
Prompting in Practice

A Hands-On Workshop for Public-Sector LR and HR Professionals

National PELRA 55th Annual Training Conference



Disclaimer, Use, Distribution, and AI Assistance Notice

This session provides professional education on prompting techniques for generative AI tools. Nothing in this presentation constitutes legal advice. Attendees should consult their own legal counsel on the application of public records, discovery, and data privacy obligations in their jurisdiction.


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 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.

Patrick Glynn is the founder and principal consultant of BoldPath Consulting. He brings over **30 years of experience working with and advising Wisconsin local governments**, including service as a county human resources director and more than a decade as a public-sector HR and management consultant. His work spans labor relations, compensation and classification, workforce planning, management and leadership training, generative AI, and organizational assessment, with a consistent focus on solutions that are fair, defensible, and workable in real public-sector environments.

Patrick's approach is grounded in an **understanding of how local governments actually operate**: political scrutiny, limited resources, union dynamics, public accountability, and the need for decisions that can be explained and sustained over time. He is known for a collaborative, candid style that emphasizes clarity over spectacle and long-term institutional health over short-term wins. In recent years, he has also **helped public organizations carefully consider the role of generative AI in government work**, treating it as a governance and management issue rather than a technology trend, and integrating it into existing professional, ethical, and organizational frameworks.



Patrick Glynn

Founder/Principal Consultant

(920) 522-2413

www.boldpathconsulting.com

Before We Begin

Session Schedule

10:00	Opening and Orientation
10:08	The Ambiguity Penalty
10:18	Practices 1-5: Building a Structured Prompt
10:45	Verification, Records, Accountability
10:55	Transition
11:00	Break
11:15	Midstream Recap
11:20	Context Management
11:35	Live Demonstrations
12:05	Iteration: From Draft to Done
12:20	Closing and Next Steps

Logistics

- Devices welcome: manage notifications
- Questions welcome throughout
- Materials packet: training handout, use case compendium, governance framework
- “Office Hours”: announced during closing



The Skills You Already Have Are the Skills This Session Teaches

- You write assignments that specify scope, format, and audience
- You review staff work for accuracy before it leaves your desk
- You distinguish fact from inference in every document you sign
- These habits produce effective AI results

This is not a technology session.
It is a professional skills session in a new context.



Your documents



Your standards



Your discipline

The Ambiguity Penalty



**If you did not give the model the fact, it may invent one.
If you approve it, the organization owns it.**



What the Penalty Looks Like in Labor Relations and HR

Grievance Analysis

An unsupported conclusion about the merits becomes a discoverable document the union can request

Disciplinary Summary

A credibility characterization creates a record the arbitrator may weigh against the employer

CBA Summary

Paraphrased contract language misstates a provision and undermines the employer's position

Board Briefing

Inference presented as established fact exposes officials to unsupportable public statements

Each of these becomes a permanent record the moment it leaves your desk.



Three Foundation Practices

1

Prioritize Clarity Over Cleverness

Effective prompting functions like a well-structured staff assignment. Name the deliverable. Identify the audience. Specify focus areas. Clarify constraints.

2

Specify the Output Format

Format directs the system to organize information for your workflow. A grievance summary under structured headings produces dramatically different utility than narrative prose.

3

Define the Audience and Purpose

A briefing for elected officials requires different depth and framing than a procedural memo for supervisors. Specifying both prevents output that is accurate but professionally unsuitable.

These practices are intuitive for professionals who already write board memos and staff assignments. The task is to apply them in a new context.



Prompt Evolution

BEFORE

Look at this grievance and the contract and tell me if we violated the overtime provisions.



AFTER

Using only the attached grievance filing and collective bargaining agreement, produce a structured analysis with four sections: (1) Allegation -- one sentence restating the union's claimed violation in neutral language; (2) Relevant Contract Language -- quote the specific articles and sections at issue, using exact contract language; (3) Documented Facts -- bullet list of confirmed facts from the grievance file and be formnse management, and management records only, with no inference; (4) Analytical Gaps -- bullet list of factual questions that must be resolved before a Step 2 response can be finalized. Do not characterize the merits of the grievance, recommend a response, or offer a legal opinion.

