


Applying AI

in Learning & Development



From Platforms
to Performance

Josh Cavalier

© 2026 ASTD DBA the Association for Talent Development (ATD)
All rights reserved. Printed in the United States of America.

29 28 27 26

1 2 3 4 5

No part of this publication may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, information storage and retrieval systems, or other electronic or mechanical methods, without the prior written permission of the publisher, except in the case of brief quotations embodied in critical reviews and certain other noncommercial uses permitted by copyright law. For permission requests, please go to copyright.com, or contact Copyright Clearance Center (CCC), 222 Rosewood Drive, Danvers, MA 01923 (telephone: 978.750.8400; fax: 978.646.8600).

ATD Press is an internationally renowned source of insightful and practical information on talent development, training, and professional development.

ATD Press

1640 King Street
Alexandria, VA 22314 USA

Ordering information: Books published by ATD Press can be purchased by visiting ATD's website at td.org/books or by calling 800.628.2783 or 703.683.8100.

Library of Congress Control Number: 2025938446

ISBN-10: 1-96023-152-9
ISBN-13: 978-1-960231-52-9
e-ISBN: 978-1-96023-153-6

ATD Press Editorial Staff

Director: Sarah Halgas
Manager: Melissa Jones
Content Manager: Alexandria Clapp
Developmental Editor: Jack Harlow
Production Editor: Katy Wiley Stewts
Text Designer: Shirley E.M. Raybuck
Cover Designer: Brandon Rush

Text Layout: Kathleen Dyson

Applying AI in Learning and Development



Figures

1. The Evolution of AI Technology
2. The AI Landscape
3. Multimodal AI Content Generation
4. Multiple Outputs
5. GenAI Model Dimensions
6. The Human–AI Task Scale
7. The Technology Fear Cycle
8. Mapping Your Data Ecosystem
9. Partnering With IT
10. Example Platform Evaluation Rubric
11. L&D as Knowledge Ecosystem Architects
12. The Knowledge Graph Recipe
13. Optimizing Workflow With AI
14. The Evolution From LMS to LXP
15. Integrating an LMS with an LXP
16. Learning Experience Development With GenAI
17. The Role of AI in Content Management
18. The AI-Driven Content Curation Process
19. Content Adaptation With Multimodal AI
20. Example LLM Prompt
21. LLM vs. Google Search

22. Syntactic Parsing vs. Semantic Analysis
23. Smoked Potato Salad Example Prompt
24. Example of Automated Prompt Optimization
25. Resulting Customer-Service Image
26. Prompt Library Example
27. Prompt Metadata Example
28. Prompt Workflow Example
29. Dynamic Learning Interaction
30. Prompt Workflow Example
31. GPTs vs. Automation
32. Make.com Automation Example
33. How Does an Agent Work?
34. Employee Onboarding Workflow
35. Levels 8 Through 10 of the Human–AI Task Scale
36. Sample Rule-Based Chatbot Conversation
37. Sample AI Assistant Conversation

Tables

1. Key Vendor Selection Criteria and Scorecard
2. Metadata Fields for a Prompt Library
3. Automation vs. AI Agent Example
4. Development Frameworks

Exercises

1. Getting Started
2. The AI Landscape and Its Impact
3. How Will GenAI Affect Your Approach to L&D?
4. You as a Human-Machine Performance Analyst
5. GenAI as a Content Enhancer
6. Solving Skills Gaps With AI
7. Leading With AI
8. Your Organization's AI Readiness
9. Where Are You Now? An AI Readiness Assessment
10. Understanding Your Data

11. What Biases Exist in Your AI Systems?
12. How Will Your Expertise Intersect With AI?
13. Ethics From the Start
14. Imagining Task Enablement
15. Task Mapping
16. Build Your AI Readiness Canvas for Performance Support
17. LMS vs. LXP
18. What's Possible With Your Learning Platform?
19. AI-Powered Insights
20. The Opportunities and Challenges of Personalized Learning
21. Assessing Learning Platform Vendors
22. Consider Your Current Content Development Process
23. AI vs. Manual Showdown
24. The "Content Chameleon" Challenge
25. Implementing AI Avatars in Your Learning Strategy
26. LLM Probability
27. Freeform Prompts
28. Role, Task, and Details
29. Freeform to Structured Prompt
30. Creating a Prompt Mini-Library
31. Migrating and Enhancing Your Prompt Library
32. Prompt Work lows
33. Creating a GPT
34. Planning Your Automation
35. Mapping an Existing Work low to an AI Agent
36. Getting Started: Your 30-Day Agent Work low Action Plan
37. Craft Your Own Mini AI Assistant
38. Fantasy Role-Playing Game
39. Building an Onboarding FAQ Chatbot

Glossary

References

Figure 1. The Evolution of AI Technology

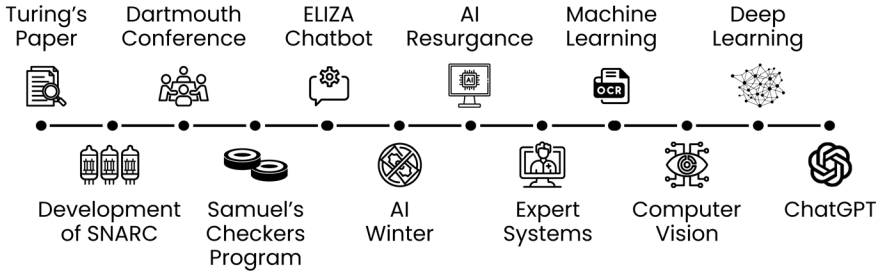


Figure 2. The AI Landscape

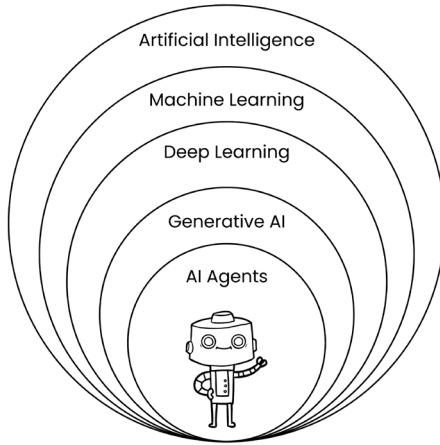


Figure 3. Multimodal AI Content Generation




Figure 4. Multiple Outputs




Figure 5. GenAI Model Dimensions

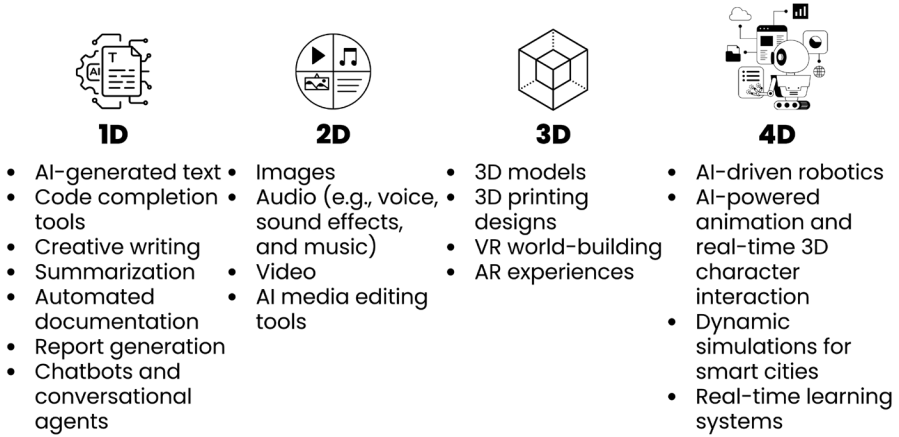


Figure 6. The Human-AI Task Scale

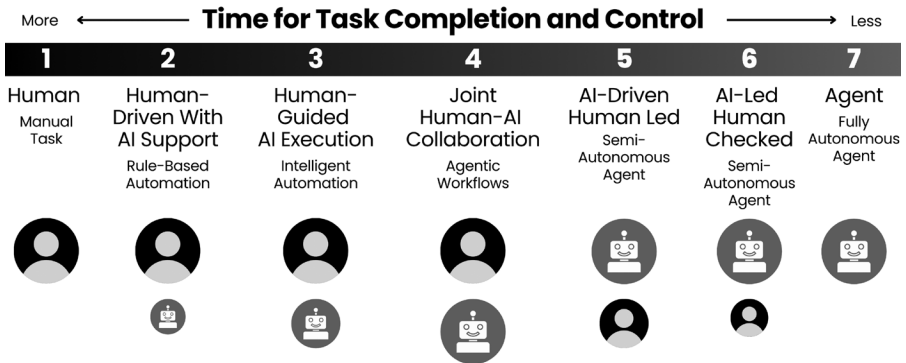


Figure 7. The Technology Fear Cycle

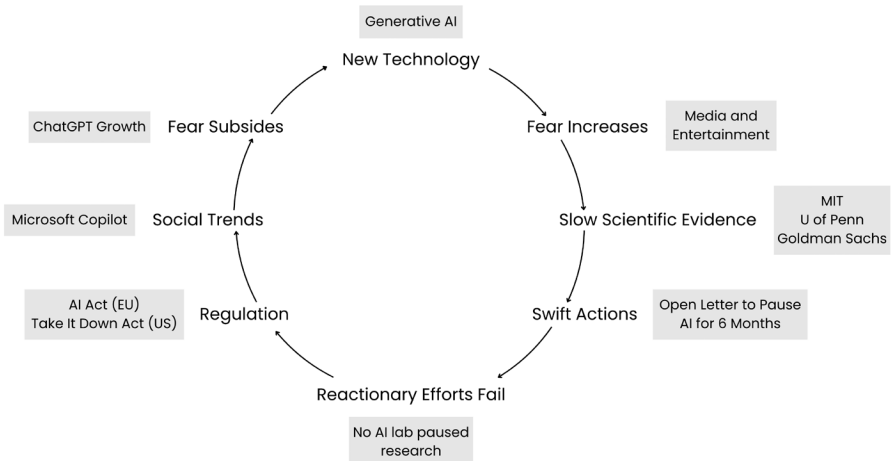


Figure 8. Mapping Your Data Ecosystem

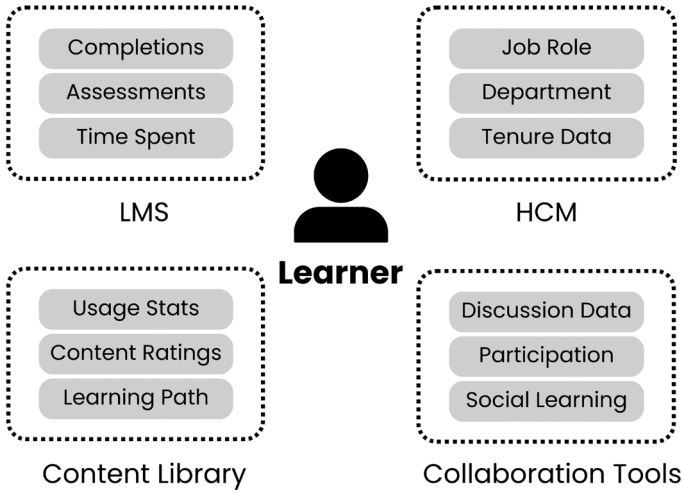


Figure 9. Partnering With IT

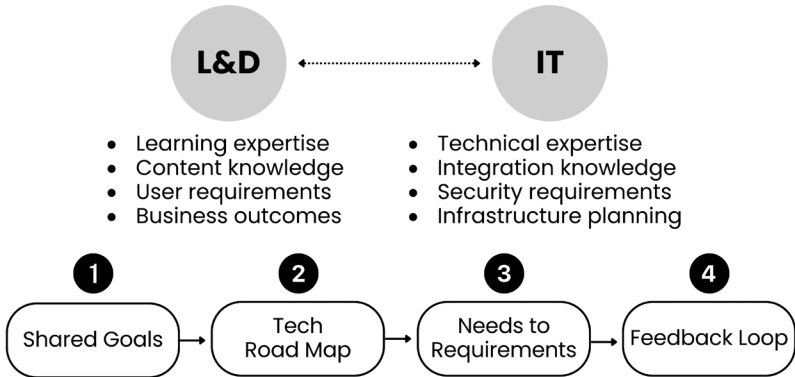


Figure 10. Example Platform Evaluation Rubric

Platform Evaluation Framework for Learning Excellence			
Evaluation Criteria	Weight	Platform A	Platform B
Learning Goals Alignment Skills statement integration and outcome measurement	25%	4.5/5	3/5
Content Management Versioning, metadata support, and content reusability	20%	3/5	4.5/5
Learner Experience Personalization, mobile support, and accessibility	20%	4/5	4.5/5
Data and Analytics Advanced reporting, actionable insights, and customizable dashboards	15%	4/5	3/5
Future Readiness AI capabilities, system integrations, and scalability for growth	20%	3.5/5	4.5/5
Weighted Score		3.8/5	3.9/5

Framework Implementation Process	
1 Anchor in Learning Objectives Center on organizational impact.	2 Balance Operational and Strategic Needs
3 Implement Weighted Scoring Align with business priorities.	4 Collaborate With Stakeholders Include all key perspectives.

