

PROMPTING IN PRACTICE

Participant's Guide

A Hands-On Workshop for Public-Sector Use of Generative AI
Labor Relations and Human Resources Applications

National PELRA 55th Annual Training Conference

Charlotte, North Carolina

Monday, 10:00 a.m. to 12:30 p.m. | Providence 1 Ballroom

Patrick Glynn

Founder and Principal Consultant



patrick@boldpathconsulting.com | (920) 522-2413 | www.boldpathconsulting.com

Using This Guide

This guide follows the slide deck in sequence. For each block of the workshop, it gives you a short orientation, the principle the block is built around, and the full text of the prompts demonstrated during the session. The prompts are presented so you can copy, adapt, and apply them to your own documents this week.

Bracketed terms in the prompts mark variables you substitute with your own scenario. A prompt written for a classification appeal or a grievance under one contract is still a template; the structure carries forward, only the specific language changes.

This is a working guide, not a transcript. It captures the prompting material from the session, not the live discussion, the examples raised from the floor, or the questions answered in office hours.

Part One: Prompting as a Professional Skill

10:00 a.m. to 11:00 a.m.

Part One builds a single claim: effective prompting is not a technology skill. It is the professional discipline you already apply when you assign a task to staff, review a draft for legal sufficiency, or prepare a document for adversarial scrutiny. The five practices in this part of the session give that discipline a structure.

Opening and Orientation

The techniques in this session apply to ChatGPT, Claude, Gemini, Copilot, and any other general-purpose generative AI tool. The demonstrations use ChatGPT for visual consistency; the prompting principles are platform-agnostic. No technical background is required.

The Ambiguity Penalty

An ambiguous prompt does not fail visibly. It produces a confident, plausible response. A human approves that response under time pressure. Once it leaves your desk, the organization owns it. That is the penalty.

What it looks like in labor relations and HR work:

- 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. The five practices that follow are designed to prevent it.

Practices 1 to 3: Clarity, Format, Audience

These three practices name habits you already apply when assigning work to staff. They need naming more than they need teaching.

- Clarity over cleverness. Name the deliverable. Specify the focus. Describe the constraint. A prompt should read like a well-structured staff assignment, not a casual conversation.
- Specify the output format. Tell the system how to organize the output for your workflow. The same content organized as narrative prose and the same content organized under structured headings produce dramatically different professional utility.
- Define the audience and purpose. A briefing for elected officials requires different depth, terminology, and framing than a procedural memo for supervisors. Specifying both prevents output that is accurate but professionally unsuitable.

Prompt Evolution: Grievance and Contract Review

The first three practices are visible in the difference between these two prompts for the same underlying task.

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 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.

Practice 4: Establish Guardrails and Constraints

Tell the system what it must not do with the same precision you use to tell it what to do. In labor relations and HR contexts, the standard guardrails are:

- Do not evaluate the appropriateness of the recommended action.
- Do not offer a legal opinion.
- Do not recommend any specific outcome.

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

Disciplinary Documentation Review

This is the anchor demonstration for Part One. The same file, the same question, and two very different work products.

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.

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 AI maps the evidence. The professional makes the judgment.

Practice 5: Limit the Sources of Information

The difference between telling the system to analyze this document and telling it to analyze this topic is the difference between a useful work product and a liability. For professionals who work with collective bargaining agreements, investigation files, and classification specifications, this is the single most consequential behavioral change in the session.

Three source-boundary practices carry the weight:

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

A source-bounded analysis is reviewable. A topic-based analysis is not.

Verification and Discoverable Records

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. If you cannot verify it, do not use it.

The records question is a separate concern, and it does not have a uniform answer. Open records obligations, discovery rules, and draft-document exemptions vary by state and by type of government entity. Talk to your counsel before you need to. Questions worth raising in your jurisdiction:

- Are AI-generated analyses created in the course of official duties public records?
- Could the prompts themselves be considered records?

- If a prompt asks the AI to evaluate whether discipline was appropriate, could that 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 discipline from Practice 4 is not only good practice. It is records management. Constraining the analytical scope of what you produce limits the scope of what becomes part of the record.