Establish Guardrails and Constraints

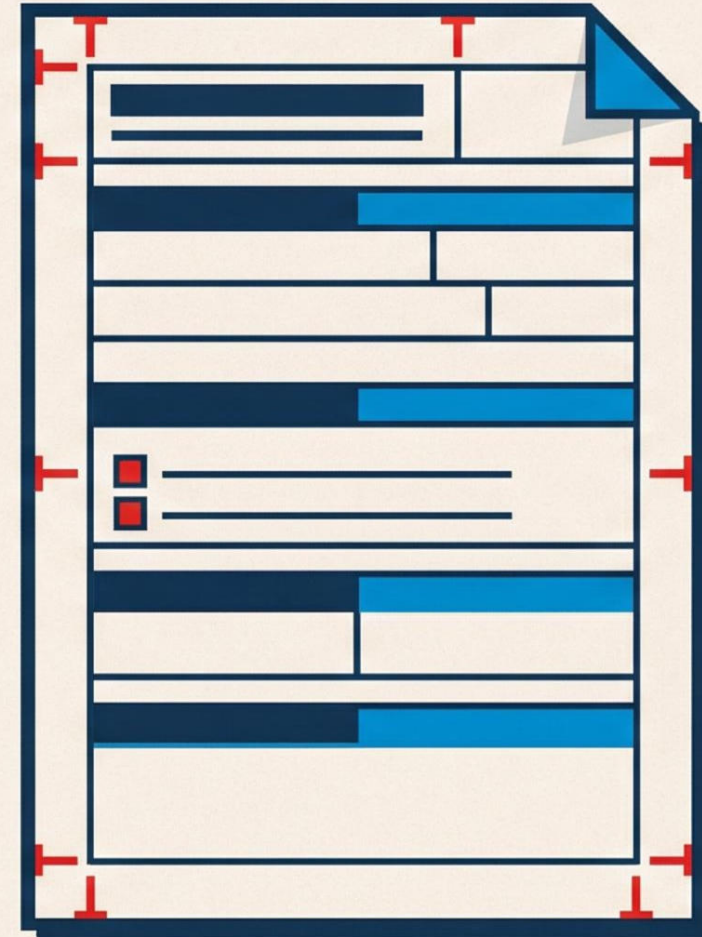
The Principle

Tell the system what it must not do with the same precision you use to tell it what to do.

In LR/HR Contexts

- Do not evaluate the appropriateness of discipline
- Do not offer a legal opinion
- Do not recommend any action

These constraints are not limitations. They keep the output in its proper analytical role and limit the scope of what becomes discoverable.



Disciplinary Documentation: The Weak Prompt

WEAK PROMPT

Look at this discipline file and tell me if we have enough to suspend this employee.

No scope. No structure. No constraints. No specified output.

The model will generalize, infer, and produce a confident answer to a question you did not carefully ask.



Disciplinary Documentation: The Improved Prompt

IMPROVED PROMPT

Review the attached disciplinary documentation, investigation file, and personnel policy. Your task is limited to: (1) identifying which specific policy provisions were cited in the recommended disciplinary action; (2) listing the documented facts from the investigation file that correspond to each cited provision; and (3) identifying any documented gaps -- information that the policy requires for the recommended level of discipline but that does not appear in the file.

Present findings in a three-column table: Policy Provision, Supporting Documentation, Documented Gaps. Do not evaluate the appropriateness of the recommended discipline, offer a legal opinion, or recommend any action. If the progressive discipline policy requires documentation of prior corrective steps, note whether those steps are present in the file.

Same task. Defined scope. Explicit constraints. Specified output.



The Output: Documentation Gap Analysis

Policy Provision	Supporting Documentation	Documented Gaps
Sec. 4.2: Attendance and Punctuality	Three written warnings dated 1/15, 3/22, 6/10; supervisor log entries	No documentation of verbal counseling step required by Sec. 2.1
Sec. 4.5: Insubordination	Witness statement from supervisor dated 7/3; employee written response	No corroborating witness statement referenced in investigation notes
Sec. 2.1: Progressive Discipline Sequence	Written warnings present for attendance; no prior discipline for insubordination	Recommended suspension skips verbal verbal and written warning for insubordination charge
Sec. 6.3: Investigation Requirements	Investigation notes, two witness interviews	Third witness identified in complaint not interviewed; no explanation documented
Sec. 3.4: Employee Notification	Pre-disciplinary hearing notice dated 7/10	None identified

The AI maps evidence. The professional makes the judgment.

