

# Ten Prompting Practices That Improve AI Results

A Professional Training Guide for HR Leaders and Public-Sector Administrators



**BoldPath Consulting, LLC**  
**Patrick W Glynn, Founder/Principal Consultant**  
**(920) 522-2413**  
**patrick@boldpathconsulting.com**

**Use, Distribution, and AI Assistance Notice**

This material is the original work of BoldPath Consulting and is provided for individual use by participants in this session. It is not intended for redistribution, reproduction, or sharing outside your organization without permission.

Generative AI tools were used to assist in developing this content and, where applicable, supporting visuals. All material has been reviewed, structured, and validated for professional use.

BoldPath Consulting provides training and advisory services on generative AI, including policy development, implementation guidance, and practical applications in public-sector environments.

## Introduction

Generative AI systems respond directly to the clarity and structure of the instructions they receive. While many demonstrations focus on clever prompts or novel tricks, the most reliable improvements come from applying the same principles professionals use when supervising staff work: defining assignments clearly, specifying expectations, and verifying the results.

The following ten practices illustrate how small, deliberate changes in prompting can significantly improve the usefulness, accuracy, and reliability of AI-generated outputs in HR, labor relations, and public-sector administration. Each section pairs a conceptual explanation with contrasting examples drawn from realistic administrative contexts.

No technical background is required. These practices build on professional habits that experienced administrators already apply when assigning work to staff, reviewing drafts, and preparing documents for multiple audiences.

## 01. Prioritize Clarity Over Cleverness

### Concept Overview

AI systems respond to the clarity and specificity of the instructions they receive. A prompt that seems reasonable in conversation often fails as a professional assignment because it leaves too much open to interpretation. When a user asks an AI system to "help with a policy" or "review this situation," the system must infer the type of assistance needed, the audience, the format, and the scope. That inference process frequently yields generic commentary rather than a usable work product.

Effective prompting functions more like a well-structured staff assignment than a casual request. It names the deliverable, identifies the audience, specifies focus areas, and clarifies constraints. Asking the system to "draft a one-page summary of the remote work policy for department heads, focusing on eligibility criteria, approval authority, and documentation requirements" tells the system exactly what to produce and for whom.

This principle applies regardless of subject matter complexity. Professionals who have spent careers writing policy, managing investigations, and preparing administrative briefings already possess the instincts required: vague assignments produce vague results, and the quality of a work product is largely determined before the work begins. The same discipline applied to directing staff applies directly to directing AI systems.

#### Weak Prompt:

"Hey, can you take a look at this remote work policy and help me summarize it? I need something for the department heads."

#### Improved Prompt:

"Using only the attached remote work policy, draft a one-page summary for department heads. Cover three topics in this order: (1) employee eligibility requirements, (2) supervisor approval authority and process, and (3) documentation and recordkeeping obligations. Use plain language; avoid legal terminology. Format as three labeled sections with no more than three short paragraphs each."

#### Why the Improved Prompt Works:

The improved prompt eliminates interpretive guesswork. It specifies the deliverable (one-page summary), the audience (department heads), the scope (three defined topics), the sequence, the tone, and the format. Each element removes a decision the system would otherwise make independently, reducing the likelihood of output that is technically responsive but professionally unusable.

## 02. Specify the Desired Output Format

### Concept Overview

The structure of a response shapes its professional usefulness more than most users initially recognize. When a prompt does not specify format, AI systems default to narrative prose. Narrative is readable, but it often requires significant reorganization before it can be used in a meeting, a briefing packet, or an investigation file.

Specifying a format directs the system to organize information in a way that fits the user's workflow. A grievance summary organized under headings such as Allegation, Relevant Contract Section, Key Facts, and Outstanding Questions can be reviewed and discussed immediately. A comparison of job descriptions presented as a side-by-side table communicates more efficiently than prose describing the same differences. A compliance checklist produces a directly actionable document rather than a narrative that must later be converted.

Format specification is particularly valuable in public-sector environments where documents routinely serve multiple audiences. A single investigative summary may be reviewed by labor relations staff, department heads, legal counsel, and elected officials. Structuring the output to serve those audiences from the outset reduces downstream editing and helps ensure that critical elements are addressed consistently across similar documents. When in doubt, specify the format explicitly; it adds seconds to the prompt and removes significant revision time from the output.

