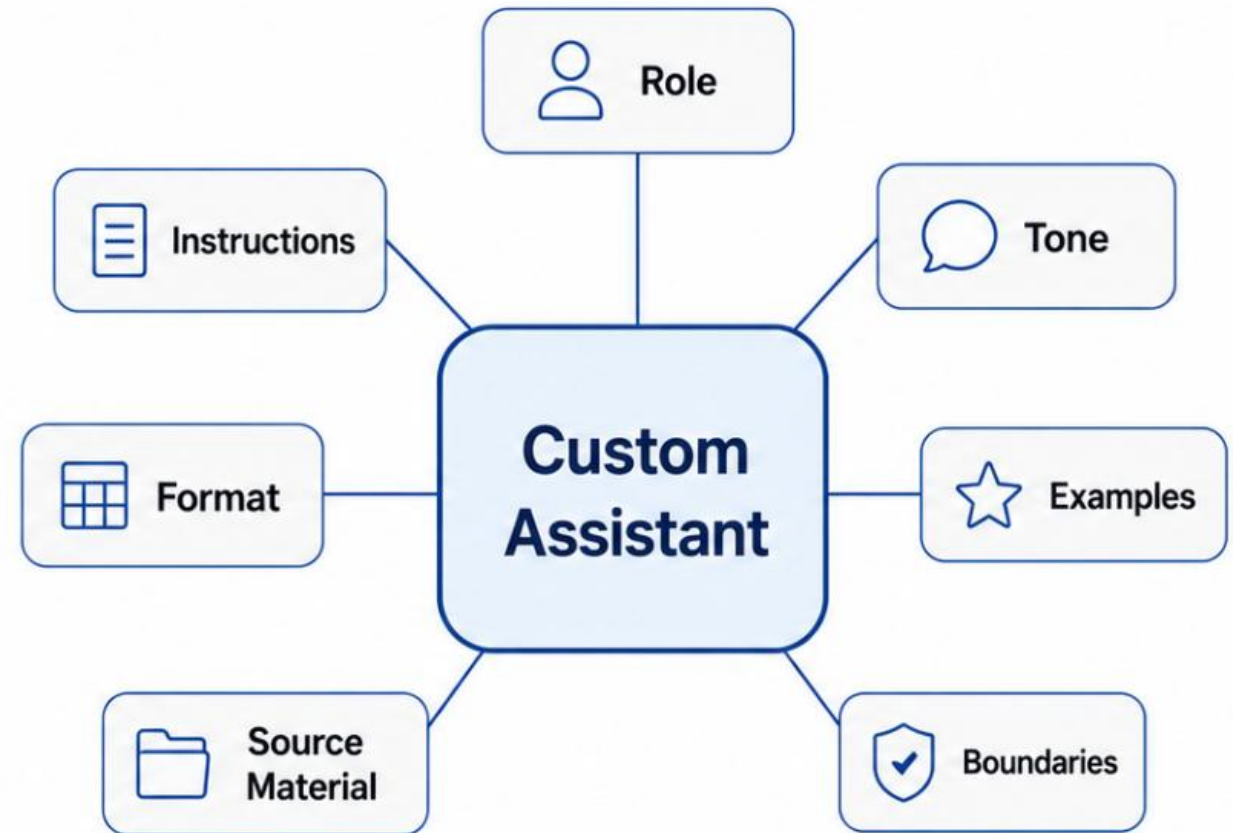
- A defined role
- Repeated instructions
- Preferred tone
- Standard formats
- Example outputs
- Source files or reference material
- Boundaries for what it should and should not do

Good uses include:

- Rewriting content in a consistent style
- Creating training materials
- Drafting FAQs
- Summarizing recurring documents
- Helping with repeated internal workflows



Best use: Use a custom assistant when you repeat the same type of task often.



A custom assistant turns repeated instructions into a reusable workflow.



Tool Choice Changes Quickly

AI tools are improving and changing fast.

A tool may change:

- Features
- Pricing
- Privacy settings
- Data controls
- Integrations
- Quality
- Availability
- Terms of use

The key skill is not memorizing every tool.
The key skill is knowing how to evaluate
the right tool for the job.

A simple evaluation lens



1. What job do we need done?

Be clear on the task or outcome.



2. What data is involved?

Consider sensitivity, volume, and ownership.



3. What tool category fits?

Match the job to the right type of tool.



4. What review is needed?

Check security, privacy, compliance,
and business impact.



5. Is this tool approved or appropriate?

Follow company guidance and policies.

Do not chase every tool. Learn how to choose the right tool.



BLUE FIELD

www.gobluefield.com

Getting Usable Results

Five Rules for Better AI Results

1. **Start with a clear task**

Know what you want the AI to help you produce.