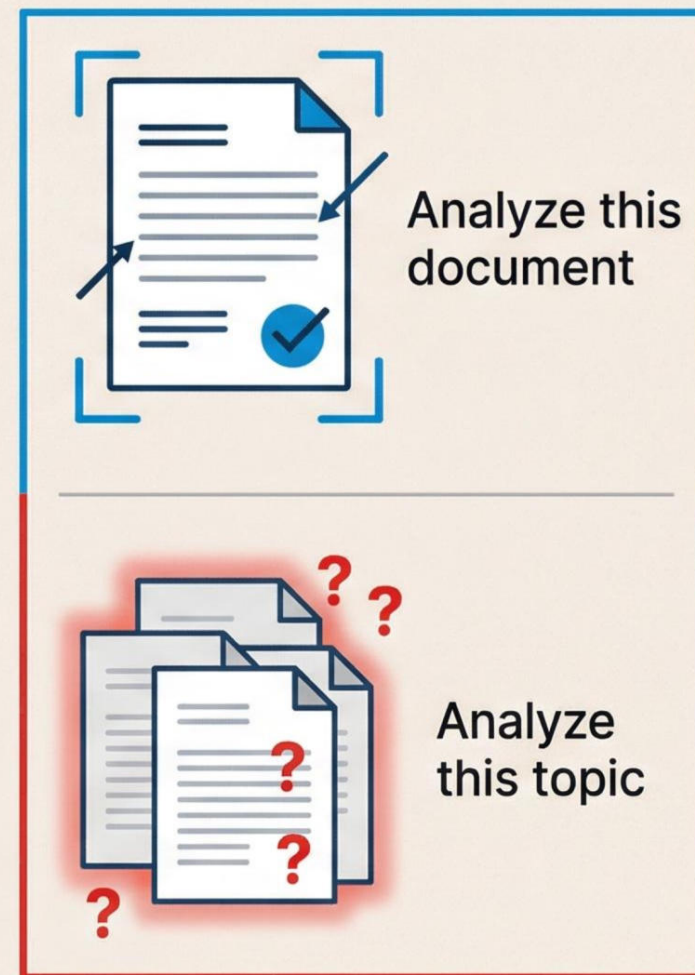


Limit the Sources of Information

The distinction between 'analyze this document' and 'analyze this topic' is the difference between a useful work product and a liability.

- Constrain to attached materials only
- Require the system to flag gaps rather than fill them
- Create an auditable trail tied to governing text

For professionals who work with CBAs, investigation files, and classification specifications, this is the single most consequential behavioral change.

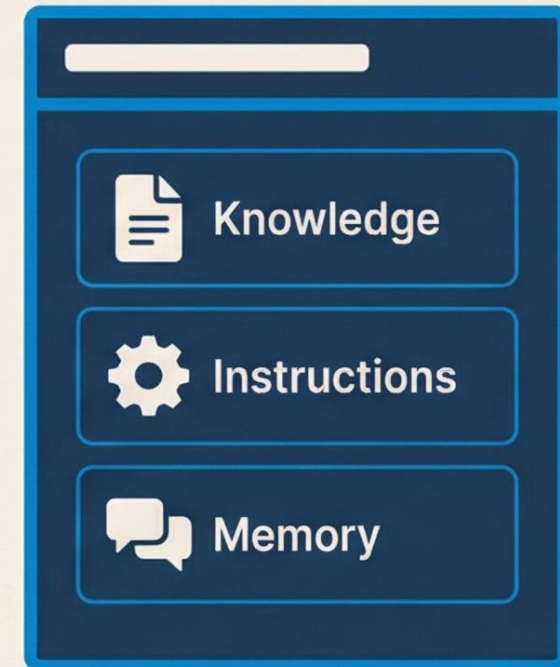


Work in a Project, Not a Single Conversation

Why Projects Change the Quality of Your Output

- **Project memory persists** across conversations -- the system retains context from prior work without you restating it each time
- **Project knowledge:** upload your CBA, personnel policies, classification specifications, and investigation templates once and reference them in any conversation within the project
- **Project instructions:** your custom instructions, analytical standards, and constraints apply automatically to every conversation -- no need to rebuild your professional context from scratch
- **Workstream separation:** a grievance analysis project does not bleed into your classification review project -- context stays contained and relevant
- **Reduced prompting burden:** when the project already knows your jurisdiction, your standards, and your governing documents, your prompts can focus on the task rather than re-establishing the baseline

A single conversation is disposable. A project is a persistent professional workspace.