#### Weak Prompt:

"Can you write up a summary of this grievance for me? I have a meeting tomorrow and need something to share with the team."

#### Improved Prompt:

"Using only the attached grievance documentation, produce a structured summary using the following four headings: (1) Allegation — one sentence stating the claimed contract violation; (2) Relevant Contract Language — quote the specific provision at issue; (3) Key Facts Documented — bullet list of confirmed facts only, no inference; (4) Outstanding Questions — bullet list of information needed before reaching a conclusion. Keep total length to one page."

#### Why the Improved Prompt Works:

Named headings force the system to address each required element, prevent narrative drift, and produce output that decision-makers can read and act on without reformatting. The explicit distinction between confirmed facts and outstanding questions also prevents the system from blending evidence and inference, which is a common failure in unstructured summaries.

## 03. Define the Audience and the Purpose

### Concept Overview

Every professional document is calibrated to its audience. A briefing for elected officials requires different depth, terminology, and framing than a procedural memo for supervisors or a policy explanation for front-line employees. When a prompt does not specify the intended audience, the AI system must infer the appropriate level of technical detail, vocabulary, and emphasis.

Those assumptions are often wrong in ways that require substantial revision. A response written at a level appropriate for HR staff may be impenetrable for department heads. A response aimed at a general audience may omit the precision that labor relations professionals require. Neither failure reflects the system's inherent limitations so much as insufficient direction.

Specifying both the audience and the purpose before generating a document is a small addition to the prompt that produces a proportionally larger improvement in the output. When the system understands that it is preparing a summary for a county board presentation rather than an internal working document, it adjusts its tone, scope, and emphasis accordingly. This step prevents the common situation in which the generated content is substantively accurate but professionally unsuitable without significant rewriting. It also prevents the opposite problem: responses that are appropriately simplified for a general audience when the actual reader is a credentialed professional who needs full technical detail.

<p><b>Weak Prompt:</b> "Write something about the new pay plan we're implementing. I need to communicate it to people."</p>
<p><b>Improved Prompt:</b> "Draft three separate communications about the county's new step-based pay plan: (1) a two-paragraph email for all employees explaining what changes to their paycheck and when; (2) a one-page FAQ for supervisors addressing how step placement decisions are made and documented; (3) a concise briefing paragraph for the county board summarizing the fiscal impact and implementation timeline. Each piece should be written for its specific audience with no jargon in the employee version."</p>
<p><b>Why the Improved Prompt Works:</b> Naming each audience and its specific informational need produces three targeted documents rather than one that serves no audience well. The instruction to eliminate jargon for employees and frame fiscal impact for the board gives the system audience-specific criteria, replacing a single vague purpose with three precise communication goals.</p>

## 04. Establish Guardrails and Constraints

### Concept Overview

AI systems are designed to be helpful, and in the absence of explicit boundaries they tend to extend responses into areas the user did not request. In administrative and HR contexts, this tendency creates real risk. A request to analyze a disciplinary situation may produce a response that includes legal opinions the system is not qualified to offer, disciplinary recommendations that properly belong to supervisors, or conclusions that go beyond what the documented evidence supports.

Guardrails are instructions that define the limits of the response as precisely as its objectives. They prevent the system from expanding into inappropriate territory and keep generated content in its proper role as a support tool rather than a substitute for professional judgment. A prompt analyzing a grievance might instruct the system to identify relevant contract provisions and compare them to the documented facts, while explicitly stating that the system should not characterize the merits of the grievance or recommend an outcome.

In sensitive personnel, labor relations, and compliance contexts, guardrails are not optional refinements — they are professional risk management. They help ensure that AI-generated content supports informed human judgment without appearing to replace it, which matters particularly when documents may be reviewed in grievance proceedings, administrative hearings, or in response to public records requests. Constraints that seem limiting at the prompting stage often prevent significant problems at the review stage.

<p><b>Weak Prompt:</b> "Look at this discipline file and tell me what you think. Did we do the right thing here?"</p>
---------------------------------------------------------------------------------------------------------------------------

**Improved Prompt:**

"Review the attached disciplinary documentation and policy. Your task is limited to: (1) identifying which policy provisions were cited in the disciplinary action; (2) listing the documented facts from the investigation file that correspond to each provision; and (3) identifying any factual gaps — information that the policy requires but that does not appear in the file. Do not evaluate the appropriateness of the discipline imposed, offer a legal opinion, or recommend any action. Present findings in a three-column table: Policy Provision | Supporting Documentation | Documented Gaps."