2. **Give enough context**

Include the audience, purpose, source material, and constraints.

3. **Do not use sensitive information casually**

Practice with public, generic, or low-risk examples.

4. **Review the output**

Check facts, assumptions, tone, and missing details.

5. **Keep improving the result**

The first draft is the starting point, not the final answer.

Stay in charge of the process



The Prompt Formula

Prompt Element	What to Include	Example
Role	Who should the AI act as?	“Act as a writing assistant.”
Task	What do you want it to do?	“Rewrite this policy summary.”
Context	What background or source material should it use?	“Use only the text below.”
Audience	Who is the output for?	“Write this for employees.”
Tone	How should it sound?	“Clear, professional, and direct.”
Format	What should the output look like?	“Format it as a short email.”
Review instruction	What should it flag or avoid?	“Flag anything that needs verification.”

Clear instructions create better first drafts.



Small Changes, Big Differences

Example Information: There is a new check-in process at North Valley Office, it starts July 1, visitors need to check in at the front desk, visitors get a temporary badge, employees should remind visitors to return badges, and managers will get more instructions next week.

Loose Prompt	Controlled Prompt
“Write an announcement about our new policy.”	“Using only the policy text below, draft a plain-language announcement for employees. Do not add dates, requirements, or benefits not included in the source text. Put anything uncertain under ‘Needs verification.’”
Loose Output Might Say	Controlled Output Is More Likely to Say
“All visitors must present ID and sign the security log beginning July 1.”	“Starting July 1, visitors must check in at the front desk and receive a temporary badge.”
“The new policy applies to all company locations.”	“This applies to North Valley Office.”
“Managers must train employees before July 1.”	“Department managers will receive additional instructions next week.”
No uncertainty section	Includes “Needs verification” if details are missing









A better prompt gives AI less room to guess and more direction to follow.