Verification Is Not Optional

The system does not verify its own work.

The professional who submits the document is accountable for its content, regardless of who or what produced the first draft.



Review every output



Verify against originals



Your name, your responsibility

If you cannot verify it, do not use it.

AI Work Products Are (Likely) Discoverable Records

Questions You Should Be Asking

- Are AI-generated analyses created in the course of official duties public records in your jurisdiction?
- Could the prompts themselves be considered records?
- If a prompt asks the AI to evaluate whether discipline was appropriate, could the prompt and the resulting analysis be discoverable?
- Does your jurisdiction exclude draft or preparatory materials from public records disclosure, and would that exclusion apply to AI-assisted work products?
- The guardrail technique in Practice 4 limits the analytical scope of what you create -- and that may matter when someone asks to see it



These questions do not have uniform answers. Open records obligations, discovery rules, and draft-document exemptions vary by jurisdiction. Talk to your counsel before you need to.

Part One Complete

Five practices. One professional standard.

Structured prompting. The Ambiguity Penalty. Verification.

Part Two applies these to the work you do every day.

- 1 Conversation Management
- 2 Live Demonstrations
- 3 Iteration and Refinement

Return at 11:15. Part Two begins with a brief recap for anyone joining.




Part Two

11:15 a.m. to
12:30 p.m.

Applied Workflows and Iteration

From principle to practice



Context Management

How to maintain
control across long
conversations

Live Demonstrations

Classification,
turnover, and CBA
analysis in real time

Iteration

Directive feedback,
adversarial review,
and knowing when
to restart

Closing

Your toolkit for
independent
practice

Everything in Part Two builds on the five practices from Part One.
The training handout is your reference.

Where We Are

In the first hour, we covered five prompting practices that produce structured, defensible AI outputs:

- Clarity over cleverness
- Specify the output format
- Define the audience and purpose
- Establish guardrails and constraints
- Limit the sources of information

Anchored by the Ambiguity Penalty and the principle that verification is a professional obligation.

Part Two: How does prompting discipline work in practice when you manage a complex task across multiple turns?




AI Has a Finite Working Memory


AI systems process conversation within a fixed context window. In long conversations, earlier material loses influence and the system may drift from your original instructions.

This is not a bug. It is a mechanical limitation you must manage.

The longer the conversation, the more actively you must maintain control of what the system is working from.




CONTEXT WINDOW CAPACITY




8K TOKENS




32K TOKENS



128K TOKENS



200K TOKENS



1M+ TOKENS

Five Techniques for Keeping Conversations on Track

1 Maintain Working Notes

Restate key decisions, definitions, and constraints. Paste them back in when drift occurs.

2 Break Complex Tasks into Steps

Request sequential outputs rather than a complete work product in a single prompt.

3 Label Your Information

Background, Draft Text, Policy Language, Reference Notes -- labels prevent the system from blending sources.

4 Extract Before Drafting

Ask the system to list key facts or criteria before writing. Writing from an explicit list beats implicit recall.

5 Restate Active Instructions

Periodically ask the system to restate all instructions as a single clean set. Use that version to continue.

Custom Instructions: Your Persistent Professional Context

EXAMPLE

I am a labor relations officer for a mid-sized public-sector jurisdiction. Applicable statutory framework: [state statute]. I work with multiple collective bargaining agreements.

Analytical standards: distinguish fact from inference in all analysis. Do not characterize the merits of any grievance or disciplinary action. Cite contract provisions by article and section number. Quote contract language exactly; do not paraphrase.

Preferred output format: structured sections with headings.

Common constraints: do not offer legal opinions, do not recommend specific outcomes, do not expand beyond assigned scope.

Set this once. It applies to every conversation.

Available in ChatGPT, Claude, and other platforms as custom instructions or system prompts. ✨

Classification Analysis with Extraction Before Evaluation

Practice 6: Separate Extraction from Evaluation

Step 1

Extract every minimum qualification and duty criterion from the classification specification. Quote exact language. Do not apply criteria yet.