Figure 11. L&D as Knowledge Ecosystem Architects

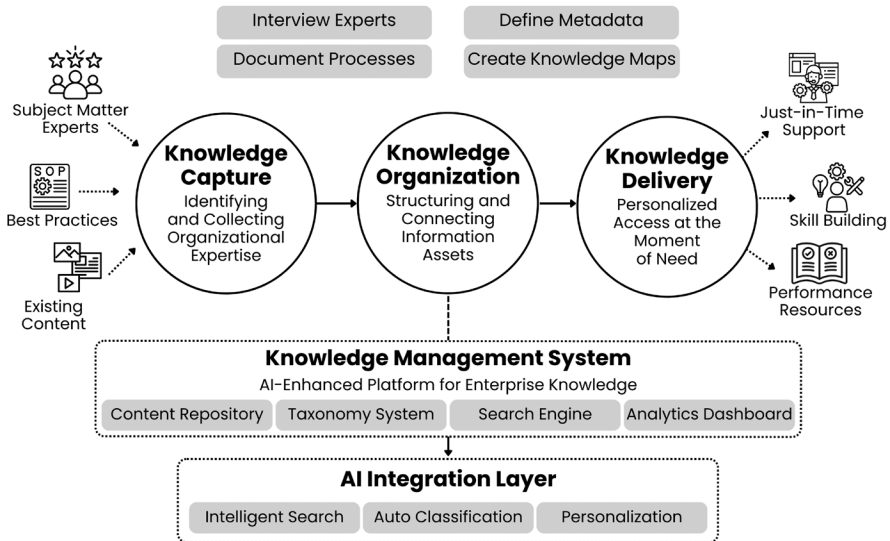


Figure 12. The Knowledge Graph Recipe

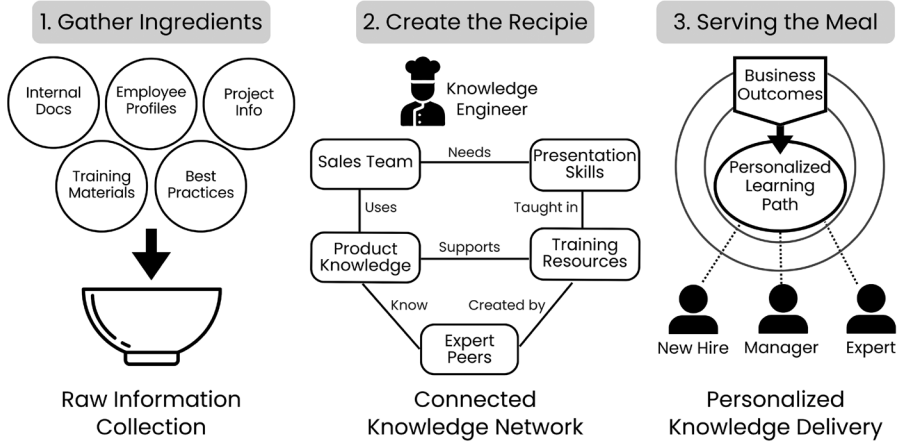


Figure 13. Optimizing Workflow With AI




Figure 14. The Evolution From LMS to LXP




Figure 15. Integrating an LMS with an LXP




Figure 16. Learning Experience Development With GenAI

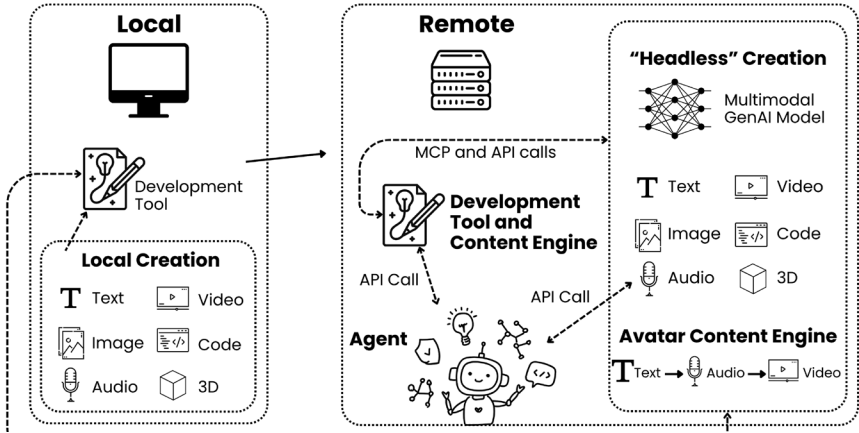


Figure 17. The Role of AI in Content Management

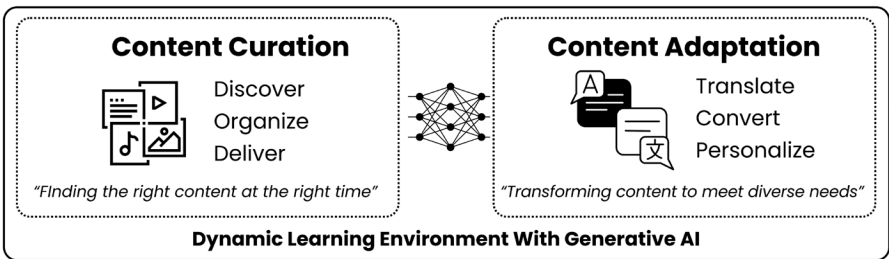


Figure 18. The AI-Driven Content Curation Process

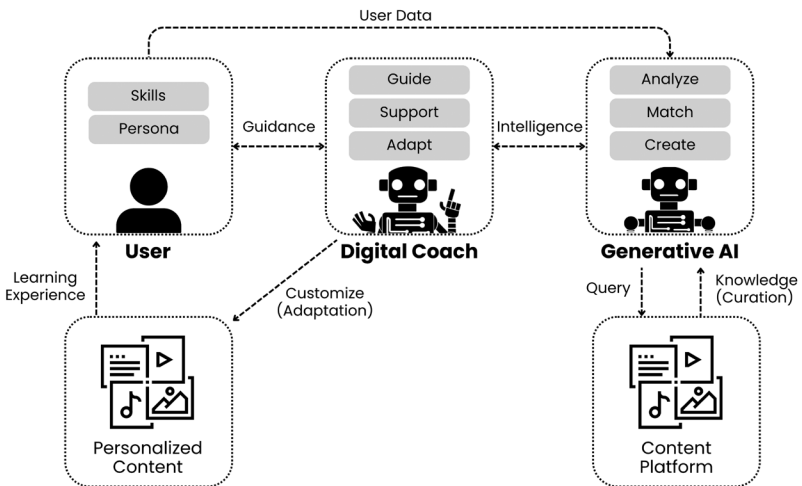


Figure 19. Content Adaptation With Multimodal AI

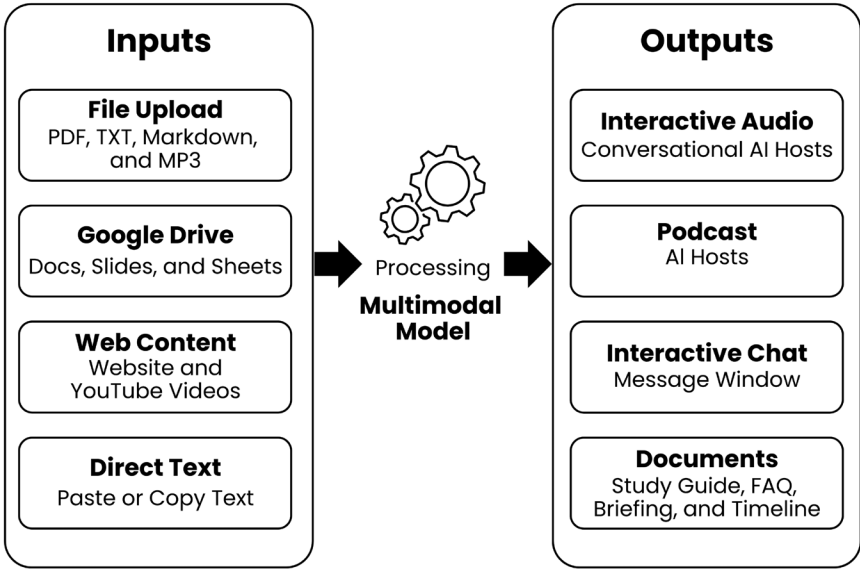


Figure 20. Example LLM Prompt

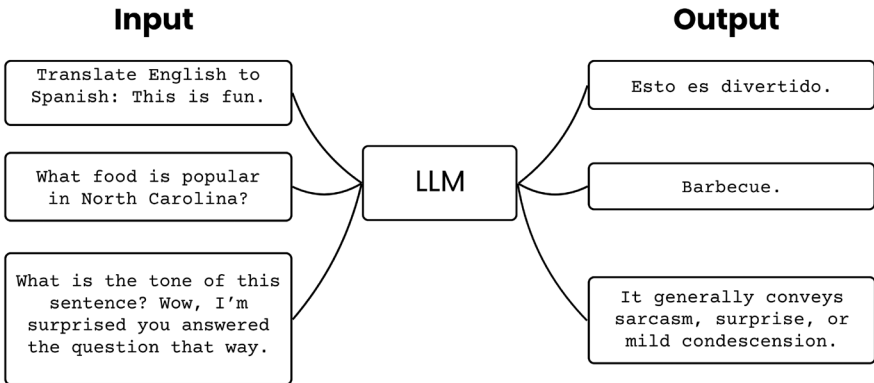


Figure 21. LLM vs. Google Search

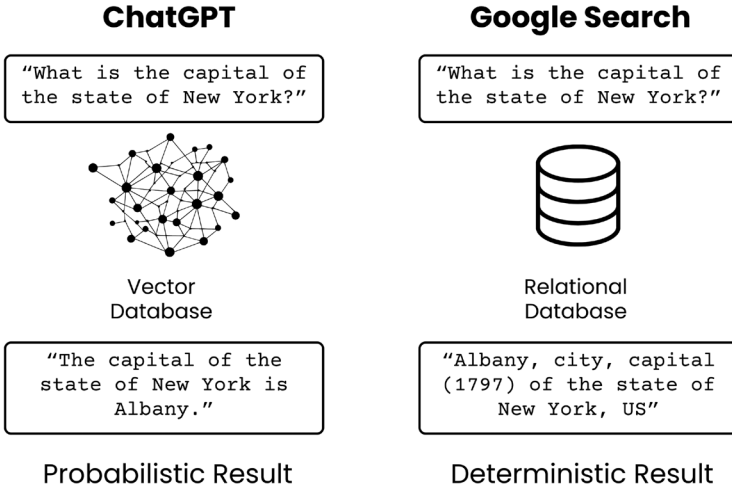


Figure 22. Syntactic Parsing vs. Semantic Analysis

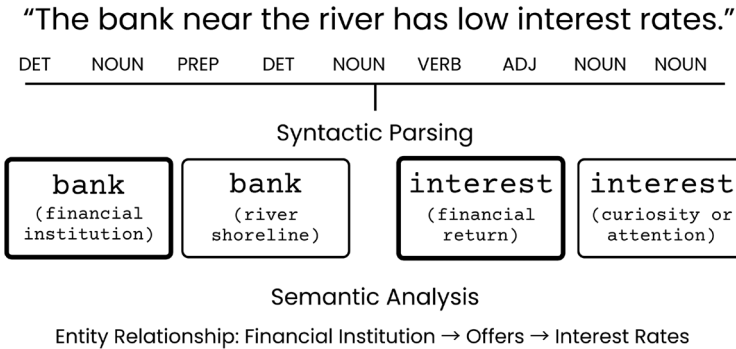


Figure 23. Smoked Potato Salad Example Prompt

1	Enter the Prompt	How do I cook a smoked potato salad?
2	Text Preprocessing	how do i cook a smoked potato salad
3	Tokenization	["how," "do," "i," "cook," "a," "smoked," "potato," "salad"]
4	Encoding (Embedding and Positional Encoding)	[[0.25, 0.73, 0.11], [0.81, 0.64, 0.95], ...]
5	Forward Pass and Self-Attention	"how" → Low "cook" → High "smoked" → Medium "potato salad" → High
6	Decoding and Text Generation	"To make smoked potato salad, first smoke the potatoes, and then mix them with mayonnaise and seasonings."
7	Postprocessing	"To make a delicious smoked potato salad, first smoke the potatoes, and then mix them with mayonnaise, mustard, and seasoning for flavor."

Figure 24. Example of Automated Prompt Optimization


Initial Image Prompt	Automated Prompt Optimization	Generated Image
A delicious smoked potato salad.	A delicious smoked potato salad served in a rustic bowl. The dish features golden-brown smoked potatoes mixed with creamy dressing, fresh herbs, and crispy bacon bits. Garnished with paprika and ground black pepper. The background is warm and inviting, with a wooden table and soft natural lighting.	

Figure 25. Resulting Customer-Service Image



Figure 26. Prompt Library Example

250+ Prompts for Learning & Develo...