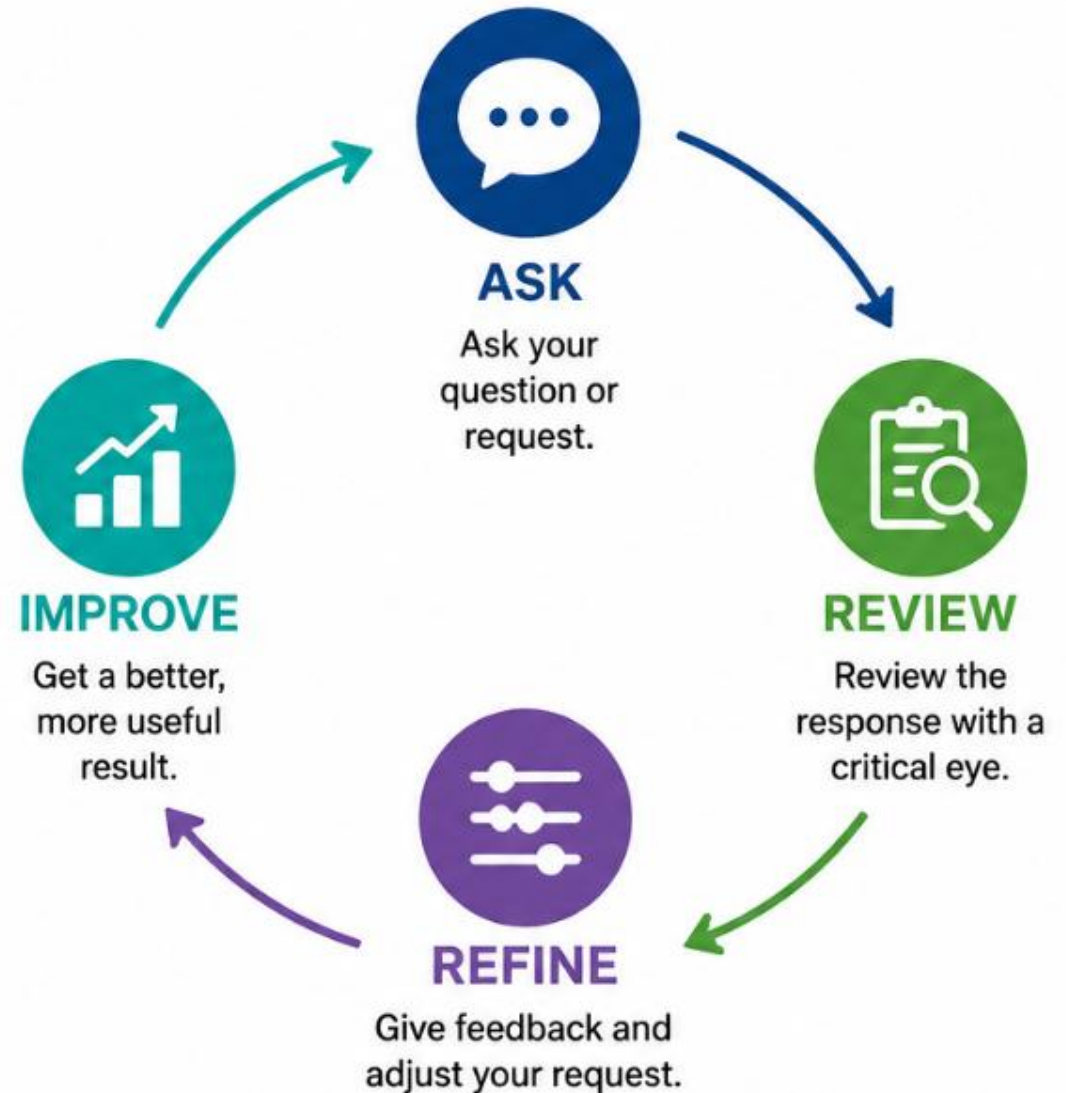


Conversation Mode Is the Advantage

AI works best when you treat the first answer as the beginning of the process.

You can follow up with:









-  Make this shorter.
-  Make it clearer.
-  Change the tone.
-  Give me three versions.
-  Turn this into a checklist.
-  Explain it for a beginner.
-  What is missing?
-  What needs verification?



Use AI to Improve Your AI Use

AI can help you get better at working with AI.

Ask AI to help you:

-  **Improve** a weak prompt
-  Make a prompt **safer** or **clearer**
-  Create a **reusable** prompt
-  **Compare tools** for a task
-  Identify **risks** in an AI workflow
-  Design a **custom assistant** or GPT
-  Create a **review checklist**
-  Build a simple **test plan**



Example prompts

- “ Make this prompt clearer and more specific.”
- “ What information should I provide to get a better answer?”
- “ Create a reusable prompt for this task.”
- “ What should I verify before using this output?”
- “ What risks should I watch for in this workflow?”

Use AI not just for answers, but to improve the way you ask.



Hallucinations, Scams, and Guardrails

Hallucinations – why do they happen?

This happens in part because AI is a **prediction engine**.
Missing information + prediction = possible hallucination

AI predicts a likely response based on:

- Your prompt
- The current conversation
- Patterns from training data
- Any source material provided

Common hallucinations include:

- Fake citations
- Wrong dates
- Made-up policies
- Invented quotes
- Unsupported requirements
- Incorrect summaries
- Assumed details presented like facts

The danger is not that AI sounds ridiculous.
The danger is that it sounds believable.



Sensitive Information



Do not paste sensitive information into AI tools unless the tool is approved for that use.

TAKEAWAY:



If you would not want it retained, disclosed, or reused, **do not** paste it into an unapproved AI tool.

Sensitive information may include:



People Data

- Customer or client information
- Employee records
- Protected health information



Financial Data

- Financial identifiers
- Bank account details
- Payment information
- Tax information



Legal / Confidential Data

- Legal strategy
- Non-public contracts
- Confidential business information



Security / Proprietary Data

- Passwords or credentials
- Security details
- Proprietary data
- System access information







When in doubt, leave it out.



Fake Media, Scams, and AI-Built Tools

AI can create things that look **real** or **functional** before they have been verified.

Fake or Misleading Content


-  Fake emails
-  Fake voices (voice cloning)
-  Fake images
-  Fake videos (deepfakes)
-  Fake invoices
-  Fake internal requests



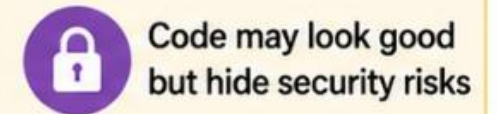
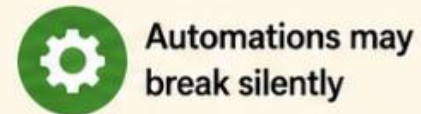
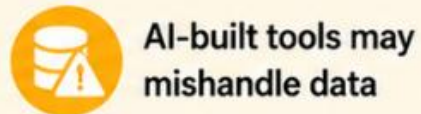
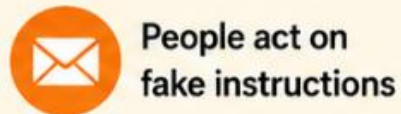
Looks real \neq Verified

Looks functional \neq Safe

AI-Built Tools and Outputs

-  AI-generated code
-  Scripts and programs
-  Automations and agents
-  Dashboards and reports
-  Websites or apps

Common Risks



If AI creates something people will trust or use, it needs review.