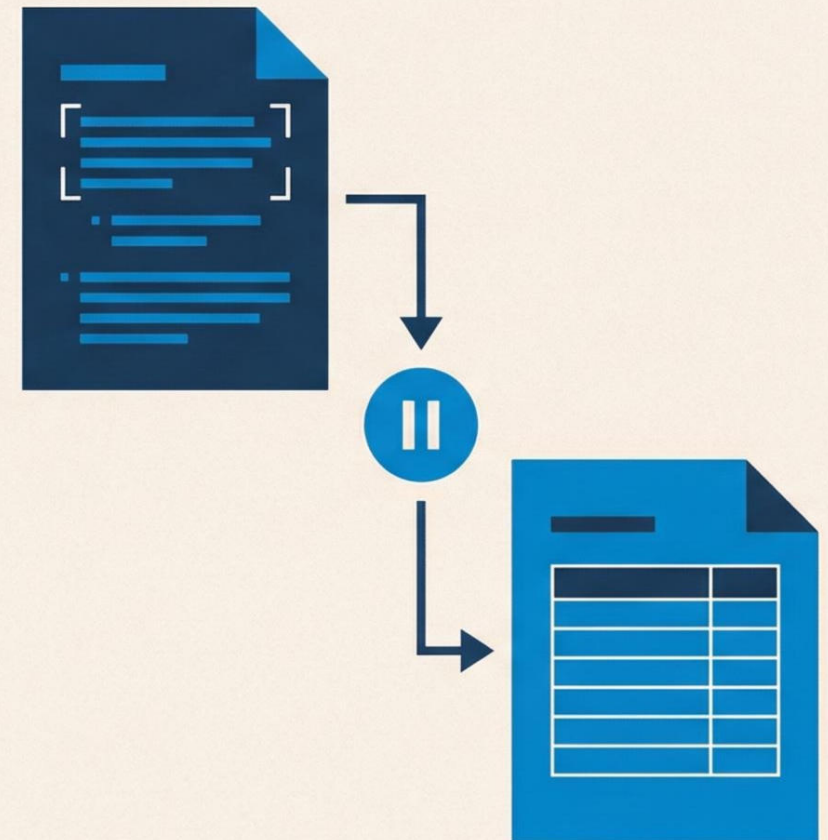
Checkpoint

Ask the system to restate its understanding of the task (Practice 10).

Step 2

Compare each criterion against the position description.
Note: met, not met, or cannot be determined.

The two-step separation creates a reviewable trail the classification appeal board can follow.



Classification Analysis: Step 1

STEP 1 PROMPT

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. Do not summarize or paraphrase. Present each criterion as a numbered item with the exact specification in quotation marks.

**The system extracts. It does not yet evaluate.
That separation is the point.**



Classification Analysis: Step 2

STEP 2 PROMPT

Before proceeding, restate your understanding of this task in two sentences. Then: using the criteria you extracted in Step 1, compare each criterion against the documented duties in the attached position description. For each criterion, note whether it is clearly met, clearly not met, or cannot be determined from the available information. For any criterion marked cannot be determined, state what additional information would be needed.

Present as a table with columns: Criterion, Position Description Evidence, Determination.

Extraction first, then evaluation. The same discipline you apply to staff work. ✨