**Why the Improved Prompt Works:**

The explicit prohibition against evaluating the discipline or recommending action keeps the output in its proper analytical lane. The system's job here is to map evidence, not adjudicate. In environments where AI-generated documents may be seen by union representatives or appear in appeal records, limiting scope is as important as defining it.

## 05. Limit the Sources of Information

### Concept Overview

AI systems are trained on large general datasets and will draw on that background knowledge unless instructed otherwise. In many contexts, that breadth is an asset. In legal, contractual, and policy-specific contexts, it is a liability. A system tasked with analyzing whether a personnel action complies with policy may supplement the actual policy language with general HR principles that do not apply to the organization's specific provisions, producing an analysis that appears authoritative while actually reflecting external assumptions rather than the governing documents.

Bounding the sources of information is a straightforward technique for managing this risk. When a prompt instructs the system to rely only on the attached documents and to flag questions that those documents do not resolve, the resulting analysis stays grounded in the materials actually under review. This approach is essential when working with collective bargaining agreements, local ordinances, personnel policies, and position descriptions, where specific language governs the outcome and general principles may actively mislead.

Source limitation also improves transparency. When the system is constrained to the documents provided, its analysis can be traced directly to identifiable text. That traceability matters when conclusions must be explained to employees, union representatives, department heads, or elected officials, and when the reasoning behind a decision must withstand adversarial scrutiny. If the system identifies a gap — a question the documents do not answer — that gap itself is valuable information requiring human follow-up.

**Weak Prompt:**

"Tell me whether we can require employees to use accrued leave during an unpaid suspension."

**Improved Prompt:**

"Using only the attached personnel policy and collective bargaining agreement — and no outside sources — identify every provision that addresses leave usage during disciplinary suspensions. Quote the relevant language from each provision. Then identify any questions this scenario raises that the attached documents do not clearly resolve. Do not apply general employment law principles or practices from other jurisdictions."

**Why the Improved Prompt Works:**

Without source limitation, the system may draw on general HR practice, statutes from other jurisdictions, or common CBA language that differs from the organization's actual agreement. Constraining the analysis to the attached documents ensures findings reflect the governing text. Flagging gaps that the documents do not resolve is equally valuable: it identifies where human legal or labor relations judgment is required.

## 06. Extract Information Before Attempting Analysis

### Concept Overview

A common failure in AI-assisted analysis is asking the system to evaluate whether something meets a standard without first establishing what that standard actually is. When prompts ask whether a situation "meets policy requirements" or "qualifies under the contract," the system may apply general assumptions about what those standards typically say rather than what the governing documents state. The resulting analysis can appear specific and well-reasoned, yet actually reflect approximations.

A more reliable approach separates extraction from analysis in two explicit steps. The first prompt asks the system to extract the relevant criteria or policy language directly from the source document, without yet applying those criteria to anything. The second prompt presents the facts of the situation and asks the system to compare them to the extracted standards. This sequence prevents the system from blending general knowledge with document-specific content in ways that are difficult to detect in the output.

The value of this approach extends beyond accuracy. It creates an auditable analytical trail: the criteria are stated and sourced before the analysis begins, making it straightforward for reviewers to evaluate whether the correct standards were applied. In classification, compensation, and disciplinary contexts — where the defensibility of a decision often depends on demonstrating that documented criteria were applied systematically — this structural transparency strengthens both the analysis and the record.

#### Weak Prompt:

"Does this person qualify for the senior analyst classification based on their duties?"

#### Improved Prompt:

"Step 1: From the attached classification specification for Senior Analyst, extract and list every minimum qualification and every duty criterion used to distinguish this class from Analyst. Quote the exact language for each criterion. Do not apply these criteria to anything yet. Step 2 (after I confirm Step 1): I will provide the position description we are evaluating, and I will ask you to compare each criterion from Step 1 against the documented duties, noting whether each criterion is met, not met, or cannot be determined from the available information."

#### Why the Improved Prompt Works:

Separating extraction from analysis forces the criteria to be stated explicitly before any judgment is applied. This prevents the system from silently substituting general classification principles for the organization's actual standards. It also creates a reviewable record: when the criteria are listed before the comparison, a skeptical audience can verify that the right standards were applied to the right facts.