Guardrails and Human Review

The rules and controls that make AI use safer, more consistent, and more accountable.



Guardrails can include:



Approved tools Use only tools that are reviewed and approved.



Data rules Define what data can be used and how it is protected.



Source limits Require AI to use provided or approved sources only.



Human review People review important or sensitive output.



Access controls Limit who can use tools and what they can access.



Logging Record activity for transparency and accountability.



Escalation paths Define when and how issues get escalated.



Testing Test prompts, outputs, and automations.



Vendor review Review vendor practices for security and privacy.



Human Review

Before relying on or sharing AI output, ask these questions:

1



Is it accurate and complete?

2



What source supports it?

3



Did it add anything unsupported or make assumptions?

4



Is sensitive or confidential information involved?

5



Is the tone appropriate for the audience and purpose?

6



Is this safe to use, publish, or act on?



FIELD

So, what are agents?

Prompting vs. Agents

What Are Agents, and How Are They Different From Prompting?

Prompting	Agents
One-off task	Recurring workflow
You provide the instruction each time	The system follows a defined process
AI gives an answer or draft	AI works through multiple steps
Depends on your judgment in the moment	Built with review points
Best for drafting, summarizing, and explaining	Best for monitoring, routing, preparing, and checking

Prompt	Agent
“Summarize this document.”	“Every morning, review new documents, summarize key items, flag missing information, and prepare a review packet.”

Prompting helps complete a task. Agents help manage a process.



Agents = ARR

ARR = Autonomous, Recurring, Reviewable



A well-designed agent may include:



Agents are workflows, not just answers.



Example of a Prompt vs. Agent

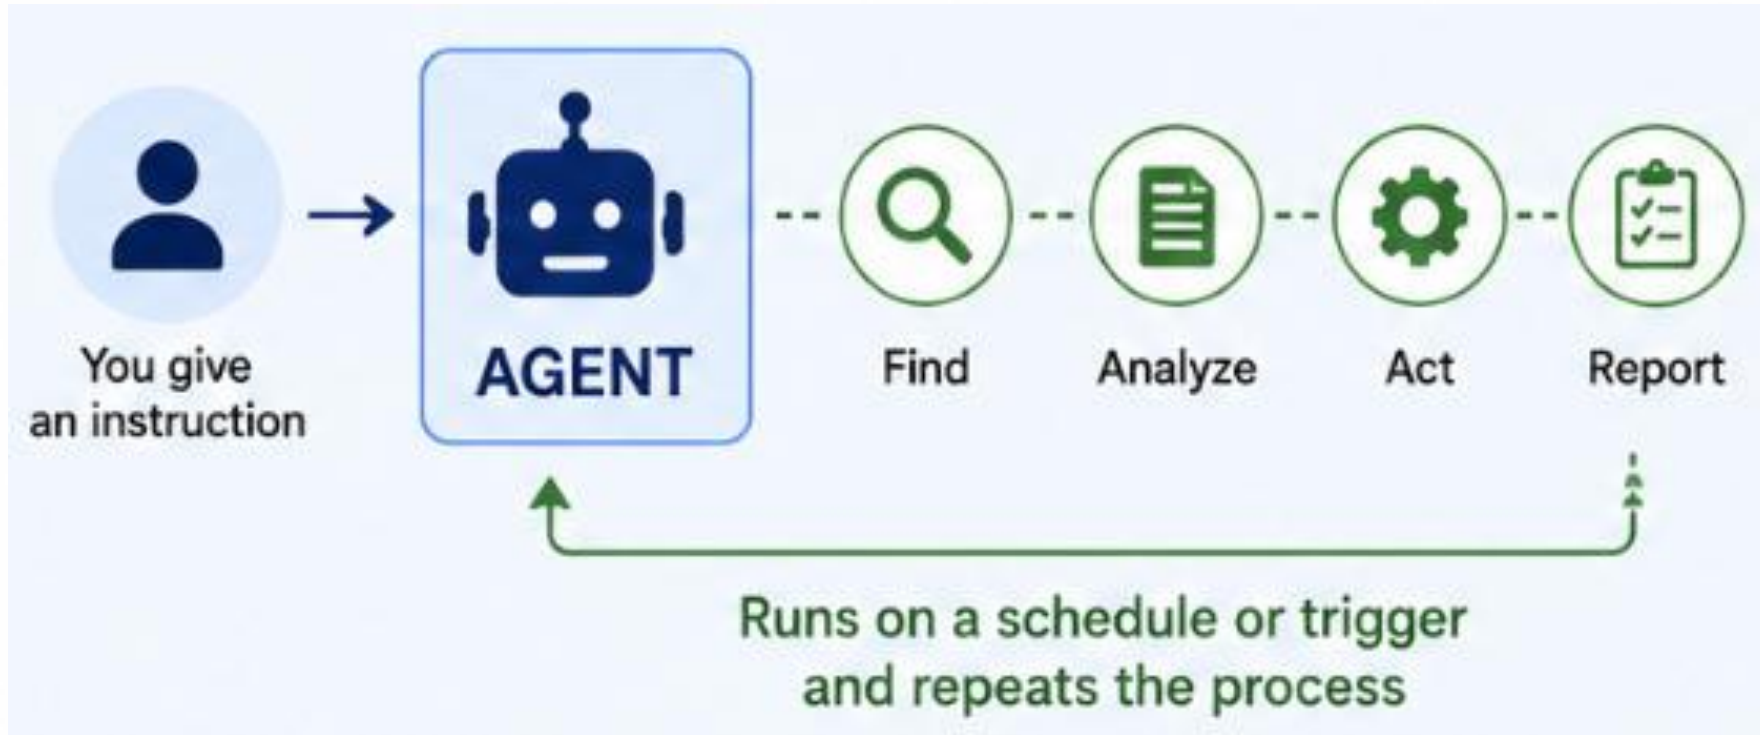
Prompt	Agent
“Summarize this invoice and tell me if anything looks unusual.”	“Every Monday morning, review new invoices, compare them against approved vendors and expected amounts, flag anything unusual, prepare a summary report, and send it to a human for approval.”