Turnover Analysis with Labeled Uncertainty

Practice 7: Label What You Know, What You Infer, and What You Need

■ Observed Patterns

What the data directly shows, with figures

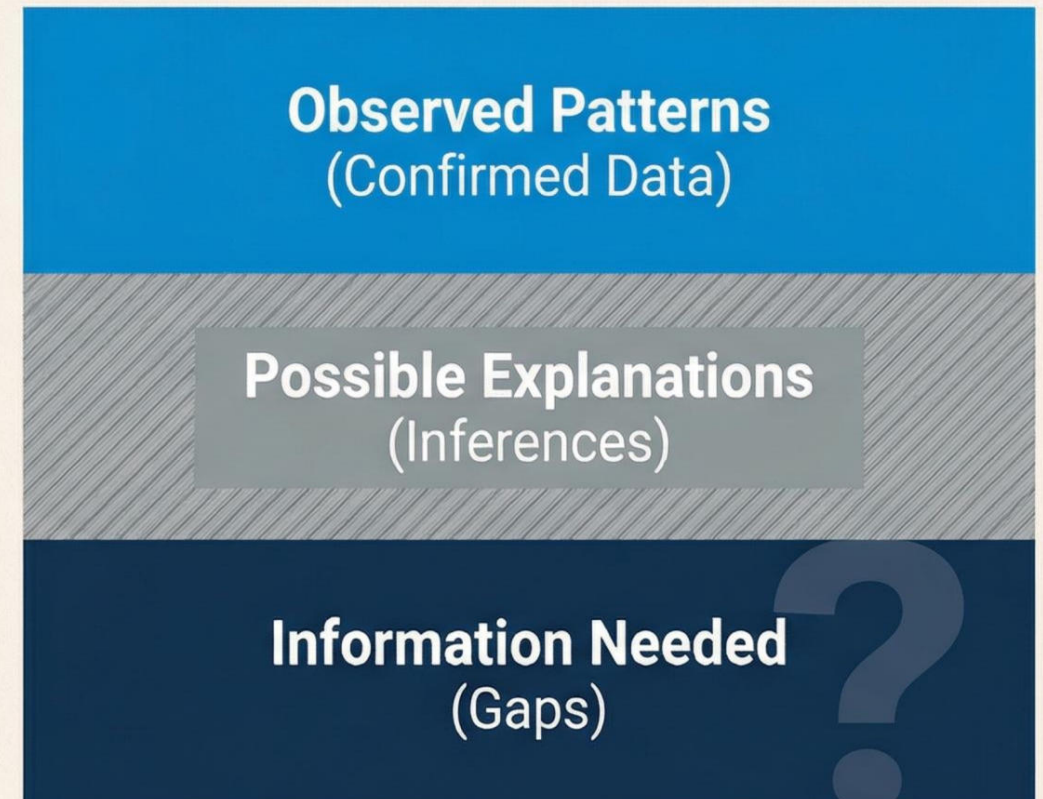
■ Possible Explanations

Inferences the data is consistent with but does not confirm -- explicitly labeled as inference

■ Information Needed

Additional data required to confirm or rule out each explanation

For elected officials, the distinction between what data shows and what it suggests is a credibility requirement.



Turnover Analysis: The Prompt

PROMPT

Analyze the attached turnover report covering the past three fiscal years. Structure your findings in three clearly labeled sections: (1) Observed Patterns -- what the data directly shows, with specific figures and percentages; (2) Possible Explanations -- inferences that the data is consistent with but does not confirm, explicitly labeled as inference; (3) Information Needed -- what additional data or analysis would be required to confirm or rule out each.

Do not present inferences as conclusions. This analysis will be presented to an elected governing body.

Three sections. Three levels of certainty. One professional standard.

Turnover Analysis: The Output Structure

■ OBSERVED PATTERNS

Overall three-year separation rate: 14.2%. Highway Department: 22.1% (highest). Separations concentrated in 2-5 year tenure bracket (47% of all departures). Voluntary resignations: 68% of total separations.

■ POSSIBLE EXPLANATIONS *labeled as inference*

The concentration in the 2-5 year bracket is consistent with, but does not confirm, compensation competitiveness as a factor. The Highway Department rate may reflect regional labor market conditions for CDL-certified positions.

■ INFORMATION NEEDED

Exit interview data by department. Comparable-jurisdiction compensation data for CDL and highway positions. Tenure-at-separation broken out by voluntary vs. involuntary.

What the data shows. What it suggests. What you still need.

Collective Bargaining Agreement: Staged Analysis

Practice 8: Stage Complex Analysis

Step 1: Map the Structure

List every article title and page range. Stop and wait.

Step 2: Summarize a Targeted Provision

Summarize a specific article selected based on the grievance scenario.

Step 3: Apply to the Question

Compare the summarized provision against documented grievance facts.

Staging prevents the system from compressing a complex contract into a misleading overview.



CBA Staged Analysis: The Prompt Sequence

STEP 1

I need to analyze the economic provisions of the attached collective bargaining agreement in stages. List every article that contains an economic provision. For each article, note the article number, title, and page range. Stop after completing Step 1 and wait for my instructions.