## 07. Ask the System to Identify Assumptions and Uncertainty

### Concept Overview

AI-generated responses are presented in a consistent, confident register that does not naturally signal the difference between well-supported conclusions and inferences drawn from incomplete information. For professionals who use these outputs to inform decisions, the appearance of certainty can be genuinely misleading. A workforce analysis that presents turnover trends without acknowledging data limitations, or a classification review that draws comparisons without noting the assumptions behind them, can lead decision-makers to treat tentative findings as established conclusions.

A straightforward prompting technique addresses this directly: instruct the system to distinguish explicitly among what the data shows, what the evidence suggests, and what additional information would be needed to increase confidence in any conclusion. Applied to a turnover analysis, this might mean separating documented exit patterns from plausible explanations and identifying what additional data —

exit survey results, benchmark comparisons, compensation data — would be required to confirm or refute each explanation.

This practice aligns AI output with the analytical standards that public-sector professionals are already expected to meet. Audit-ready reports, board briefings, and labor relations documents all benefit from a clear distinction among fact, inference, and speculation. Building that structure into the prompt produces it in the output, rather than requiring the user to impose it after the fact. It also reduces the risk of presenting borrowed confidence — where the system's authoritative tone obscures the actual evidentiary basis for a finding.

<p><b>Weak Prompt:</b> "Look at this turnover data and tell me what's causing people to leave."</p>
<p><b>Improved Prompt:</b> "Analyze the attached turnover report covering the past three fiscal years. Structure your findings in three labeled sections: (1) Observed Patterns — what the data directly shows, with figures; (2) Possible Explanations — inferences that the data is consistent with but does not confirm, labeled as inference; (3) Information Needed — what additional data or analysis would be required to confirm or rule out each explanation. Do not present inferences as conclusions."</p>
<p><b>Why the Improved Prompt Works:</b> Requiring three labeled sections forces the system to separate documented findings from interpretive judgment before submitting the analysis. This structure prevents the most common AI analysis failure: confidently stated inferences presented as if they were established facts. For boards, elected officials, and senior administrators who may act on these findings, the distinction between what data shows and what it suggests is not a formatting preference — it is a professional obligation.</p>

## 08. Break Complex Tasks into Smaller Steps

### Concept Overview

Large, complex documents frequently yield inconsistent results when analyzed or summarized in a single prompt. The system may give disproportionate attention to early sections, compress or skip later material, or produce a summary that reflects the document's general shape rather than its actual content. For professionals working with lengthy collective bargaining agreements, multi-part position descriptions, or comprehensive policy manuals, these inconsistencies can result in material gaps that are not obvious from reading the output.

Dividing complex tasks into sequential, bounded steps produces more complete and accurate results. A classification analysis might begin with a prompt identifying the major duty categories in a position description, followed by a separate prompt comparing each category against the relevant classification criteria, with a final prompt integrating those comparisons into an overall finding. Each step is narrower and more tractable than the full analysis taken at once, and the final product reflects appropriate attention to each component.

This approach also creates natural checkpoints. After each step, the user can review the intermediate output and correct any mischaracterizations before they propagate into subsequent analysis. In high-stakes contexts — classification appeals, disciplinary investigations, compensation equity reviews, contract administration — this step-by-step structure improves both accuracy and defensibility. The additional time required is modest; the reduction in revision and rework is substantial.

<p><b>Weak Prompt:</b> "Can you summarize this whole union contract for me? I need to know what it says."</p>
-------------------------------------------------------------------------------------------------------------------

**Improved Prompt:**

"I need to analyze this collective bargaining agreement in stages. Step 1: List every article title and its page range so I can see the full structure. Stop after completing Step 1 and wait for my instruction. I will then ask you to summarize specific articles one at a time before we produce any overall summary."

**Why the Improved Prompt Works:**

Starting with structure before content prevents the system from compressing a 100-page agreement into a misleading overview that loses critical detail. Staging the work allows the user to prioritize the articles most relevant to a current issue, review each section for accuracy before building on it, and produce a final summary grounded in verified intermediate work rather than a single pass across a complex document.