Agents are process automation tools.



Why Agents Need Stronger Guardrails

Agents can do more than respond to a single prompt, and most often perform independently.



More autonomy requires more oversight.



What Can I Start Doing Today?

Start With Low-Risk, Useful Tasks

Good first uses of AI include:

- Rewrite dense language in plain English
- Draft internal checklists
- Summarize public or non-sensitive notes
- Create training quizzes
- Turn procedures into FAQs
- Improve email drafts
- Create multiple versions of a message
- Ask for clarity edits
- Ask what needs verification

Start with content that is:


- Public
- Generic
- Non-sensitive
- Easy to review
- Low consequence if the first draft is wrong

Start where the risk is low
and the value is easy to see.












Use Tools You Already Have and Inventory Current Use

Before buying something new, understand what is already available.



AI May Already Be Built Into

-  Microsoft 365 / Copilot
-  Google Workspace / Gemini
-  Adobe tools
-  Zoom or Teams summaries
-  Search tools
-  Email drafting features
-  Document management systems
-  Business software
-  Browser-based assistants



Build a Prompt Library and Pick a Few Pilots

Turn useful AI experiments into repeatable habits.

Start a simple prompt library for recurring tasks:

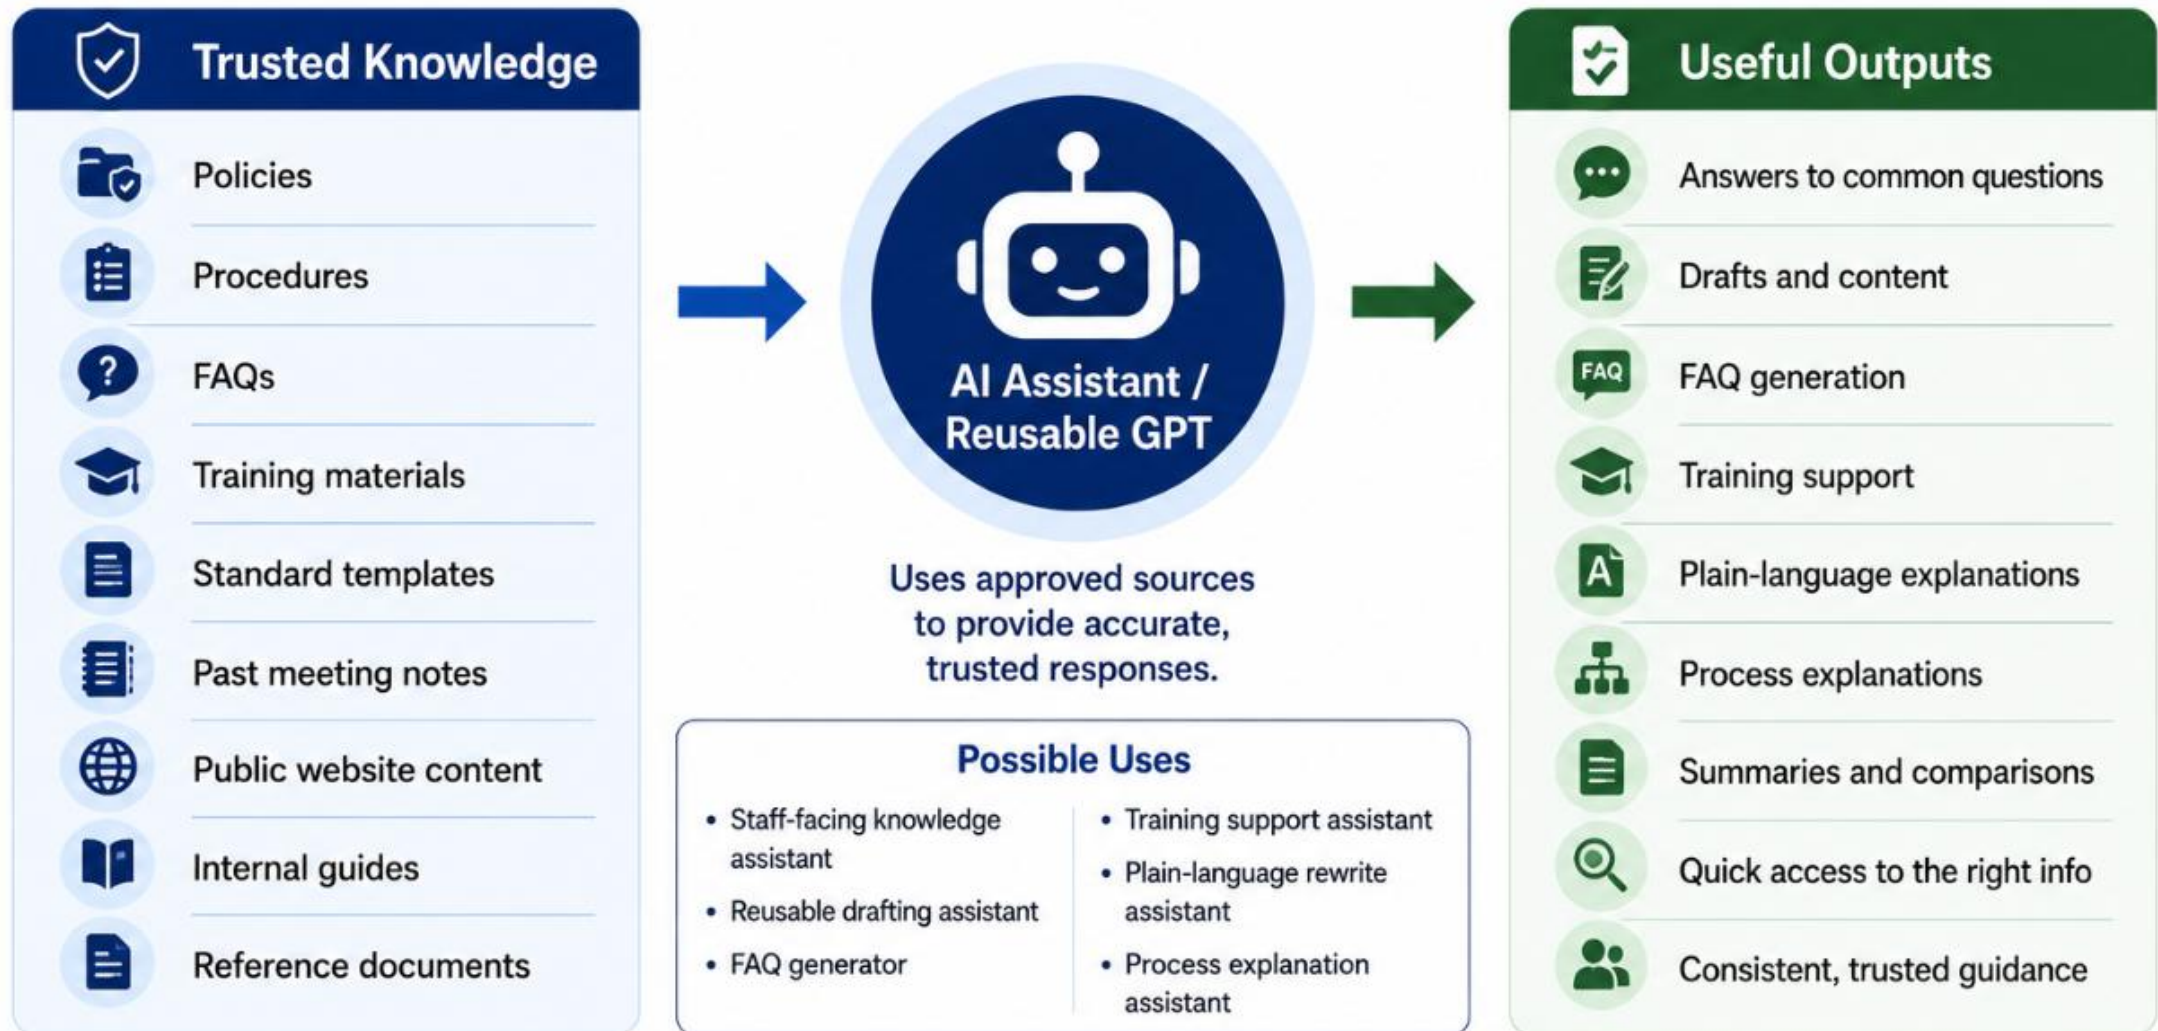
- Plain-language rewrite prompt
- Meeting summary prompt
- Checklist creation prompt
- Training quiz prompt
- Needs-verification prompt
- Email improvement prompt
- Policy or procedure summary prompt
- Tool comparison prompt