Review the inventory before going deeper.

STEP 2

Summarize Article [X] -- [Title]. Quote key terms, formulas, and dollar amounts exactly as written. Do not paraphrase contract language.

Verify accuracy of the summary against the contract.

STEP 3

Compare the provisions in Article [X] against the following documented facts from the grievance: [facts]. Note where the facts align with contract language and where they diverge.



From Draft to Done

Practice 9: Iterate with the Same Discipline You Use to Supervise Staff Work


Effective iteration is not:

'This is not quite right. Can you make it better?'

Effective iteration is:

- Naming specific changes
- Directing structural reorganization
- Requiring evidentiary standards

If you would not accept that feedback from a supervisor, do not give it to the system.



Directive Iteration: Three Examples

EXAMPLE 1 Reduce the summary to no more than three paragraphs. Lead with the fiscal impact figure, not the program description. Remove all qualifying language and state only what the data confirms. Add a final sentence identifying the decision the board is being asked to make.

EXAMPLE 2 The grievance summary blends documented facts with inferences. Revise it: under **Key Facts**, include only items that appear in the investigation file. Move everything else to a new section titled **Unresolved Questions**. Do not change the Allegation or Contract Language sections.

EXAMPLE 3 This classification comparison is too long for the appeal packet. Condense it to a single table with four columns: Criterion, Class Spec Language, Position Description Evidence, and Determination (Met / Not Met / Insufficient Information). Remove all narrative.

Specific. Structural. Directive. The same standard you hold your staff to.



Red Team: Test Your Output Before the Other Side Does

The Technique

Ask the system to switch roles and review its own output as an adversary. The model evaluates its own work from a hostile perspective, identifying weaknesses the original prompt did not anticipate.

Personas to assign: Skeptical arbitrator. Union advocate. Hostile auditor. Opposing counsel.

Example Prompt

Review the disciplinary documentation analysis you just produced. Assume the role of the union's advocate preparing for arbitration. Identify every weakness in the employer's documentation, every gap an arbitrator might find troubling, and every argument the union could make based on the documented gaps. Be aggressive and specific.

If the red team finds the weakness, you found it first.



Fresh Slate: Knowing When to Restart

When to Start a New Conversation

- The system keeps returning to a pattern you have already rejected
- Earlier errors are contaminating later output despite corrections
- The conversation has grown long enough that drift is evident

How to Restart Cleanly

Summarize the current state of the analysis: what has been decided, what remains open, and what constraints apply. Paste that summary as the opening of a new conversation. You lose nothing except the accumulated drift.

Knowing when to restart is as important as knowing how to iterate.



Your Take-Home Toolkit

1 Ten Prompting Practices

Your day-to-day desk reference. Covers all ten practices with weak/improved prompt pairs. Keep this accessible during your first weeks of regular AI use.

2 Use Case Compendium

Fifteen ready-to-adapt prompts for grievance analysis, CBA interpretation, arbitration prep, classification, compensation, investigations, and board communications. Copy, adapt, your own documents this week.

3 AI Governance Framework

Eight-pillar policy framework with bracketed placeholders for your jurisdiction. Designed for adoption by governing bodies. Useful whether your organization has an AI policy or not.

**The Use Case Compendium is your bridge to independent practice.
Start with one scenario this week.**

Session C21 on Wednesday addresses how AI changes workforce strategy at the organizational level.



Three Questions for Monday Morning

1

What Will I Try First?

Identify one low-risk task from the Use Case Compendium you will try this week with your own documents.

2

What Do I Need to Clarify?

What tool approval, records question, or policy gap must be resolved before you can proceed responsibly?

3

What Conversation Do I Need to Have?

Who in your organization should be part of the AI governance discussion and is not yet?

The prompting discipline you learned today transfers to any platform and any subject matter.

The professional habits are yours. The tools will keep changing. The discipline does not.



Patrick Glynn

Founder and Principal Consultant
BoldPath Consulting, LLC

(920) 522-2413

patrick@boldpathconsulting.com

www.boldpathconsulting.com

[linkedin.com/in/patrickwglynn](https://www.linkedin.com/in/patrickwglynn)

Office hours: [time and location TBD]

The Use Case Compendium is your bridge to independent practice. Start with one scenario this week.