## 09. Refine Prompts Through Iteration

### Concept Overview

The first response from an AI system should generally be treated as a draft rather than a finished product. Even a well-constructed prompt will sometimes produce output that is technically responsive but not quite right in emphasis, structure, or depth. Treating the initial output as a starting point for refinement — rather than as a completed deliverable — significantly improves results without requiring expertise in prompt construction.

Iteration is not a sign that the original prompt failed; it is a normal part of an effective AI-assisted workflow. A summary that is too broad can be narrowed. An analysis that lacks sufficient specificity can be refined by adding an additional constraint. An incorrectly formatted document can be reorganized. Each refinement takes a fraction of the time required to produce the original output and moves the work product progressively closer to the professional standard required.

Public-sector administrators who manage staff work apply this discipline instinctively: they review drafts, provide specific feedback, and expect revision. The same approach translates directly to AI direction. Feedback such as "condense this to three paragraphs for a county board briefing" or "remove all inferential language and restate the findings as factual observations only" functions the same way a supervisor's margin note does — it produces a more targeted revision with minimal additional effort. The professional skill is not in writing a perfect initial prompt; it is in reading the output and knowing specifically what to correct.

**Weak Prompt:**

"This isn't quite right. Can you make it better?"

**Improved Prompt:**

"The summary you produced is too general for our purpose. Please revise it with the following changes: (1) reduce total length to no more than three paragraphs; (2) lead with the fiscal impact figure, not the program description; (3) remove all language that qualifies or hedges the findings — state only what the data confirms; (4) add a final sentence identifying the decision the board is being asked to make. Keep the table format."

**Why the Improved Prompt Works:**

Effective iteration requires specific, actionable feedback, not general dissatisfaction. The revised instruction names four concrete changes, each of which the system can execute directly. This mirrors effective supervision: a good revision note tells the writer exactly what to change and why, not simply that the current version is insufficient. Iterating with precision produces usable output faster than restarting from a new prompt.

## 10. Confirm That the Task Was Understood

### Concept Overview

Even carefully constructed prompts can be interpreted in unexpected ways. A request for a "summary of the grievance" may yield a narrative of events rather than a structured analysis. A request to "review the position description for classification accuracy" may produce an evaluation of writing quality rather than a classification analysis. These mismatches are not always apparent from the output itself, which often appears coherent regardless of whether it addresses the intended question.

A simple preemptive technique reduces this risk: ask the system to restate its understanding of the task before producing the substantive response. This instruction costs almost no additional time and functions as a brief alignment check. If the system's restatement diverges from the intended assignment, the misalignment can be corrected before the system produces a full response based on an incorrect interpretation.

This technique is most valuable for complex, multi-part assignments and for instructions that involve simultaneous constraints and conditions — for example, prompts that specify both what to include and what to exclude, or that require a specific analytical structure while also limiting source material. It functions as the AI equivalent of confirming instructions with a staff member before work begins: a brief investment in clarity that prevents larger corrections later. For high-stakes documents — those that will be shared with elected officials, labor representatives, or legal counsel — the few seconds required to verify task understanding are well spent.

#### Weak Prompt:

"Here's the position description. Go ahead and analyze it."

#### Improved Prompt:

"Before producing your analysis, briefly restate in two or three sentences what you understand this task to require: what document you are working with, what analytical question you are answering, what criteria you are applying, and what constraints limit your response. Wait for me to confirm your restatement before proceeding."

#### Why the Improved Prompt Works:

Requiring a restatement before analysis exposes interpretive mismatches before they become fully drafted outputs. The four elements named in the restatement (document, question, criteria, constraints) are the same elements that determine whether the analysis will be useful. Confirming them takes thirty seconds; correcting a fully developed response built on a misunderstanding takes considerably longer, and in some contexts creates a record of analysis the organization did not intend to produce.

### Final Observation

Generative AI is a powerful tool, but its usefulness depends entirely on how it is directed. The ten practices described in this guide do not require technical expertise. They require the same professional habits that lead to effective supervision, clear communication, and defensible analysis.

When professionals approach prompting as a structured assignment rather than a casual conversation, the quality of AI-generated outputs improves substantially. The prompts that produce the most reliable results are not the most elaborate; they are the most clearly reasoned, written by professionals who know what they need and have stated it with precision.

That skill already exists in the professionals this guide is written for. The task is simply to apply it in a new context.