

A PRACTITIONER'S PROGRAM

HR AI Literacy Curriculum

A 4-module learning program for HR professionals at all experience levels. Designed for facilitated delivery over 4-6 weeks, or adapted for self-paced completion.

4

MODULES

8

HOURS TOTAL

All

PEOPLE TEAM MEMBERS

Program at a glance

What this curriculum delivers and how it is structured.

TOTAL TIME

~8 hours

Across all modules

FORMAT

Workshop + lab

Async reading + hands-on practice

AUDIENCE

All People Team

HRBPs, HR Ops, HR Tech, leadership

PREREQUISITE

None

Built for HR professionals at every experience level

The four modules

Each module builds on the previous. Together they take an HR professional from zero understanding to drafting their own use case spec.

01	How AI works <i>What HR professionals actually need to know</i>	90 min <i>Workshop</i>
02	Using AI in your daily HR work <i>Prompts, labs, and the verify-before-use discipline</i>	2 hours <i>Workshop + lab</i>
03	Governance, ethics, responsible use <i>Legal risk, scenarios, risk assessment</i>	90 min <i>Workshop</i>
04	Building and scaling AI use cases <i>Prioritization, design, deployment, scale</i>	2 hours <i>Workshop + project</i>

M O D U L E

01

How AI works

What HR professionals actually need to know

90 minutes

Workshop format

Foundational concepts: how LLMs work, hallucination, the rule-based to agentic spectrum, and why prompt quality drives output quality.

Learning objectives

After this session, participants will be able to:

1

Explain the difference between traditional automation and generative AI.

2

Describe why AI "hallucinates" and how to mitigate it in HR contexts.

3

Identify which HR tasks are good candidates for AI assistance and which are not.

Session outline

90 minutes

0:00–0:10	Opening	What do you already believe about AI? Quick poll + discussion.
0:10–0:30	Concept	How LLMs work — without the jargon. Facilitator presentation.
0:30–0:45	Concept	Hallucination, confidence, and why AI makes things up.
0:45–1:05	Exercise	Spot the hallucination — identify errors in AI-generated HR content.
1:05–1:20	Discussion	Where should HR use AI? Where should we be cautious?
1:20–1:30	Wrap-up	Homework: try one AI tool before next session.

Three concepts every HR professional needs

What LLMs actually are

Large language models predict the next most likely word given context. They don't "know" things the way a database knows things — they pattern-match against training data. This is why they can sound authoritative while being wrong.

The rule-based to agentic spectrum

Most HR AI today sits in the middle. Rule-based automation → AI Copilot → AI Agent. True agentic AI, where the model takes multi-step actions autonomously, is emerging and brings new governance requirements.

Prompt quality drives output

The single highest-leverage skill for HR professionals using AI is prompt construction. Vague prompts produce generic output. Specific, well-structured prompts with good context produce usable output.

M O D U L E

02

Using AI in your daily HR work

Prompts, labs, and the verify-before-use discipline

2 hours

Workshop + hands-on lab

Prompt engineering for HR. Live demos. Each participant runs two prompts from the library on their own work and iterates.

Learning objectives

After this session, participants will be able to:

1 Write effective prompts for HR use cases using the four-elements framework.

2 Use the prompt library for at least three HR workflows.

3 Identify when to use AI, when to verify AI output, and when not to use it.

Session outline

2 hours including lab

0:00–0:15	Share-back	What did you try from Module 1 homework? What happened?
0:15–0:45	Concept	Prompt engineering for HR: the four elements of a good prompt.
0:45–1:10	Demo	Live demo: onboarding plan, performance review draft, policy Q&A.
1:10–1:50	Lab	Hands-on: each participant runs 2 prompts from the library on real work.
1:50–2:00	Debrief	What surprised you? What didn't work?

The four elements of an effective HR prompt

The single highest-leverage skill HR professionals can build.

1**ROLE**

Tell the AI who it is. "You are an experienced HRBP..." produces better output than just asking a question.

2**CONTEXT**

Provide the specific details the AI needs. Role, tenure, team size, situation. Don't make the AI guess.

3**TASK**

State clearly what you want it to produce. One output per prompt. "Draft a..." is clearer than "Help me with..."

4**CONSTRAINTS**

Define length, tone, format, and what to avoid. "300 words, no bullet points, avoid generic phrases" significantly improves output.

The verify-before-use rule

A heuristic introduced in Module 2 and reinforced throughout the program.