Start with a few, safe, repeatable wins.



Internal Knowledge Assistants and Reusable GPTs

One of the best AI opportunities is making existing knowledge easier to find and use



Reviewable Automation and Better Visibility


AI can help prepare, monitor, summarize, and surface information for people to review.

Useful opportunities include:

- Monitoring incoming requests
- Summarizing new items
- Flagging missing information
- Routing work
- Preparing review packets
- Tracking repeated issues
- Summarizing trends
- Supporting dashboards and reports
- Highlighting exceptions or anomalies

Important distinction:

- AI should prepare and surface information.
- People should review and decide.




- AI is used for recommendations
- Humans then make the decisions.

The goal is better visibility and faster preparation.








What to Prepare Now

Organizations that benefit from AI usually prepare three things



INFORMATION

Prepare trusted, clean, well-organized information.


-  Clean up key documents
-  Organize trusted sources
-  Remove outdated material
-  Define sensitive information
-  Clarify authoritative sources



PEOPLE






Prepare people with the right skills, awareness, and habits.

-  Train staff on prompting and effective use
-  Build review and verification habits
-  Increase privacy and security awareness
-  Clarify appropriate use and boundaries
-  Encourage feedback and improvements











PROCESSES

Prepare practical processes that support responsible AI.

-  Identify repetitive workflows
-  Define review steps and approvals
-  Document current processes
-  Start small pilots with clear outcomes
-  Measure results and learn quickly







Practical Preparation Steps

-  Identify repetitive workflows
-  Organize trusted documents
-  Define sensitive information
-  Clarify approved tools
-  Create review steps
-  Build a prompt library
-  Start with small pilots
-  Evaluate vendors carefully

Prepare before you scale.



Working With IT and Engineering

What to Discuss	Why It Matters
 Workflow	What work is repeated, slow, or difficult?
 Data	What information would AI need, and is it sensitive?
 Systems	What tools, documents, or software are already involved?
 Review	Who checks the output before it is used?
 Risk	What happens if the AI output is wrong?
 Success	What improvement would make the project worthwhile?



Questions to ask back

- ✓ Will our data train your model?
- ✓ Where is our data stored and for how long?
- ✓ Who can access it?
- ✓ What requires human approval?
- ✓ What logs, documentation, and training are included?



Wrapping up...

Final Takeaways

Remember:

1



AI is a powerful tool

It can draft, summarize, explain, organize, compare, and improve work.

2



It is not a search engine, truth machine, or human expert

It can be wrong or incomplete.

3



Better prompts produce better starting points

Clear instructions lead to better results.

4



The first draft is not the final answer

Always review, edit, and verify.

5



Sensitive information should stay out of unapproved tools

Protect privacy, security, and trust.

6



Hallucinations can sound believable

Check facts and sources.

7



Agents need stronger guardrails

Define limits, approvals, and monitoring.

8



Start small, use low-risk tasks, and build from what works

Learn, measure, and improve.

Use AI as an accelerator, rely on people for judgement.



BLUE FIELD

www.gobluefield.com

Thank you!