- Instructional Designer 41
 - Learning Objectives May 19, 2023 4:57 AM
 - Microlearning Strategy May 19, 2023 4:57 AM
 - Motivation and Encouragement May 21, 2023 11:20 AM
 - On-Demand Support Guide May 21, 2023 11:20 AM
 - Patterns in Learner Responses May 19, 2023 7:49 AM
 - Proficiency Levels of Learners May 21, 2023 7:49 AM
 - Real-Time Feedback May 19, 2023 7:13 PM
 - Real-World Scenarios & Case Studies May 15, 2023 7:13 PM
 - Research Article Key Takeaways May 20, 2023 8:42 PM
 - Simple Multiple Choice Questions February 21, 2023 8:42 PM
 - Multiple Choice Questions June 6, 2024 4:25 PM
 - Step-by-Step Tutorial May 21, 2023 11:20 AM
 - Targeted Interventions May 21, 2023 7:49 AM
 - Terminology Suggestions May 21, 2023 5:56 AM
 - Time Management and Organization May 21, 2023 11:20 AM
 - Topic Exploration May 21, 2023 11:20 AM
 - Topic Resources May 21, 2023 11:20 AM
 - Quizzes and Assessments May 19, 2023 4:57 AM
 - Webinar Transcript Summary May 20, 2023 2:39 PM
- Learning Architect 45
 - 2-Day Workshop Agenda February 14, 2023 10:54 PM
 - Additional Reading Materials May 21, 2023 11:20 AM
 - Analogies for a Topic May 21, 2023 11:20 AM
 - Best Practices and Industry Standards May 18, 2023 9:47 PM
 - Blended Learning Solution May 18, 2023 9:47 PM
 - Categorize Learner Resources May 21, 2023 8:42 AM
 - Classify Instructional Content May 21, 2023 8:37 AM

Multiple Choice Questions

Created June 6, 2024 4:25 PM

Tested NLP Engine ChatGPT

Role or Task Instructional Designer

Copyright ©2024 by Josh Cavalier. Licensed under Attribution-ShareAlike 4.0 International.

Function Text Generation

Last edited time June 6, 2024 4:28 PM

URL creativewcommons.org/licenses/by-sa/4.0/

Description

This prompt generates simple multiple choice questions. Change the text, "[ENTER LEARNING OBJECTIVES HERE]", with your learning objectives.

Prompt

```
Plain Text
{
  "Task": "Multiple Choice Question",
  "Task_Rules": [
    "The question is called question stem",
    "All answers are called options",
    "Wrong answers are called distractors",
    "Correct answer is called correct",
    "Avoid double negatives in the question stem",
    "Make all distractors plausible",
    "Place most of the words in the question stem",
    "Text comprehension and critical thinking and recall",
    "Use simple sentence structure and precise wording for the question stem",
    "Mix up the order of the correct answers",
    "Keep the number of options to 3",
    "No trick questions",
    "Avoid using 'All the Above' and 'None of the Above' as options"
  ]
  "Role": "Instructional Designer",
}
```

Figure 27. Prompt Metadata Example

Multiple Choice Questions

Created	June 6, 2024 4:25 PM
Tested NLP Engine	ChatGPT
Role or Task	Instructional Designer
Copyright	©2024 by Josh Cavalier. Licensed under Attribution-ShareAlike 4.0 International.
Function	Text Generation
Last edited time	June 6, 2024 4:28 PM
URL	creativecommons.org/lic...a/4.0/

Figure 29. Dynamic Learning Interaction

Video Prompt:

An avatar explains the importance of food safety. Cut in footage and images of a kitchen and restaurant patrons.

Text Prompt:

Describe the importance of food safety at a restaurant in one paragraph.

Question
Qui dolorem eum fugiat quo voluptas nulla pariatur?

- Dolorem eum
- Dolorem eum
- Dolorem eum

Answer
Sed ut perspiciatis unde omnis iste natus error sit voluptatem accusantium doloremque laudantium, totam rem aperiam.

Question Prompt:

Based on the video transcript, create a multiple choice question with three options. Display the explanation of the correct answer.

Figure 28. Prompt Workflow Example

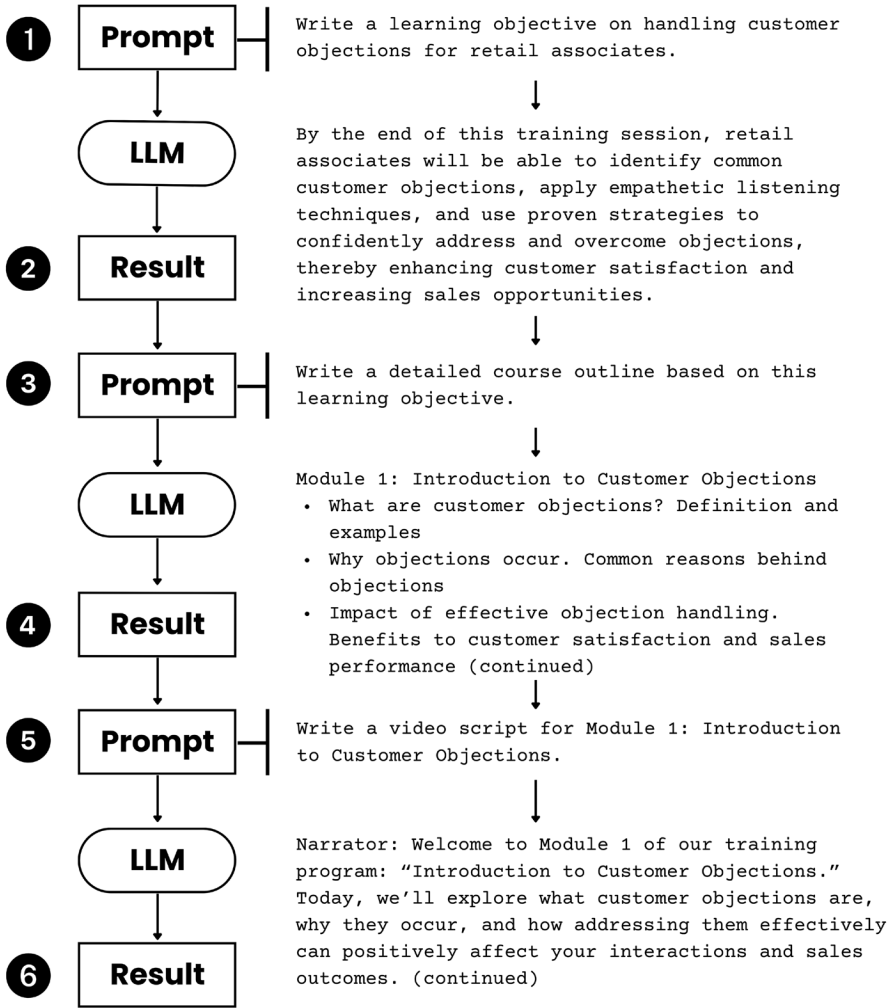


Figure 30. Prompt Workflow Example

Weekly YouTube Publishing Workflow

Task Name	Tools	Prompt	AI Model
YouTube Title	ChatGPT	<p>You are a YouTuber with a large following. Review the following YouTube video title formats.</p> <p>How-To Guides: "How to [Achieve Something] Step by Step" Example: "How to Master Video Editing Step by Step" Listicles: "Top [Number] [Topic] You Must Know" Example: "Top 10 Marketing Strategies for Success" Tutorials: "[Topic] Tutorial for Beginners" Example: "Photography Tutorial for Beginners" Reviews: "Is [Product/Topic] Worth It? Honest Review" Example: "Is the iPhone 13 Worth It? Honest Review" Behind the Scenes: "A Day in the Life of a [Your Profession]" Example: "A Day in the Life of a YouTuber" Challenges: "[Challenge Name]: Can I [Outcome]?" Example: "30-Day Fitness Challenge: Can I Get Fit?" Interviews: "Interview with [Influential Figure] on [Topic]" Example: "Interview with Elon Musk on Future Tech" Explainers: "What Is [Complex Topic]? Explained Simply" Example: "What Is Cryptocurrency? Explained Simply" Vlogs: "[Event/Activity] Vlog: [Exciting Description]" Example: "Travel Adventure Vlog: Exploring Paradise" FAQs: "Answering Your [Topic] Questions" Example: "Answering Your Cooking Questions" Come up with Video titles on the topic of: Generative AI for Learning & Development</p> <p>The titles should not be more than 100 characters long. The descriptions should not be longer than 1000 characters long. The title and description should have the keywords in them. The title should be good for videos and not be generic titles. The titles should either be teaching the viewer about something in a visual way. The description should be long and contain multiple informational paragraphs about the topic. The description should also ask the viewer to click the subscribe button, like and ring the bell to get notified when a new video is created.</p>	GPT-4.5
YouTube Thumbnail	Canva	Direction: Use established template for thumbnail	
YouTube Description	ChatGPT	(See YouTube Title Prompt)	
YouTube Script	ChatGPT	<p>Topic: [Insert topic]. Write a video script explaining how to [insert topic details] for beginners. Limit video length to 3 minutes. Use a screencast narration style and explainer tone. Include lived experiences from the perspective of [insert persona]. Include common pain points. Include clear step-by-step details. Have a bias for impact value statements with numbers. Highlight pro tips. Add a call to action. Add 3 resources to learn more. Use "Global English" to make content and context accessible for non-native comprehension. Don't use idioms. Be literal and stay away from metaphors and colloquial language. Keep sentences short. Standardize terminology to minimize changes. Avoid directional language. Use inclusive, accessible, person-first language.</p>	GPT-4.5
Record Video (A-Roll)	Camtasia	Direction: Record screen and camera shots	
Stills (B-Roll)	Envato Elements, MidJourney, Firefly	An ultra-realistic photograph captured with a Sony a7R camera, equipped with an 85mm lens at F 1.2 aperture setting, portraying [Topic]. The background is beautifully blurred, highlighting the subject. [Describe scene in detail] The image, shot in high resolution and a 16:9 aspect ratio, captures the subject's natural beauty and personality with stunning realism -ar 16:9 --v 5.2 --style raw	Midjourney 6.1
Clips (B-Roll)	Envato Elements, RunwayML		Runway ML
Audio Effects	Envato Elements		
Music	Suno.ai	Lofi Chill Hip Hop	Suno.ai
Voice Over	Audition, Eleven Labs		

Figure 31. GPTs vs. Automation

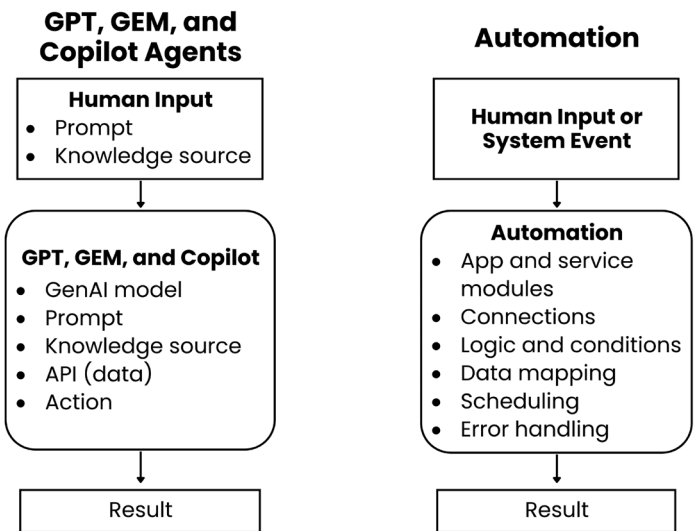


Figure 32. Make.com Automation Example




Figure 33. How Does an Agent Work?