POLICY DETAILS

Always verify against source system before sharing.

LEGAL REQUIREMENTS

Always verify with Legal before relying on.

FACTUAL CLAIMS

Spot-check against a known source.

tone and content

Always read before sending. AI doesn't know your audience.

M O D U L E

03

Governance, ethics, responsible use

Legal risk, scenarios, and the risk assessment template

90 minutes

Workshop format

*The session that generates the most discussion and the most resistance. Budget flex time.
Keep an AI skeptic on the facilitation team.*

Learning objectives

After this session, participants will be able to:

1

Apply the organization's AI use policy to real HR scenarios.

2

Identify the key legal risk areas in HR AI: bias, privacy, employment law.

3

Complete a basic risk assessment for a new use case.

Key legal risk areas

Regulatory references current as of May 2026. Verify current law before relying on specific references.

Bias and disparate impact

AI trained on historical HR data can encode and amplify past discrimination. Annual adverse impact testing is not optional — it is legally and ethically required.

State and local AI hiring laws

Illinois AI Video Interview Act, NYC Local Law 144, Colorado AI Act, and emerging legislation. Several jurisdictions now require disclosure, impact assessments, or consent.

GDPR and CCPA

Employees in covered jurisdictions have rights regarding automated decision-making. Know which employees are covered and how use cases interact with their rights.

The "AI as a manager" trap

AI that monitors, evaluates, or scores employees without their knowledge — even for legitimate purposes — creates legal and trust exposure. Transparency is protective.

Scenario practice

Groups of 3-4 discuss: Is this use appropriate? What are the risks? What would you change?



SCENARIO

An AI tool automatically rejects resumes with gaps of more than 6 months without showing them to a recruiter.



SCENARIO

An HRBP uses an AI writing assistant to draft a PIP for an employee, then reviews and edits it before sharing with the manager.



SCENARIO

People Analytics builds an attrition risk model that scores all employees monthly and shares the scores directly with managers.



SCENARIO

An AI chatbot answers benefits questions for employees. It sometimes gives incorrect information but flags that users should confirm with HR.

M O D U L E

04

Building and scaling AI use cases

From idea to deployment to scale

2 hours

Workshop + project

Participants apply the prioritization matrix to their team's top ideas and draft a one-page use case spec for their highest-ranked use case.

Learning objectives

After this session, participants will be able to:

1 Apply the prioritization matrix to identify your team's highest-value AI use cases.

2 Define a use case with clear success metrics and governance requirements.

3 Draft a basic implementation plan for a low-complexity use case.

The use case lifecycle

From idea to scale, with iteration baked in.



↺ *iterate*

One-page use case spec template

What participants complete for their team's top-ranked use case.

Use case name	Human review gate
Problem it solves (1 sentence)	Success metric(s)
Who benefits and how	What could go wrong
AI approach (Agent / Copilot / Automation / Analytics)	Governance requirements (from risk assessment)
Data needed	First step to move forward

Participants present specs to HR Technology for the optional "HR AI Champion" designation.

Assessment & recognition

How participants earn recognition from the People Team.

STANDARD COMPLETION

HR AI Practitioner

Awarded to participants who complete all four modules and the hands-on components.

Requirements:

- Attend all 4 modules
- Complete Module 2 lab
- Participate in Module 3 scenarios

OPTIONAL EXTENSION

HR AI Champion

An informal internal designation that supports peer enablement across the People Team.

Additional requirement:

Complete a full use case spec and present it to the HR Technology team.

Facilitator notes

Lessons from delivering this program. Skip these at your peril.

MODULE 3

Generates the most discussion and sometimes the most resistance. Budget flex time.

MODULE 2 LAB

The highest-value session in the program. Do not let it get cut for time.

CASE STUDIES

Adapt Module 3 scenarios to your organization's actual context. Generic examples land less well than ones that feel real.

FACILITATION

Keep an AI skeptic on the facilitation team if possible. The best discussions happen when the room has range.

THIS CURRICULUM IS PART OF

HR AI Transformation Playbook

A practitioner's field guide for HR and People Teams building responsible AI programs. The full playbook also includes use cases, prompt libraries, governance templates, agent design patterns, evaluation frameworks, and a transformation roadmap.

REPOSITORY

github.com/ellehelvig/hr-ai-transformation-playbook

CONNECT

linkedin.com/in/elle-helvig