AI/ML SOLUTIONS FOR COUNTY GOVERNMENT FRAUD

How AI/ML Can Detect, Prevent & Investigate Fraud Across Key Fraud Types



FRAUD CATEGORY	COMMON EXAMPLES	DATA SIGNALS & FEATURES	AI/ML APPROACHES	HOW AI/ML HELPS	BUSINESS IMPACT
1. INTERNAL FRAUD (EMPLOYEES / OFFICIALS)	<ul style="list-style-type: none"> Ghost employees Overtime abuse Cash skimming Unauthorized overrides 	<ul style="list-style-type: none"> Overtime hours vs peers Duplicate bank accounts Unusual manual overrides Lifestyle vs income mismatch 	<ul style="list-style-type: none"> Anomaly Detection (Isolation Forest, Autoencoders) Peer Group Analysis Graph Analytics (shared accounts) 	<ul style="list-style-type: none"> ✓ Flags unusual employee behavior ✓ Detects hidden relationships ✓ Prioritizes high-risk employees 	<ul style="list-style-type: none"> Reduce internal theft Lower financial losses Improve accountability
2. PROCUREMENT & VENDOR FRAUD	<ul style="list-style-type: none"> Bid rigging / collusion Overbilling Fake vendors Change-order inflation 	<ul style="list-style-type: none"> Vendor win rate patterns Invoice anomalies Shared addresses / phones Purchases split below threshold 	<ul style="list-style-type: none"> Graph / Network Analysis Clustering (vendor behavior) Supervised Classification (if labels available) 	<ul style="list-style-type: none"> ✓ Detects collusion & conflicts ✓ Finds unusual vendor behavior ✓ Scores vendor risk in real time 	<ul style="list-style-type: none"> Prevent overpayments Ensure fair competition Strengthen procurement integrity
3. PUBLIC ASSISTANCE & BENEFITS FRAUD	<ul style="list-style-type: none"> False eligibility Identity theft Duplicate enrollments Provider billing fraud 	<ul style="list-style-type: none"> Income vs reported data Identity match / reuse Household composition Service utilization patterns 	<ul style="list-style-type: none"> Classification (XGBoost, RF) Anomaly Detection Entity Resolution / Record Linkage 	<ul style="list-style-type: none"> ✓ Identifies ineligible applicants ✓ Finds duplicate identities ✓ Stops improper payments 	<ul style="list-style-type: none"> Protect public funds Ensure benefits go to eligible residents Improve program trust
4. REVENUE & TAX FRAUD	<ul style="list-style-type: none"> Property tax evasion False exemptions Skimming payments Altered records 	<ul style="list-style-type: none"> Property value vs comparables Payment timing anomalies Fee waivers & adjustments Missing or altered records 	<ul style="list-style-type: none"> Regression (value prediction) Outlier Detection Time Series Anomaly Detection 	<ul style="list-style-type: none"> ✓ Detects undervalued properties ✓ Finds unreported transactions ✓ Identifies abnormal fee waivers 	<ul style="list-style-type: none"> Increase revenue capture Improve compliance Reduce revenue leakage
5. GRANT & FUNDING FRAUD	<ul style="list-style-type: none"> Misuse of funds Falsified reports Inflated metrics Cost misallocation 	<ul style="list-style-type: none"> Spend categories vs grant scope Unusual spending patterns Report text similarity Performance metric trends 	<ul style="list-style-type: none"> NLP / Text Analytics (similarity, topic modeling) Classification Clustering (spend patterns) 	<ul style="list-style-type: none"> ✓ Detects off-scope spending ✓ Finds copy/paste reporting ✓ Improves grant oversight 	<ul style="list-style-type: none"> Ensure grant compliance Avoid penalties Improve outcomes
6. CYBER & PAYMENT FRAUD	<ul style="list-style-type: none"> Phishing / BEC attacks Payment redirection Credential theft Fraudulent transfers 	<ul style="list-style-type: none"> Email content & urgency Login behavior / device info Vendor bank changes Transaction velocity & size 	<ul style="list-style-type: none"> NLP (phishing, BEC detection) Behavioral Anomaly Detection Real-time Fraud Scoring 	<ul style="list-style-type: none"> ✓ Detects suspicious emails ✓ Flags unusual logins & actions ✓ Stops fraudulent payments fast 	<ul style="list-style-type: none"> Prevent financial loss Protect sensitive data Strengthen cybersecurity
7. PERMITS, LICENSING & INSPECTION FRAUD	<ul style="list-style-type: none"> Bribes to pass inspections Unauthorized approvals Falsified reports Improper licensing 	<ul style="list-style-type: none"> Approval rates vs peers Time-to-approval patterns Inspection outcome anomalies Repeat applicants / addresses 	<ul style="list-style-type: none"> Anomaly Detection Peer Comparison Models Process Mining 	<ul style="list-style-type: none"> ✓ Finds suspicious approval patterns ✓ Detects falsified inspections ✓ Improves process compliance 	<ul style="list-style-type: none"> Ensure public safety Improve compliance Increase public trust

CROSS-CUTTING AI/ML ENABLERS



Data Integration
Combine data across departments (ERP, HR, Tax, Grants, Permits, etc.)



Entity Resolution & Graphs
Unify identities & relationships across systems



Model Monitoring & Feedback Loops
Continuously learn from investigations and outcomes



Explainability & Compliance
Transparent models, audit trails, fairness & regulatory alignment



Human-in-the-Loop
AI prioritizes cases; humans make final decisions



OUTCOME
Smarter detection.
Faster investigations.
Stronger communities.



BLUE FIELD

www.gobluefield.com