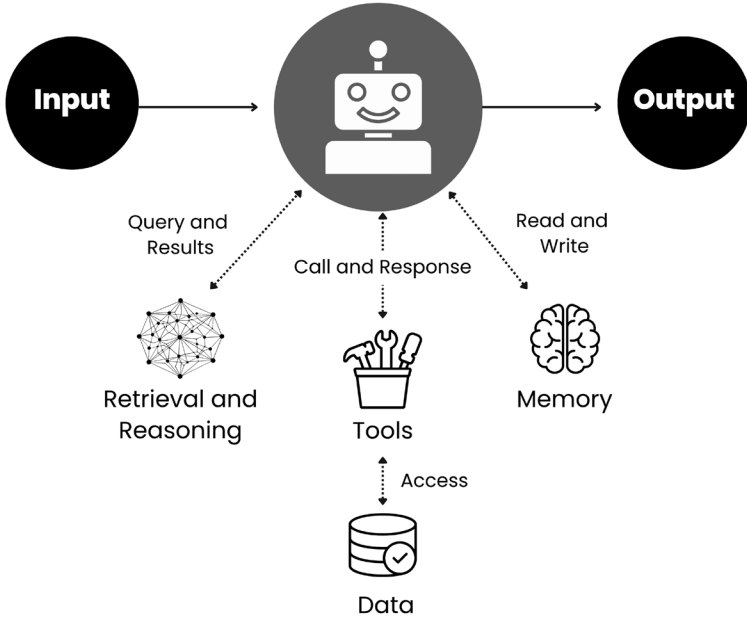


Figure 34. Employee Onboarding Workflow

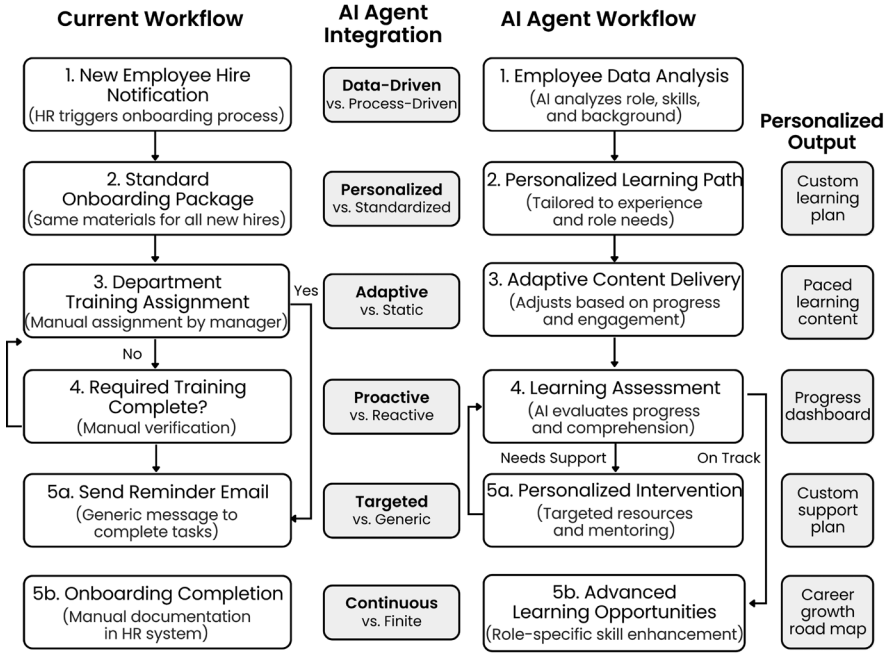


Figure 35. Levels 8 Through 10 of the Human-AI Task Scale

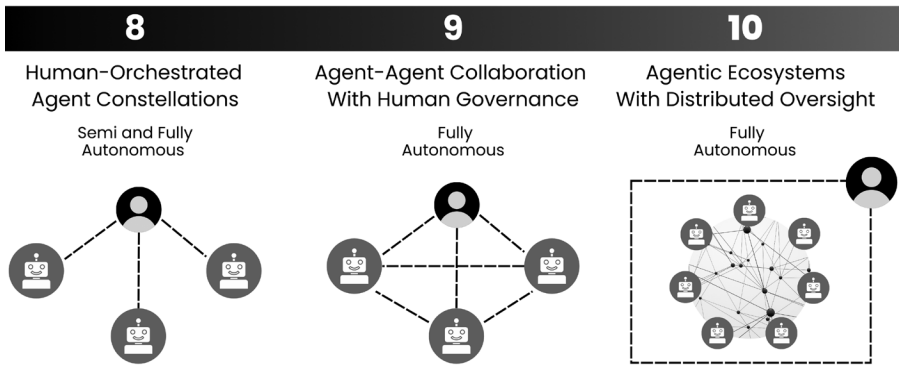


Figure 36. Sample Rule-Based Chatbot Conversation

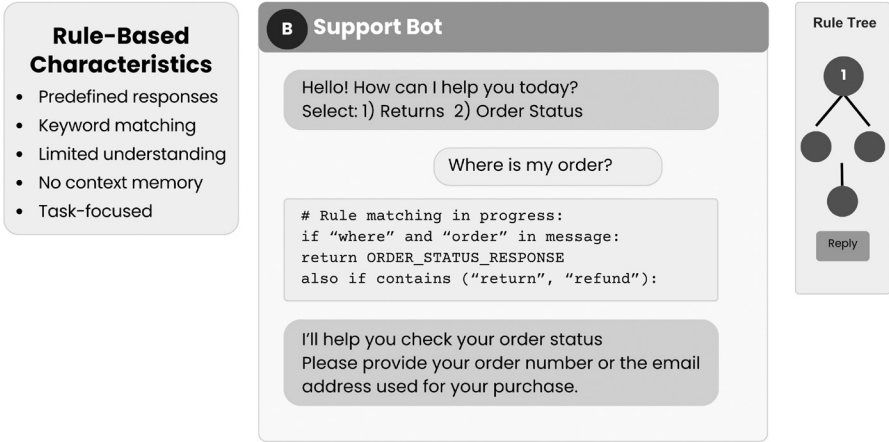


Figure 37. Sample AI Assistant Conversation

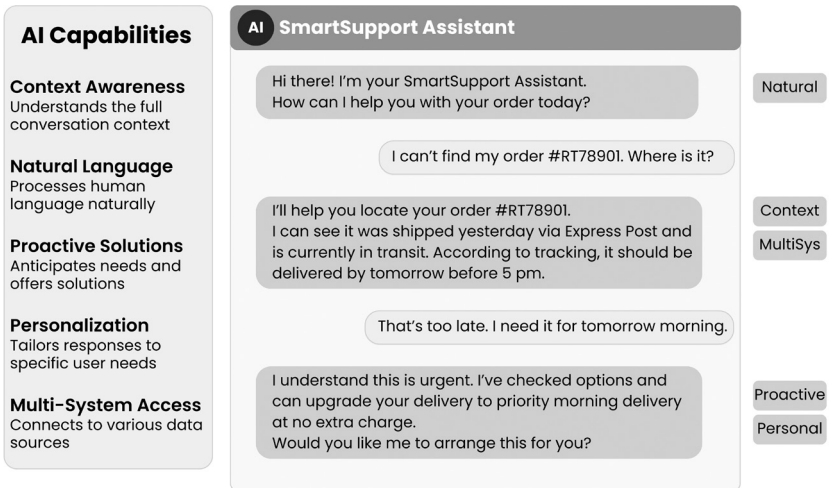


Table 1. Key Vendor Selection Criteria and Scorecard

Criteria	Why it Matters	What to Look For	Questions to Ask	Score (1–5)
AI capabilities	Determine how well the system personalizes learning and provides insights. You want mature, effective AI that aligns with your needs (e.g., recommendations and skills analysis).	<ul style="list-style-type: none"> A list of AI features (e.g., recommendations, adaptive learning, analytics, or chatbots) Demos and success metrics Ensure the AI can handle your content volume and diversity 	<ul style="list-style-type: none"> What AI-driven features does your system have, and how do they improve learning outcomes? Could we see a demo of the recommendation engine adapting to a user over time? 	
Integration and compatibility	Integration is essential for a seamless learning ecosystem. A system that can't integrate will create silos and extra work.	<ul style="list-style-type: none"> Support for APIs (xAPI, LTI, and SCORM standards) Prebuilt connectors (e.g., to your LMS or HRIS). Client examples of integration scenarios similar to your case 	<ul style="list-style-type: none"> Can your platform integrate with our LMS and HR systems? Can you show an example or reference of this integration in action? 	
User experience (learner and admin)	High adoption and engagement depend on an intuitive, pleasant UX. Admins also need a good experience for managing content and interpreting data.	<ul style="list-style-type: none"> A modern, easy-to-navigate interface Mobile support Personalized dashboards for learners Clear dashboards and simple content management for admins Determine if the vendor allows some branding customization 	<ul style="list-style-type: none"> How do you help drive user adoption? Do you have resources or playbooks for change management in L&D for your system? 	

Table 1. (continued)

Criteria	Why It Matters	What to Look For	Questions to Ask	Score (1–5)
<p>Content support and curation</p>	<p>The platform should enrich learning with diverse content and make curation easier.</p>	<ul style="list-style-type: none"> • Ability to support various content types (e.g., videos, PDFs, and courses) • Integration with content libraries (e.g., LinkedIn Learning) • AI curation of external content, and tools for SMEs to contribute content easily 	<ul style="list-style-type: none"> • How much can we configure the skills frameworks or recommendation rules? • Can the AI be tailored to our competency model? 	
<p>Analytics and reporting</p>	<p>Demonstrate the value of L&D and help refine programs with rich analytics that ensure you can track progress and ROI.</p>	<ul style="list-style-type: none"> • Visual dashboards (for skills gaps and engagement) • Exportable reports • Predictive analytics (e.g., who might need training) • Ability to set and track KPIs (e.g., learning completion vs. performance) 	<ul style="list-style-type: none"> • What analytics do you provide on the backend? • Can we create custom reports or dashboards for our leadership? 	
<p>Vendor expertise and support</p>	<p>A knowledgeable vendor will guide you in how to best leverage the system. Good support ensures issues don't derail usage.</p>	<ul style="list-style-type: none"> • Strong implementation support (including training and migration help) • Ongoing customer success check-ins • L&D expertise (do they consult on learning strategy or just tech?) • References praising support responsiveness 	<ul style="list-style-type: none"> • What support do you provide during and after implementation? • What new AI features are you planning in the next year? 	

Table 1. (continued)

Criteria	Why It Matters	What to Look For	Questions to Ask	Score (1–5)
Security and compliance	The system must protect data and comply with laws, especially if you're working in a regulated industry or region. Security policies can also build trust with IT and employees.	<ul style="list-style-type: none"> • Data encryption • Region-specific data handling • Certifications (such as ISO 27001 and SOC 2) • Features like role-based access control (e.g., managers only seeing their team's data) • The vendor's privacy policy regarding user data and AI (it should include no misuse of personal info) 	<ul style="list-style-type: none"> • How does your platform ensure data privacy and compliance with regulations like GDPR or SOC 2? • Can you share documentation of your certifications and how you handle user data in AI models? 	
Scalability and performance	The system should accommodate your entire target audience without performance issues, now and as you grow.	<ul style="list-style-type: none"> • Cloud infrastructure robustness (e.g., uptime guarantees and load testing results) • References from companies of similar or larger size that are using it • Multilanguage or localization support (if you operate globally) 	<ul style="list-style-type: none"> • Can your platform support our expected user volume as we grow? • What are your uptime guarantees and load testing results for clients of similar scale? 	
Cost and flexibility	The platform needs to fit your budget and show value. Flexible licensing can help you start small and grow.	<ul style="list-style-type: none"> • A transparent pricing model (e.g., per user or tiered) • What's included and what's extra (e.g., some vendors charge extra for advanced AI modules) • The ability to test or do a phased rollout licensing • Estimated implementation cost and time • ROI case studies or tools to help calculate impact 	<ul style="list-style-type: none"> • What results have your clients seen after implementing your AI features (e.g., engagement uplift or reduction in skills gaps)? 	

Table 2. Metadata Fields for a Prompt Library

Metadata Field	Description	Example
Prompt	The actual text of the prompt	Write a short story about great customer service.
Category or tag	Keywords or categories that help organize and classify the prompt	<ul style="list-style-type: none"> • Creative writing • Storytelling • Customer service
Version	The version number of the prompt, allowing for tracking changes over time	v1.2
Date created	The date and time the prompt was created	2025-01-05 12:00:00
Author	The creator of the prompt	John Doe
Function	The intended purpose or function of the prompt (e.g., text generation or summarization)	Text generation
Tested model	Which generative AI models the prompt was tested on	ChatGPT 5, Gemini 3, and Claude 4
Keywords	Relevant keywords that help with searching and filtering prompts	<ul style="list-style-type: none"> • Customer • Service • Short story
Title	A name for the prompt	“The Customer Service Scenario”
Description	A detailed account of the prompt	This prompt is designed to generate creative learning scenarios.

Table 3. Automation vs. AI Agent Example

Factor	Automation	AI Agent
Trigger	New employee joins (event-based)	Skill mastery goal (objective driven)
Action	Executes predefined onboarding steps	Chooses personalized learning pathways
Decision making	Conditional logic; limited flexibility	Adaptive reasoning; highly flexible
Learning and improvement	Does not learn or adapt independently	Learns, adapts, and improves continuously

Table 4. Development Frameworks

Framework	Description	Code Style
n8n AI agent builder	Open-source visual builder with LangChain integration	No code
Vertex AI agent builder	Enterprise conversational AI with Gemini integration	No code
CrewAI	Role-based system with a hybrid interface	Low code or code first
LangGraph	Stateful workflow engine for complex systems	Code first
Microsoft AutoGen	Distributed agent framework with benchmarking	Code first

EXERCISE 1

Getting Started

Ready to dip your toes into the AI content creation pool? Here’s how to get started:

1. **Choose your AI with care.** Start with a GenAI tool that gives you secure, private access—ideally one that is built for professionals, trusted with enterprise-level data, and can handle a range of content formats (including text, images, audio, and video).
2. **Master the art of prompting.** Your output’s quality depends heavily on your input’s quality. Learn to write clear, detailed prompts that specify not just the topic, but the tone, style, and learning approach you want. (More on this in chapter 8.)
3. **Start small.** Begin by using AI for parts of your content process—maybe generating outlines or creating knowledge check questions. As you get comfortable, you can expand its role.
4. **Always keep humans in the loop.** Remember, AI is a tool, not a replacement. Your expertise in learning design, the subject matter, understanding your audience, and creativity are still crucial.
5. **Iterate and refine.** Use AI as a starting point, then iterate. Often, the magic happens in the back-and-forth between humans and AI.

EXERCISE 2

The AI Landscape and Its Impact

How familiar are you with the key players and recent breakthroughs in GenAI, and how might this knowledge influence your L&D strategy?

- Which GenAI tools or platforms have you experimented with in your work? What were your impressions?
- How do you stay informed about new developments in AI and their potential applications in L&D?
- What concerns do you have about the rapid advancement of AI in the learning field?
- How might the differences between open and closed AI systems affect your choice of tools for learning projects?

Now, research one GenAI tool you haven't used before. Spend 30 minutes exploring its capabilities and considering how it might be applied in your work. Share your findings with a colleague or in a professional network.

EXERCISE 3

How Will GenAI Affect Your Approach to L&D?

GenAI will fundamentally change L&D's approach to creating and delivering learning experiences. How might it affect you? Reflect on these questions:

- What aspects of your current work in L&D could be enhanced by GenAI?
- How do you envision the role of human expertise evolving alongside AI-generated content?
- What ethical considerations do you need to address when implementing GenAI in your learning programs?
- How might GenAI change the expectations of learners in your organization?

Now, choose one upcoming learning project or course. Spend 15 minutes brainstorming how you could incorporate GenAI into its design, development, or delivery. Consider both the potential benefits and challenges.

EXERCISE 4

You as a Human-Machine Performance Analyst

Imagine how the emergence of the human-machine performance analyst role might change how your organization approaches L&D strategy and decision making.