Part One Closing

Five practices. One professional standard.

- 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, not an optional step. Part Two applies these practices to the work you do every day.

BREAK | 11:00 to 11:15 a.m.

Part Two: Applied Workflows and Iteration

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

Part One answered what to put into a prompt. Part Two answers how to sustain that discipline across a complex task that spans multiple prompts, and how to refine an output once the system returns a draft.

Context Management

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 manage.

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 for LR and HR Professionals

Most major platforms (ChatGPT, Claude, and others) allow a persistent set of instructions that applies to every conversation. Set once. The system treats this as standing context for every task. An example tailored to a labor relations officer:

CUSTOM INSTRUCTIONS 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.
```

Persistent Workspaces

Platforms like ChatGPT and Claude also allow persistent document storage and per-project instructions (Projects, workspaces). For recurring work such as a specific contract, a specific classification study, or a standing committee, a dedicated workspace preserves the documents and the ground rules so you do not rebuild context every time.

Live Demonstrations

Three applied demonstrations follow, each connecting a prompting technique from Part One to a specific type of LR or HR task. The prompts below are the ones used in the session. They are structured as reusable templates.

Demonstration 1: Classification Analysis

Practice 6: Separate extraction from evaluation.

The two-step pattern keeps the extraction of criteria separate from the evaluation of a position against those criteria. The separation creates a reviewable analytical trail that a classification appeal board can follow, and it surfaces gaps in the underlying documentation rather than letting the system fill them in.

STEP 1 PROMPT | EXTRACT THE CRITERIA

```
Extract every minimum qualification and duty criterion from the attached classification specification. 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 language in quotation marks.
```

STEP 2 PROMPT | COMPARE AGAINST THE POSITION

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.

Demonstration 2: Turnover Analysis

Practice 7: Label what you know, what you infer, and what you need.

The three-section output structure distinguishes what the data shows from what it suggests. For elected officials and senior administrators who may act on the findings, that distinction is a credibility requirement, not a formatting preference.

PROMPT | THREE-SECTION OUTPUT

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

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

Demonstration 3: Collective Bargaining Agreement Staged Analysis

Practice 8: Stage complex analysis.

Staging prevents the system from compressing a complex contract into a misleading overview. Each step is reviewable before the next one begins. The three-step pattern below works for any structured document (CBA, personnel rules, administrative code) where you need a targeted provision pulled from a larger body of text.

STEP 1 PROMPT | MAP THE STRUCTURE

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.

STEP 2 PROMPT | SUMMARIZE A TARGETED PROVISION

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

STEP 3 PROMPT | APPLY TO THE QUESTION

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.

Iteration: 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 names specific changes, directs structural reorganization, and requires evidentiary standards. If you would not accept that feedback from a supervisor, do not give it to the system.

Directive Feedback: Three Examples

Each of these is a real iteration prompt. Copy them and adapt the specifics to your own output.

EXAMPLE 1 | TIGHTENING A BOARD SUMMARY

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 | SEPARATING FACTS FROM INFERENCE

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 | CONDENSING FOR AN APPEAL PACKET

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

Ask the system to switch roles and review its own output as an adversary. The model evaluates its own work from a hostile perspective and surfaces weaknesses the original prompt did not anticipate. Personas worth assigning include a skeptical arbitrator, a union advocate, a hostile auditor, or opposing counsel.

RED TEAM PROMPT | ARBITRATION PREPARATION EXAMPLE

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

Restart a conversation when any of the following is true:

- 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.

The restart itself is straightforward. Ask the system to 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.

Closing: Three Questions for Monday Morning

The session ends with three questions designed to move you from workshop attendance to independent practice. Carry them home.

1. What will I try first?

Identify one low-risk task you will try this week using one of the prompts from this guide with your own documents. A gap analysis against a disciplinary file, a three-section turnover summary, a staged CBA review. Pick one.

2. What do I need to clarify?

What tool approval, records question, or policy gap must be resolved in your organization before you can proceed responsibly? Identify the person who can answer it.

3. What conversation do I need to have?

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

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

patrick@boldpathconsulting.com | (920) 522-2413

www.boldpathconsulting.com