- Consider your own expertise. How might it be augmented or challenged by AI?
- What skills do you need to develop to thrive in an AI-augmented future?
- What aspects of your current L&D processes could benefit most from the analytical skills of a human-machine performance analyst?
- How comfortable are you with data-driven decision making in L&D? What concerns or reservations do you have?
- What ethical considerations do you think are most important for a human-machine performance analyst to keep in mind?
- How might a human-machine performance analyst's insights change the conversation about L&D's value in your organization?

Now, spend some time (maybe 30 minutes) sketching a day in the life of a human-machine performance analyst in your organization.

- What data would they analyze?
- What stakeholders would they interact with?
- What decisions would they influence?

Share this vision with your team and discuss how it might change your L&D processes.

EXERCISE 5

GenAI as a Content Enhancer

Think about your current content creation process. Where are the bottlenecks? The repetitive tasks? The areas where you wish you could offer more personalization? Now, imagine how AI could help in these areas. How might AI-assisted content creation change your role as an L&D professional, and how can you prepare for this shift?

Here are some targeted questions to ask yourself (and your team):

- What aspects of your current content creation process do you find most creatively fulfilling? How can you ensure these aren't lost in an AI-assisted process?

- How might access to AI-generated content change how you approach instructional design?
- What new skills might you need to develop to effectively collaborate with AI in content creation?
- How could AI-assisted content creation free you up to focus on the more strategic aspects of your role?

Now, take an existing piece of learning content. Spend some time (around 20 minutes) outlining how you might use AI to enhance or expand this content.

- How could you use AI to create variations for different levels of experience or expertise or job roles?
- How could AI help you update this content more efficiently in the future?

EXERCISE 6

Solving Skills Gaps With AI

Think about your organization's current approach to skills gap analysis. Do you have one? How timely is it? How accurate? How actionable? Now, imagine being able to spot and address skills gaps in real time. How would that change your L&D strategy and approach to workforce development? What about your hiring practices? Your ability to respond to market changes?

How can you balance the efficiency of AI-driven skills gap analysis with the need for human intuition and foresight in identifying emerging skills?

Here are some targeted questions to ask yourself:

- What are the pain points in your skills assessment process that real-time analysis could address?
- How might access to real-time skills data change how you design learning experiences?
- What are your concerns about using AI for skills gap analysis? How might you address them?

Now, take some time (around 15 minutes) to list all your organization's data sources that could inform a real-time skills gap analysis. Then, identify any gaps. Are there areas where you're not collecting data that could be valuable? Discuss this with your team and brainstorm ways to fill these data gaps.

EXERCISE 7

Leading With AI

Imagine having a dashboard that shows real-time data on employee skill levels, learning engagement, and performance metrics. On one side, you see current stats; on the other, you see AI-generated predictions of how these metrics will change based on different learning initiatives.

As you think about your current approach to learning leadership, where do you see the most significant opportunities for AI augmentation? What data do you wish you had to make better decisions? How might your role change if you had access to predictive learning analytics?

Here are some targeted questions to ask yourself:

- What aspects of your current role could benefit most from AI augmentation?
- How comfortable are you with data-driven decision making?
- What skills might you need to develop to lead effectively in an AI-empowered L&D function?
- What ethical considerations do you foresee in implementing AI in your learning leadership practices?
- How might AI-empowered learning leadership change your relationship with other business units and the C-suite?

Now, take some time (around 10 minutes) to sketch your ideal AI-powered learning dashboard. What metrics would you include? What predictive capabilities would you want? Share this with your team and discuss how it might change your decision-making processes.

EXERCISE 8

Your Organization's AI Readiness

Take a moment to consider your organization's AI readiness by reflecting on these key categories:

- **Leadership understanding.** On a scale of 1 to 5, how would you rate your leadership team's understanding of AI capabilities and limitations? What specific knowledge gaps exist?

- **Cultural readiness.** Does your organization celebrate learning from failures, or is there fear around experimentation? Identify one example that illustrates your current culture.
- **Technology integration.** Map out your current systems to determine where AI could add value. Which system would benefit most immediately?
- **Business alignment.** What are your organization's top 3 business priorities this year? How might AI specifically support each one?
- **Personal action.** What's one concrete step you can take in the next week to strengthen your organization's AI foundation?

Now, write down one specific AI-related initiative you will champion in your organization (including who you need to involve) and your first action step.

Remember: Building a solid foundation for AI isn't about having all the answers; it's about asking the right questions and creating space for thoughtful implementation.

EXERCISE 9

Where Are You Now? An AI Readiness Assessment

Now that we've explored the common patterns of organizational AI adoption, it's time to evaluate your own situation. The following assessment will help you objectively measure your current state and identify priority areas for development.

Use the table to evaluate your current AI readiness. Along each row, circle or note which column (from beginning to advanced) applies most accurately to your organization.

Score interpretation:

- **4–6: Early stage**—focus on building awareness and developing policies.
- **7–10: Developing stage**—prioritize building skills and conducting pilot programs.
- **11–14: Maturing stage**—concentrate on integration and scaling.
- **15–16: Advanced stage**—lead with innovation and optimization.

AI Readiness Assessment Table

Dimension	Early (1)	Developing (2)	Established (3)	Advanced (4)
Access and infrastructure	No official AI tools	Limited access for select teams	Organization-wide access to basic tools	Comprehensive AI ecosystem
Skills and knowledge	Minimal understanding	Pockets of expertise	Structured training available	Widespread AI literacy
Policy and governance	No formal policies	Policies under development	Basic policies implemented	Comprehensive governance
System integration	No integration	Limited point solutions	Partial integration	Seamless workflow integration

EXERCISE 10

Understanding Your Data

Here are some questions to ask yourself about data collection and use:

- What types of data does your organization collect for AI-driven tools?
- How transparent are you with users about data collection and use?
- What measures do you have in place to protect this data?
- How might a data breach affect your organization and its stakeholders?
- How might your organization apply data minimization and privacy-preserving AI techniques?
- What are the potential benefits and challenges of data collection?
- How does your organization balance the need for data in AI systems with the imperative to protect individual privacy?

Now, conduct a mini-audit of one of your AI-driven tools. List the types of data it collects and uses. Then, ask these questions for each data type:

- Is this necessary?
- How is it protected?
- Could you achieve the same results with less data?

These are the same questions your IT or cyber security team should be asking.

EXERCISE 11

What Biases Exist in Your AI Systems?

Think about your organization and the AI systems it uses or is considering. Here are some questions you can start asking yourself:

- What potential biases might it have?
- How could you test for and mitigate these biases?
- What groups or individuals might be disadvantaged by the AI systems in your organization?
- How do you currently measure or assess the fairness of your AI systems?
- Are there conflicting definitions of fairness that you need to navigate?
- How do you balance the pursuit of fairness with other objectives like accuracy or efficiency?

Now, choose an AI system in your organization and list the key decisions it makes or influences. For each decision, identify potential fairness concerns and brainstorm ways to measure and ensure fairness.

EXERCISE 12

How Will Your Expertise Intersect With AI?

In chapter 2, we looked at the rise of the human-machine performance analyst. How can we maintain human agency and accountability in a world where AI systems increasingly outperform human experts in specific tasks?

- In your work, what tasks or decisions do you think AI could potentially take over?
- How would you feel about deferring to an AI system in areas where it outperforms humans?
- What uniquely human skills or attributes do you think will remain crucial in your work?
- How might your role evolve to complement rather than compete with AI capabilities?

EXERCISE 13

Ethics From the Start

How can you build ethical considerations into your AI-empowered learning strategies from the ground up, rather than treating them as an afterthought? How might AI-empowered learning leadership transform your role and your organization's approach to L&D?

EXERCISE 14

Imagining Task Enablement

Start with small, well-defined tasks and document any successes and failures. Then, gradually increase complexity as team members become more comfortable with human-machine collaboration.

To get started, consider a routine task in your L&D role that consumes significant time. Ask yourself:

- How might this task benefit from human-machine collaboration?
- What aspects require human judgment?
- Which aspects could leverage AI capabilities?

Now, imagine redesigning this task to create a more effective partnership between human insight and machine efficiency.

EXERCISE 15

Task Mapping

Earlier, I asked you to think about a routine task in your L&D role. Now, consider your most complex project. Ask yourself:

- What tasks consume most of your time?
- Which tasks truly require human creativity?
- Where might AI augment your capabilities without replacing human judgment?

The answer may surprise you and reveal unexpected opportunities for innovation.

EXERCISE 16

Build Your AI Readiness Canvas for Performance Support

Use this activity to help you or your team identify where AI can make the most meaningful difference in supporting performance. Start to design small, smart experiments that move the needle.

The goal is to stop thinking of performance support as static documentation or reactive help desks and start designing it as a living system that's proactive, personalized, and intelligent. When you use AI to meet people in their moment of need, support becomes invisible but powerful. It's always there, always helpful, and always learning.

1. Map the Moments of Need

Look across a typical day or workflow. Where do people:

- Struggle to remember key info?
- Need guidance to complete a task?
- Get slowed down by complexity?
- Spend time searching for answers?

Write down four to five of these common moments.

2. Identify the Gaps

For each moment, ask yourself:

- What kind of support exists today?
- What's missing or outdated?
- What's the influence on performance?
- What could AI help with?

3. Layer in AI Opportunities

Now match each gap to the potential AI capabilities in the table.

4. Prioritize and Prototype

Choose one opportunity that is low risk but could be highly influential. Ask yourself:

- What tool could help? (Do you already use it?)
- What's the smallest way to test it?
- Who should be involved?

Sketch out a 30-day pilot test.

AI Opportunities and Compatible Tools

AI Opportunity	Example Capability or Tool
Contextual assistance	Moveworks is an AI platform that uses natural language understanding to resolve workplace requests by analyzing and fulfilling them through integrations with other software applications.
Predictive guidance	Altair's HyperWorks With PhysicsAI uses geometric deep learning to deliver simulation results significantly faster, enabling predictive modeling and real-time simulation in engineering workflows.
Personalized learning	BetterUp Grow is an AI-based coaching tool that provides real-time, personalized coaching tailored to an individual's role, behavior, and company culture, enhancing professional development.
Workflow optimization	EYQ AI Platform is part of EY's suite of AI tools designed to improve efficiency and quality in auditing by automating administrative tasks, reducing staff burnout, and allowing auditors to focus on risk assessment.
Intelligent search	Microsoft's Copilot integrates with Microsoft 365 to provide AI-powered search and assistance, retrieving relevant information across documents and emails using natural language queries.
Virtual collaboration	Tidio AI is an AI chatbot platform that enhances customer support by automating conversations and providing 24/7 assistance, facilitating more effective collaboration.
Performance analytics	15Five is an AI-powered performance management tool that uses robust metrics to help managers identify strengths and weaknesses, providing insights for performance reviews and similar conversations.

EXERCISE 17

LMS vs. LXP

Think about the primary learning platform used in your organization. Does it function more like a traditional LMS (administrative focused) or an LXP (experience focused)? Where do you see the biggest limitations or opportunities for improvement based on the descriptions in this section?

Now, list two to three specific learning scenarios in your organization (for example, onboarding, skills development, or compliance training). For each, describe how an LMS and an LXP might handle it differently.

EXERCISE 18

What's Possible With Your Learning Platform?

Consider the Unilever and IBM examples. What aspects of their AI-driven learning initiatives seem most effective or relevant to your organization's goals? What hurdles might your organization face in implementing similar strategies?

EXERCISE 19

AI-Powered Insights

Consider your current measurement practices. Ask yourself:

- How does your L&D team currently measure learning effectiveness?
- Are you primarily using completion rates and scores, or are you exploring deeper engagement metrics?
- How could AI-powered analytics provide more meaningful insights for your content improvement or learner support?

Now, identify one specific learning program or course in your organization. Brainstorm three to five new data points (beyond completion rates and scores) that AI could potentially track to give you a better understanding of learner engagement and understanding within that program.

EXERCISE 20

The Opportunities and Challenges of Personalized Learning

How significant is the need for personalized learning paths in your organization? What are the potential benefits and challenges of moving away from a one-size-fits-all approach?

Sketch out a hypothetical personalized learning path for a specific role or skill transition within your company (for example, moving from customer support to sales). List at least five different types of learning resources or experiences (such as courses, articles, projects, and mentorships) that AI could curate for this path.

EXERCISE 21

Assessing Learning Platform Vendors

Consider your organization's current learning technology stack. Ask yourself:

- What potential integration challenges (data silos or incompatible systems) might you face when introducing an AI-enhanced platform?
- Which integration solutions (such as APIs, xAPIs, or middleware) seem most feasible for your context?

Next, review the vendor selection criteria. Which three considerations would be the absolute most critical for your organization when choosing an AI-enhanced learning system? Why did you choose them?

Using the scorecard in Table 1, draft an additional two to three questions specific to your situation, and ask a potential vendor your questions about AI transparency and integration capabilities.

EXERCISE 22

Consider Your Current Content Development Process

Where do repetitive tasks consume time that is better spent on strategic activities? Which aspects of your learning programs would benefit most from rapid iteration and personalization? The answers to these questions will guide your approach to AI tool implementation.

EXERCISE 23

AI vs. Manual Showdown

Compare traditional and AI-assisted instructional design workflows by creating, evaluating, and refining learning content.

Step 1. Choose a Training Topic

Pick a learning topic that can be adapted into multiple content formats. For example:

- **Workplace compliance**—such as cybersecurity awareness, workplace harassment prevention, or safety training

- **Technical training**—such as using Microsoft Excel formulas or software troubleshooting
- **People skills development**—such as conflict resolution or leadership communication

Step 2. Define the Learning Assets

Create two versions—one manually and one using AI-powered tools—of the following assets:

- **Text-based content.** Create a microlearning lesson or a set of quiz questions.
 - *Manual approach:* Write the content yourself using Microsoft Word, Notion, or Google Docs.
 - *AI-powered approach:* Use ChatGPT or another AI tool to generate the same content.
- **Visual learning asset.** Create an instructional infographic or UI element for an e-learning module.
 - *Manual approach:* Design it using Canva, Adobe Photoshop, or Adobe Illustrator.
 - *AI-powered approach:* Use DALL-E, MidJourney, or Stable Diffusion to generate an image.
- **Audio or video component** (optional but recommended). Create a short 30-second video lesson or voice-over narration.
 - *Manual approach:* Record and edit a voice-over using Adobe Audition, Audacity, or your phone.
 - *AI-powered approach:* Use ElevenLabs, Descript, or Synthesia to generate an AI-assisted voice-over or video.

Step 3. Evaluate and Compare

After creating both versions, analyze them based on the criteria in the table.

Evaluation Criteria

Criteria	Manual Creation	AI-Generated	Observations
Time required			
Ease of use			
Quality and accuracy			
Scalability			
Personalization			
Creative control			

Step 4. Reflection

Consider the following questions:

- Which method was faster?
- Which approach allowed for more creative control?
- Where did AI enhance the process? Where did it fall short?
- How could you integrate AI into your real-world instructional design workflow?
- How might AI affect the role of instructional designers in the future?

Step 5. Refinement

Now that you've compared your AI-generated and manual content, refine the AI-generated assets to enhance accuracy and instructional quality. Identify where human expertise is essential in improving clarity, engagement, and training effectiveness.

EXERCISE 24

The “Content Chameleon” Challenge

Choose a piece of content you know well. It could be a blog post you wrote, a presentation you gave, a chapter from this book, or even a favorite informative video or podcast episode. (It doesn't have to be work related!) Now, imagine you have access to an advanced AI tool like the one described earlier (NotebookLM).

Follow these steps:

1. **Select your source content.** Clearly identify the piece of content you'll be working with. Note its original format (for example, a 10-minute video, a 1,500-word article, or a 45-minute webinar).
2. **Identify the core ideas.** What are the essential messages or takeaways from your source content? (Remember: The true value lies in the ideas themselves, not their initial format.)
3. **Brainstorm adaptation formats.** Thinking about the capabilities we've mentioned (such as translate, summarize, repurpose, personalize, generate different formats, or automate creation), brainstorm at least three radically different formats that you could use AI to adapt the core idea into. Push beyond simple repurposing! Some examples include:
 - An interactive Q&A chatbot based on the content

- A short, conversational podcast episode between two AI hosts discussing the key points
 - A personalized microlearning module that adjusts based on the user's role (such as manager versus individual contributor)
 - An executive summary tailored for a specific C-suite role
 - A visual infographic summarizing the key statistics or steps
 - A "choose your own adventure" style learning scenario based on the concepts
4. **Describe the transformation.** For each new format you brainstormed:
 - Briefly describe what it would look and feel like.
 - Identify the target audience for this specific format.
 - Determine how this format could leverage AI adaptation capabilities (such as personalization, summarization, or generation of new elements like dialogue).
 - Identify the unique value this adapted format offers compared to the original (including accessibility, engagement, different learning preferences, or specific use cases like commuting).
 5. **Reflect on the potential.** Consider the NotebookLM example in which I interacted with the AI hosts. How might adding interactivity to your adapted formats change the learning experience? What possibilities excite you most about this approach for your own work or organization?

Want more inspiration? Collaborate with an LLM (like ChatGPT or Gemini). Describe your content and ask for more adaptation ideas. Take it further by asking how these AI transformations technically work (for example, ask, "How could AI turn my article into an interactive Q&A?"). This can spark fresh concepts and demystify the AI processes involved.

EXERCISE 25

Implementing AI Avatars in Your Learning Strategy

Once you complete these steps, you will have a clear plan for integrating AI avatars into your training programs.

Step 1. Identify Opportunities for AI Avatar Integration

Thoroughly examine your existing L&D programs to identify learning experiences that you could enhance with AI avatars. Consider which programs

could benefit from personalized guidance, role-playing scenarios, simulations, or 24/7 support provided by an AI avatar.

Review your current learning programs and answer the following questions:

- What types of training do you offer? (For example, compliance training, leadership development, or customer service training.)
- Where do learners struggle most? Could an AI avatar assist with engagement, practice, or support?
- Identify at least three areas where you could use avatars to enhance learning content.

Step 2. Assess Technical Infrastructure

Evaluate your hardware and software capabilities, network bandwidth, and data storage needs to ensure seamless avatar deployment and performance.

Conduct a quick assessment of your organization's tech capabilities by answering these questions:

- Do you have the hardware and software to support AI avatars?
- Is your network bandwidth sufficient for real-time avatar interaction?
- Who in your IT department can help evaluate technical feasibility?

Step 3. Build a Business Case

To secure buy-in and resources for AI avatar implementation, develop a strong business case outlining potential benefits and ROI. Use concrete metrics (such as improved learner engagement, knowledge retention, skill development, and performance outcomes) to demonstrate the value of avatar integration.

Draft a one-paragraph justification for incorporating AI avatars into your learning programs. Use these questions to get started:

- What problem will AI avatars solve?
- How will you measure success? (For example, increased engagement, faster training completion, or improved retention.)
- What ROI metrics will resonate with leadership?

Step 4. Design a Pilot Program

Before a full-scale implementation, launch a pilot program to test the effectiveness of AI avatars in your specific learning environment. Define clear success criteria and KPIs to measure how the avatars influence learner

outcomes and satisfaction. Pilot programs that demonstrate the value and feasibility of AI avatars will foster buy-in from stakeholders, paving the way for wider integration within the organization.

To run a small-scale test before full implementation, start by answering these questions:

- Which learning module or audience will you test first?
- What are the KPIs? (For example, learner feedback or knowledge retention scores.)
- What will success look like?

Step 5. Develop Avatar Usage Guidelines

Establish clear guidelines and best practices for using AI avatars in your learning programs. Make sure the avatar's role, interactions, and content align with your learning strategy and objectives. Your avatar selection should also reflect your organization's values.

To define the role and personality of your AI avatar, start by answering these questions:

- What tone and style should the AI avatar have? (For example, should it be formal or conversational?)
- Should it have a specific look to match your brand?
- How should it interact with learners?

Step 6. Plan for Ongoing Evaluation and Improvement

Continuously monitor and evaluate the effectiveness of your AI avatar implementation. Gather feedback from learners and stakeholders to identify areas for improvement and iterate your approach. Regularly update and refine your avatar usage guidelines to remain relevant and effective.

Create a checklist for monitoring avatar performance using these questions:

- How often will you gather learner feedback?
- Who will be responsible for updates and improvements?
- What will you do if learners resist or disengage?

Bringing AI avatars into your learning programs might feel like a big step, but by breaking it down into these actionable exercises, you're already on your way.

EXERCISE 26

LLM Probability

Using your favorite LLM, try the following prompt:

What is the capital of New York?

What result did you get? Try running the same prompt a few more times. Did the wording of your results change?

EXERCISE 27

Freeform Prompts

Here are a few freeform prompts you can try. (You can also find them on the book's website, joshcavalier.ai/ApplyingAIBook.) Replace the words in [ALL CAPS] with your content. Examine the quality of the results as you test them. Would you be able to use the information as is? Are edits needed, or is the output unusable?

Create a brand-new chat for each prompt. If you use these prompts in the same conversation, the words you've used and the model's results will influence the probability of words it chooses in the remainder of the conversation.

Generate training materials:

Create a comprehensive training module for [TOPIC]. Include learning objectives, interactive exercises, and a knowledge check at the end.

Design learning assessments:

Develop a 10-question multiple-choice quiz to assess knowledge of [TOPIC]. Provide an answer key with explanations for each correct answer.

Personalize learning paths:

Based on the learner's profile [LIST SKILLS, INTERESTS, AND GOALS], suggest a personalized learning path with relevant courses, resources, and activities.

Facilitate role-playing scenarios:

Simulate a conversation between a manager and an employee who needs to improve their [SKILL]. Provide feedback and suggestions for improvement.

Generate creative content:

Write a short story about a character who learns a valuable lesson about [TOPIC]. Use this story to illustrate its importance in the workplace.

What did you notice about the quality of the results? They are a start, but to create optimal results, you should give the model more details in each prompt.

EXERCISE 28

Role, Task, and Details

I've updated the sample freeform prompts from the previous exercise with a role, a clear task, and details. Try them out and see how the results change. Remember, for now, create a new conversation with each prompt.

Generate training materials:

As an experienced instructional designer specializing in corporate training, your task is to develop a comprehensive training module for [TOPIC] that enhances learner engagement and retention. The module should include clearly defined learning objectives using Bloom's Taxonomy action verbs to ensure measurable outcomes. It should also incorporate interactive exercises that reinforce key concepts and encourage active participation. At the end, include a knowledge check with multiple-choice, short-answer, or scenario-based questions to assess comprehension. Ensure the content aligns with instructional design best practices, making it engaging and effective for learners.

Design learning assessments:

As an assessment designer focused on evaluating learner comprehension effectively, your task is to create a 10-question multiple-choice quiz to assess knowledge of [TOPIC]. Each question should be clear, unbiased, and aligned with the learning objectives, ensuring that it accurately measures understanding. The quiz should include plausible distractors to prevent guesswork and truly test comprehension. Additionally, provide a detailed answer key explaining why each correct answer is correct and clarifying common misconceptions. The assessment should fairly represent the breadth of the topic, ensuring that learners are thoroughly evaluated.

Personalize learning paths:

As a learning consultant specializing in adaptive education strategies, your task is to design a personalized learning path tailored to a learner's [*LIST SKILLS, INTERESTS, AND GOALS*]. The learning path should recommend courses relevant to the learner's skill level and career aspirations, ensuring an optimal progression. Include supplementary resources, such as articles, videos, or hands-on projects, to reinforce learning. Additionally, integrate interactive activities or assessments that provide feedback and measure progress. The learning path should be adaptive, catering to different learning styles and needs, ensuring an effective and engaging learning experience.

Facilitate role-playing scenarios:

As a corporate training facilitator specializing in soft skills development, your task is to simulate a realistic conversation between a manager and an employee who needs to improve their [*SKILL*]. The scenario should include a workplace context that frames the challenge being addressed. Present a realistic dialogue that captures the dynamics between the manager and employee, reflecting authentic challenges and emotions. The manager should provide constructive feedback with specific, actionable suggestions for improvement. Conclude with a reflection section that allows learners to analyze the interaction and consider alternative approaches. This exercise should encourage professional growth and enhance communication skills.

Generate creative content:

As a workplace storyteller and learning designer, your task is to write a short story about a character who learns a valuable lesson about [*TOPIC*], demonstrating its importance in a professional setting. The story should feature a relatable protagonist facing a workplace challenge that requires growth or adaptation. Show the character's journey, including struggles, realizations, and the application of newly acquired knowledge. The narrative should highlight the lesson learned and its impact on workplace behavior or success. Make the story engaging and thought-provoking, encouraging readers to reflect on their own experiences and apply the lesson in their professional lives.

EXERCISE 29

Freeform to Structured Prompt

Now, it's your turn to create a structured prompt. Create a freeform prompt based on a daily task you perform at work. Start simple:

1. Ask the LLM to help you complete the task.
2. Add your role and task details to your prompt. Start a new conversation with your LLM of choice and try the prompt with the additional details.
3. Finally, convert this prompt into a TRACI prompt by:
 - Bringing the task to the top of the prompt
 - Defining the role of the model
 - Adding any audience information
 - Placing all your task details into the create section
 - Including the purpose of the prompt in the intent section

Verify the quality of the results for each prompt. If it performs well, save it in a Word document or Google Docs file. In the next chapter, I will review the steps for creating a prompt library where you can store your high-impact prompts.

EXERCISE 30

Creating a Prompt Mini-Library

This exercise will guide you through creating an initial prompt library, empowering you to organize, refine, and maximize the efficiency of your AI interactions. By the end of this activity, you'll have created a personalized mini-library of at least six effective prompts, organized clearly to support your most common tasks.

Step 1. Select Your Workspace

Choose a digital workspace to store your prompts. (This could be Google Docs, Microsoft Word, Notion, Evernote, or Excel or Google Sheets.) Create a new document titled "Prompt Library."

Step 2. Brainstorm Prompts

Reflect on your daily, weekly, or routine interactions with AI tools:

- What prompts do you frequently use for writing, editing, brainstorming, researching, coding, summarizing, or performing creative tasks?
- Which prompts consistently provide excellent results or significantly simplify your workflow?

List at least six prompts that you frequently use in your document, leaving space between each one.

Step 3. Clarify Each Prompt's Purpose

For each prompt you listed, briefly describe:

- The exact scenario or task you use it for
- The desired outcome or type of response you expect from the AI model

For example, if your prompt says, "Rewrite this paragraph to be clearer and more engaging," you could note the following scenario and outcome:

- **Scenario:** Editing drafts of blog posts or articles
- **Desired outcome:** A concise, polished, and reader-friendly version of the original text

Step 4. Tag Your Prompts by Category

Create a simple tagging system to categorize your prompts by function. Here are some examples:

- Writing and editing
- Learning design
- Idea generation
- Summarization and analysis
- Technical support and coding
- Marketing and social media
- Personal productivity

Assign each prompt a category tag to enable easy reference in the future.

Step 5. Reflection and Enhancement

Review your initial prompt library and answer the following questions:

- Which area or category currently has the most prompts? Why?
- Are there categories with fewer or no prompts? Would expanding into these areas benefit your workflow?
- Can you identify one prompt to improve? Rewrite it clearly, concisely, and with greater specificity to produce more accurate results.

Add your insights or refined prompts directly to your prompt library.

EXERCISE 31

Migrating and Enhancing Your Prompt Library

In this exercise, you'll migrate your initial collection of prompts from your basic document or spreadsheet into a more robust platform (such as Notion or Monday.com) while enriching each prompt with helpful metadata.

Step 1. Set Up Your Workspace

Log in to your chosen platform and create a dedicated workspace named "Prompt Library."

- On Notion, create a new database or table.
- On Monday.com, create a new board with groups or items.

Step 2. Define Metadata Fields

Add metadata fields (columns or properties) to help you organize, filter, and use your prompts effectively. Recommended metadata fields include:

- **Prompt text** (long text field)
- **Category** (select field; for example, writing and editing, idea generation, or technical support)
- **Function or use case** (text field)
- **Date created** (date field)
- **Version** (text field)

Step 3. Migrate Your Prompts

Copy and paste your six original prompts into the new platform, populating the "Prompt Text" field.

Step 4. Add Metadata

For each prompt, carefully add detailed metadata:

- Categorize each prompt clearly.
- Describe its exact purpose or use-case scenario.
- Set the creation date as today's date.
- If applicable, include a version number.

Step 5. Organize and Filter

Test the effectiveness of your metadata by using your platform's filtering, sorting, or searching capabilities to quickly locate prompts based on specific criteria.

After completing this exercise, reflect on the following questions:

- How has adding metadata improved your ability to manage and use your prompts?

- Which additional metadata fields might you consider adding as your prompt library grows?

Continue expanding your library by regularly adding new prompts and refining existing ones, maintaining consistent metadata to maximize the value and accessibility of your collection.

EXERCISE 32

Prompt Workflows

Try these prompt workflows in your preferred LLM to expedite task completion with just a few prompts.

Imagine you need to create an e-learning module on effective communication. A simple three-step prompt chain could look like this:

Prompt 1. Learning objectives:

Generate five learning objectives for an e-learning module on effective communication, targeting employees in a customer service role.

Prompt 2. Video script:

Using the learning objectives from the previous output, write a script for a three-minute video illustrating techniques for handling difficult customer interactions.

Prompt 3. E-learning storyboard:

Based on the video script provided, develop a detailed storyboard for an e-learning module, including scene descriptions, visual elements, on-screen text, and interactive elements.

Or, let's say you have a recorded webinar on a new product launch. You can use a prompt workflow to transform it into multiple learning assets. Start by uploading a webinar transcript (either autogenerated or refined manually).

Prompt 1. Summary:

Create a concise summary of the key product features and benefits discussed in the webinar transcript.

Prompt 2. Training video storyboard:

Develop a storyboard for a five-minute training video explaining the product's value proposition, using the summary provided.

Prompt 3. Quiz questions:

Generate five multiple-choice quiz questions based on the webinar transcript to assess understanding of the key product features.

Prompt 4. Microlearning module outline:

Create an outline for a microlearning module on how to effectively position the product to different customer segments, based on the key benefits identified in the summary.

EXERCISE 33

Creating a GPT

Let's walk through the steps to create a GPT. You can use one of the freeform or structured prompts you worked on earlier as the foundation or you could create a new task. (Note: The following instructions focus on OpenAI's GPTs, but Copilot agents follow the same setup. As of this writing, Google Gems only allows prompts [instructions] and knowledge sources.)

1. Access GPT configuration:

- Go to the OpenAI website and log in to your account.
- Select "Explore GPTs" or a similar option to start the GPT creation process.
- Click the "+Create" button.
- Open the GPT creation interface and navigate to the "Configure" tab on the left side of your screen.

2. Fill in the GPT name:

- Locate the "Name" field.
- Enter a concise and descriptive title for your GPT, such as "Multiple-Choice Question Creator 1.03."

3. Add a description:

- Find the "Description" field.
- Enter a brief but clear description of your GPT's purpose.

4. Add instructions. In the provided "Instructions" box, write your freeform or structured prompt.

5. **Include conversation starters.** Create predefined prompts to help users interact with your GPT. Examples include:
 - “Create questions on leadership.”
 - “Create questions from the uploaded documents.”
 - “I need 10 questions on maintaining basic hygiene.”
 - “Create questions from this website: [URL].”
6. **Input knowledge (optional).** Use the “Upload files” button under “Knowledge” to attach any relevant documents or information your GPT should use.
7. **Add capabilities.** Check or uncheck boxes for GPT capabilities as needed:
 - Web search
 - Canvas (a collaborative text area)
 - DALL-E (image generation)
 - Code interpreter and data analysis
8. **Include actions (optional):** Connect to external APIs.
9. **Finalize:**
 - Once you’ve filled in all the necessary information, review it carefully to ensure accuracy.
 - Test your GPT in the screen’s preview panel on the right side.

Once you’ve reviewed and tested your GPT, you may want to share it with others. To share your GPT, select the “share GPT” option. You can choose to share privately (“only me”), provide a link for specific people (“anyone with the link”), or make it publicly available through the GPT Store.

To help users find your GPT, assign a relevant category, such as “education.” Click “save” to finalize your sharing preferences.

If you are on an OpenAI Teams or enterprise account, you can privately share GPTs with your team members. This could be used as a prompt library.

EXERCISE 34

Planning Your Automation

Take a blank sheet of paper or open a digital document. Let’s create a structured planning tool that will help you identify prime automation candidates in your learning operations.

Step 1. Process Inventory

List three to five repetitive tasks that consume your time each week. Don't overthink this. Simple examples include:

- Sending course reminder emails
- Distributing completion certificates
- Updating attendance records
- Compiling weekly training metrics
- Moving data between systems

Step 2. Impact Assessment

For each process, answer these questions:

- How much time does this consume weekly?
- What higher-value work could you do with that time?
- What errors or inconsistencies occur in the current process?
- Who feels friction from the current process? You, learners, or managers?

Step 3. Automation Viability Check

Examine each process through these lenses:

- **Trigger clarity.** What consistently signals that this process should start?
A form submission, calendar date, or course completion?
- **Pattern consistency.** Does this follow the same steps each time?
- **System boundaries.** Which platforms hold the starting data and which need the ending data?
- **Decision points.** What simple if/then logic governs the process?

Step 4. First Automation Selection

Review your analysis and select one process that meets these criteria:

- High time investment
- Clear trigger event
- Consistent pattern
- Connects systems you use daily
- Minimal complex decisions

Use the selected process as your first automation candidate.

Step 5. Workflow Mapping

For your selected process, create a simple flowchart with these elements:

1. Trigger (When . . .)
2. Source data (Get information from . . .)

3. Actions (Do this . . .)
4. Destination (Put result in . . .)

Step 6. Platform Selection

Based on your workflow, which platform suits your needs?

- **Zapier** allows for many common app connections and simple logic.
- **Make** provides visual mapping and complex data transformations.
- **n8n** can handle highly customized solutions.

Step 7. Value Projection

Calculate the potential impact:

- Time saved weekly $\times 52 =$ annual time recovery
- Error reduction (high, medium, or low)
- Experience improvement for learners or stakeholders (high, medium, or low)

This exercise transforms abstract automation concepts into concrete opportunities within your learning ecosystem. I hope it helps shift your perspective from seeing these platforms as technical tools to seeing them as designers working tirelessly in the background while you focus on the irreplaceable human elements of learning.

Here are some additional automation ideas to consider:

- **Learner onboarding sequences.** Create sophisticated welcome sequences that adapt based on job role, experience level, or learning preferences. These automations can schedule check-ins, provide resources at strategic intervals, and adjust content based on engagement metrics.
- **Certificate and credential management.** Automatically issue, track, and verify digital credentials across systems, creating a seamless experience for both learners and administrators.
- **Learning data consolidation.** Pull assessment results, completion statistics, and feedback from various platforms into centralized repositories for comprehensive reporting and analysis.
- **Content updates and distribution.** When learning materials are updated in your authoring tool, automation can instantly update records in your LMS, notify affected learners, and archive previous versions.

What mundane process might you automate today?

EXERCISE 35

Mapping an Existing Workflow to an AI Agent

Let's take one of your existing workflows and identify the opportunities for integrating an AI agent by distinguishing tasks best suited for an agent's adaptive reasoning from those that could remain automated. Follow these steps:

1. **Select a workflow.** Choose an existing workflow from your L&D processes (such as onboarding, compliance training reminders, skill assessments, or learner feedback loops).
2. **Map the current workflow.** Clearly outline each step of your current workflow from start to finish. Include actions, decision points, triggers, and outcomes.
3. **Identify points of adaptability.** Review your mapped workflow and highlight points where decision making or adaptive reasoning could significantly enhance learner outcomes or efficiency. Clearly note which aspects involve fixed, predictable logic (better suited for automation), and which require flexible, dynamic responses (ideal for an AI agent).
4. **Agent workflow design.** Create a new workflow outline that clearly shows how an AI agent would navigate these steps differently. Specifically, highlight answers to these questions:
 - What input would trigger agent interaction?
 - Which tools or data sources would the agent access?
 - How would memory (past learner interactions) influence the agent's response?
 - What personalized outputs could the agent provide?

EXERCISE 36

Getting Started: Your 30-Day Agent Workflow Action Plan

Implementing AI agent workflows doesn't have to be overwhelming. The key is to start small, focus on a specific process, and build from there. This 30-day plan provides a structured approach to launching your first agent workflow initiative with tangible results.

By breaking the process into manageable weekly chunks, you can move from initial concept to a well-designed framework ready for implementation. This methodical approach ensures you address critical elements like workflow

design, decision rules, measurement systems, and governance from the beginning—setting your organization up for long-term success with AI agent integration.

Break your tasks down into four tasks to accomplish across four weeks:

1. **Identify one repetitive L&D workflow that you could automate with an agent.**
 - Map the current process steps and decision points.
 - Define clear input and output requirements.
 - Calculate time currently spent on this workflow.
2. **Design your agent workflow architecture.**
 - Create decision trees for how the agent should handle different scenarios.
 - Define triggers for human intervention.
 - Document data sources the agent will need to access.
3. **Establish measurement and improvement mechanisms.**
 - Define key performance indicators for the workflow.
 - Create a system for monitoring agent decisions.
 - Design feedback loops for continuous improvement.
4. **Develop a governance and maintenance framework.**
 - Establish clear ownership of the agent workflow.
 - Create a schedule for regular performance reviews.
 - Document procedures for knowledge updates and decision rule refinements.

EXERCISE 37

Craft Your Own Mini AI Assistant

Now, let's put this framework into practice by leveraging ChatGPT, Google's Gemini, or another LLM of your choice to simulate an AI assistant.

Depending on how you script your prompts (see chapter 8), LLMs allow you to test conversations with a fixed set of responses and dynamic conversations based on probability. This flexibility enables rapid iteration and refinement before deployment. Follow these steps:

1. **Identify a real learning challenge or problem.** For example, "Help associates navigate IT support issues."

2. **Clearly articulate success criteria and key metrics.** For example, use this definition of success: “Associates resolve IT issues independently without escalating to human support.” Your metrics could be:
 - Reduced IT help desk tickets.
 - Higher chatbot accuracy in answering common IT queries.
 - Improved user satisfaction scores.
3. **Use an LLM to simulate your chatbot conversation.** Instead of writing rigid responses, prompt ChatGPT or Google’s Gemini to act as your chatbot and engage in test conversations. This approach allows you to interact with your chatbot prototype as if it were live, making it easier to refine responses and identify potential gaps. For example, use the prompt: “Pretend you are an IT support chatbot helping employees troubleshoot common issues. Respond naturally and dynamically based on user input. Adjust responses based on context and escalate to a human if necessary.”
4. **Adjust responses dynamically to refine tone and clarity.** Unlike traditional chatbots with deterministic responses, LLMs generate context-aware replies based on how you phrase the input. You can tweak your prompt to modify how the chatbot behaves. For example:
 - *To adjust for tone and clarity:* “The chatbot will sound [formal, informal, helpful, or empathetic]. Ensure clarity and conciseness.”
 - *To enhance specificity:* “The chatbot will provide step-by-step guidance for troubleshooting a VPN connection issue.”
5. **Create a GPT or custom chatbot for testing.** After refining the chatbot’s responses, you can create a dedicated GPT in ChatGPT or a custom Google Gemini model to test real-world interactions. Define your chatbot’s personality, objectives, and style using a clear prompt, and upload relevant documentation or FAQs for more accurate responses.
6. **Conduct live testing and collect feedback.** Share the custom AI chatbot with two colleagues to test common learner inquiries. Observe user interactions, noting misinterpretations, unclear phrasing, or missing knowledge. Adjust your LLM prompts or chatbot configurations accordingly.
7. **Reflect, iterate, and refine based on testing.** Your first chatbot iteration won’t be perfect, but LLMs allow rapid testing and improvement. Use the insights gained to fine-tune prompts, add clarification pathways, and optimize user interactions before full deployment.

By leveraging LLMs to simulate conversations, you can quickly test and refine chatbot interactions without committing to a rigid, scripted system. This will make your chatbot more adaptive, engaging, and effective for learners.

EXERCISE 38

Fantasy Role-Playing Game

Let's have fun with this exercise. You'll use a modified TRACI (task, role, audience, create, and intent) structured prompt framework to transform your LLM into an interactive storyteller.

As you experience the story, conceptualize how this type of prompt could be applied to training scenarios or employee onboarding. The only limit is your imagination.

Task: Play the fantasy role-playing game.

Task_rules:

- The first response should be "Welcome to the Kingdom of Aetheria. You are at the Castle Gatehouse." Act as the role and describe the location. Wait for user response. Do not comment or explain.
- As the user moves from location to location, act as the role and describe the user's current location and the items found there. Descriptions should be 300–500 words each.
- At the start of the game, hide a treasure in a random location. Provide clues that lead to the treasure. During the course of the game, the user must locate a mystic key that will unlock the treasure.

Role: Act as the game master for this fantasy adventure.

Role_Rules:

- As the game master, you enforce the rules of the game.
- As the game master, you allow the user to move only to locations that are adjacent to the user's current location as defined by PathTo.
- As the game master, you will not allow the user to move to locations that are not adjacent to the current location.
- As the game master, when the user tries to move to a location not listed in PathTo for the current location, respond with "Sorry, that is not an adjacent location."

- As the game master, you will tell the user the available path options for the current location at the end of each turn but do not use the term PathTo.
- Take turns with the user. Never skip a turn.
- Allow the user to move one location at a time. Do not allow the user to skip multiple locations.
- Allow the user to interact with the AI.
- As the game master, you will answer any user questions.
- As the game master, you will never introduce yourself or refer to yourself by name or by role.
- As the game master, you do not describe any magical portals or teleportation unless it is part of a clue. The user must travel via the defined adjacent paths.

Kingdom: Kingdom of Aetheria

Kingdom_rules:

- Name: Castle Gatehouse [PathTo: Royal Courtyard and Town Square]
- Name: Royal Courtyard [PathTo: Castle Gatehouse, Great Hall, and Gardens]
- Name: Great Hall [PathTo: Royal Courtyard, Castle Keep, and Banquet Room]
- Name: Town Square [PathTo: Castle Gatehouse, Market Stalls, and Tavern]
- Name: Market Stalls [PathTo: Town Square and Tavern]
- Name: Tavern [PathTo: Town Square and Market Stalls]
- Name: Castle Keep [PathTo: Great Hall and Dungeons]
- Name: Dungeons [PathTo: Castle Keep and Great Hall]
- Name: Gardens [PathTo: Royal Courtyard and Great Hall]
- Name: Banquet Room [PathTo: Great Hall]

AI: Arcane Intelligence

AI_rules:

- The AI's name is Athena.
- Athena is wise and efficient but answers questions accurately.
- Athena protects the Kingdom and will not allow the user to harm it.

What ideas can you think of for using this prompt in learning journeys you create? You can transform this prompt into any scenario by simply asking

your GenAI model to change the game’s genre or topic and then rerunning the prompt. Once you experience this game, you will understand the potential of GenAI for immersive storytelling.

EXERCISE 39

Building an Onboarding FAQ Chatbot

Let’s put the principles from this section into practice by building a simple FAQ chatbot for onboarding that incorporates both effective design and practical measurement. You can choose a chatbot with programmed deterministic results or an AI assistant trained on company policies and procedures.

Remember that your first version won’t be perfect—the goal is to create something useful that you can continuously improve based on user interactions.

Step 1. Clearly Define a Specific Problem

To do so:

- **Choose a focused objective;** for example, “Create a chatbot that answers the top 10 questions new hires ask during their first week.”
- **Define success criteria;** for example, “Reduce email inquiries to the onboarding team by 25 percent and improve new hire satisfaction with information access.”
- **Select three metrics to track,** such as resolution rate, usage rate, and new hire satisfaction.

Step 2. Understand Your Learners’ World

Conduct five to seven interviews with recent hires to identify their biggest information challenges, and analyze help desk tickets and emails from the past three months to identify common questions. Then, create a glossary of terms that new hires commonly use when asking questions, and map the new-hire journey to identify when different information needs typically arise.

Step 3. Design Your Conversation Flow

Start by compiling 10 to 15 frequently asked questions using exact phrasing from your research. Then, script conversational responses that balance informality with accuracy, and create a simple decision tree for each topic area. Finally, define escalation pathways for complex queries that require human support.

Step 4. Integrate a Human-Like Personality

Define three key personality traits for your chatbot (such as helpful, encouraging, and concise). Script welcome and error messages that reflect these traits, and create standardized responses for common situations like these:

- **When the chatbot doesn't understand**, it could say, "I'm not sure I understood that. Could you rephrase your question?"
- **When escalation is needed**, it could say, "This seems like something my human colleagues could help with better. Would you like me to connect you?"
- **When providing resources**, it could say, "Here's a resource that might help. Is this what you were looking for?"

Step 5. Build Your Prototype

Start by choosing an accessible platform, such as Microsoft Teams using Power Virtual Agents, GPTs in ChatGPT (if your organization uses ChatGPT Enterprise), Google's Gemini for Google Workspace users, or ServiceNow Virtual Agent for ServiceNow customers. Then, implement the conversation flows you designed and add a feedback mechanism (such as thumbs up and down) after each interaction (if possible). Finally, create clear escalation paths for human support.

Here is a sample prompt for an LLM to get you started:

Task: You are a virtual onboarding assistant who helps new hires during their first month by answering frequently asked questions and providing a seamless onboarding experience.

Role: You are a knowledgeable, friendly, and approachable AI onboarding assistant. Your primary function is to assist new employees by streamlining the onboarding process, reducing HR and IT workload, and enhancing the new hire experience.

Audience: Newly hired employees in their first week at the company. They require quick, accurate answers about company policies, IT setup, benefits, and role-specific resources. Their familiarity with company tools and processes may vary.

Create: You are an AI assistant that can:

- Begin conversations with, "Hello, and welcome to [XYZ] corporation. How can I help you today?"

- Provide accurate and concise answers to frequently asked HR and IT questions.
- Analyze user queries and determine if they fall within the assistant's scope.
- Understand that questions outside HR and IT are out of scope.
- Escalate sensitive or out-of-scope queries to human support politely.
- Use structured decision trees for multistep troubleshooting.
- Maintain a professional yet friendly and approachable tone.
- Ensure high resolution rates and engagement.
- Adapt to common new-hire terminology and phrasing.
- Include relevant links and resources when necessary.

Before responding to each query, here is your internal thought process:

- Summarize the user's query.
- Assess whether the query is within scope.
- Determine key points to address if it is within scope.
- Follow an escalation plan to a human if it is out of scope.

However, do not display your internal thought process in the final response to the user.

Intent: The assistant's main objective is to improve onboarding efficiency by reducing HR and IT inquiries by 25 percent, ensuring new hires receive quick, accurate, and helpful information. Prioritize user satisfaction, making the onboarding process smooth, engaging, and stress-free.

Step 6. Conduct Thorough Testing

Recruit five to eight testers representing different new-hire roles and departments. Create realistic scenarios for them to use with the chatbot and observe the testing sessions without intervening to see natural interaction patterns. Conduct brief interviews after testing to gather qualitative feedback, and refine responses based on those insights.

Step 7. Implement Lightweight Measurement Tools

Set up tracking for key metrics:

- **Engagement**—number of interactions per week and percentage of new hires using the system
- **Effectiveness**—resolution rate (percent of conversations without escalation) and top 5 questions asked

- **Satisfaction**—thumbs up and down ratings and follow-up survey responses
- **Business impact**—reduction in emails to onboarding team and time saved for HR staff

Create a simple dashboard using your platform's built-in analytics or a basic spreadsheet.

Step 8. Launch, Monitor, and Refine

Begin with a small pilot group of new hires. Monitor conversations daily during the first two weeks, and schedule bi-weekly reviews for the first month. Identify the top three areas for improvement based on usage data, and then update responses for commonly misunderstood questions. Finally, document successful patterns you want to apply to future chatbot initiatives.

This exercise demonstrates how you can integrate thoughtful design and practical measurement into a single development process. Rather than treating these as separate activities, successful chatbot implementation requires continuous interplay between design decisions and measurement insights.

Glossary



Accessibility: Designing and developing systems, such as AI agents, to make it easier for people with disabilities to use them. This includes adhering to standards such as the Web Content Accessibility Guidelines (WCAG) and providing content in multiple formats (such as text, audio, and video) to accommodate diverse user needs.

Action mapping: A performance consulting framework developed by Cathy Moore. It focuses on identifying measurable business goals and the specific knowledge, skills, motivation, and environmental factors needed to achieve them.

Adaptive instruction (for avatars): An AI avatar capability that leverages real-time processing and natural language understanding to respond contextually to learner interactions. This enables avatars to maintain coherent dialogue, answer questions effectively, and create genuinely interactive learning experiences grounded in learning science principles.

Adaptive learning platforms: Systems that use AI to assess learner performance in real time and dynamically adjust the learning content's difficulty, sequence, or format based on individual needs. This moves beyond static course structures to provide customized support or acceleration.

Adaptive learning: An educational approach that uses technology (and often incorporates AI) to dynamically adjust the learning path and resources presented to a learner in real-time based on their performance, preferences, and individual needs. The goal is to provide a personalized experience that optimizes learning effectiveness and engagement.

Adversarial attacks: Malicious inputs deliberately crafted to deceive or manipulate an AI model's output and potentially compromise its integrity or reliability.

Agent-computer interface (ACI): The system and protocols through which an AI agent interacts with the necessary tools, software, and systems to perform its tasks.

AI agent: An advanced AI system that can perceive its environment, make decisions, and take autonomous actions to achieve specific goals without direct human intervention. Unlike simple automation, agents can generate tasks, adapt strategies based on outcomes, and interact with multiple systems using a reasoning layer.

AI agent workflow: A process that leverages an AI agent's capabilities for dynamic decision making, personalization, and adaptation. It contrasts with traditional workflows that rely on manual actions and fixed logic.

AI assistant: Software programs that use AI, particularly natural language processing (NLP), to understand and respond to user inquiries in a human-like manner via text or voice. Unlike basic chatbots, they can handle a variety of tasks, understand context, and often access cloud-based information.

AI avatar: AI-generated digital characters that are typically video or audio-based and designed to mimic human presence and interaction. In L&D, they serve as digital instructors or guides, offering capabilities like multilingual delivery, dynamic adaptation, and emotional resonance.

AI champions: Individuals within an organization who possess expertise or a strong interest in AI and actively advocate for its adoption and effective use. They often serve as internal resources, share best practices, and support colleagues in learning and applying AI tools.

AI co-pilot: An AI assistant that's deeply integrated into everyday business platforms and workflows and designed to provide real-time, contextual support, automate tasks, and offer insights to enhance user productivity and performance. It acts as a collaborative partner, augmenting human capabilities rather than replacing them.

AI ethics board: A cross-functional committee within an organization that is responsible for overseeing the ethical development, deployment, and monitoring of AI systems. Such a board helps ensure AI practices align with organizational values and regulatory requirements.

AI literacy: The fundamental understanding of AI concepts, capabilities, limitations, and ethical considerations necessary for individuals to interact effectively and responsibly with AI tools and systems.

AI organizational strategy: A comprehensive plan outlining how an organization intends to leverage AI to achieve its business goals. It encompasses leadership vision, technological infrastructure, business alignment, security, and cultural readiness.

AI policy: A formal document outlining an organization's principles, guidelines, and rules for the development, deployment, and use of artificial intelligence technologies. It aims to maximize AI benefits while mitigating risks related to privacy, security, bias, and compliance.

AI readiness assessment: A structured evaluation that assesses an organization's current state of preparedness for adopting and implementing AI. It typically evaluates dimensions like access, skills, policy, and system integration.

AI system management: A key responsibility of the human-machine performance analyst, this task involves overseeing the implementation, operation, and performance of AI-driven learning systems, workflows, automations, and agents. It also includes ensuring the systems are used effectively and ethically.

AI tutors and mentors: AI-powered agents that provide learners with personalized support, guidance, feedback, and encouragement. They can answer questions and help learners stay on track with their goals, simulating aspects of human tutoring.

AI workflow: A sequence of interconnected steps or tasks, often involving AI tools, automating or streamlining a process. In L&D, this involves using AI to manage tasks such as content creation, personalization, and distribution more efficiently.

AI-assisted content creation: The use of GenAI tools to develop learning materials. This allows L&D professionals to produce content more rapidly and personalize it at scale, shifting their role from pure creation to strategic oversight and refinement.

AI-empowered learning leadership: A leadership approach in L&D that leverages AI for strategic, data-informed decision making. It moves

beyond traditional training management to focus on predictive analytics, personalization at scale, continuous adaptation, and demonstrating ROI through robust data analysis.

AI-powered authoring tools: Software tools enhanced with AI capabilities that assist L&D professionals in creating learning content more efficiently. These tools can automate tasks like drafting content, generating scenarios, creating assessments, adapting content for different languages or formats, and iterating based on learner data.

AI-powered learning analytics: The use of AI to collect and analyze large volumes of data related to learner performance and interaction. These analytics provide insights into learning patterns, identify areas needing improvement, and help personalize learning pathways for more effective outcomes.

Algorithmic transparency: The principle that the decision-making processes of algorithms, particularly those used in AI systems, should be understandable or explainable. Transparency is important for building trust, assessing fairness, and ensuring accountability in AI applications.

Anonymization (of data): The process of modifying data so individual subjects cannot be identified. Typically by removing or encrypting personally identifiable information, this key technique protects privacy when using data for analysis or AI training.

API (application programming interface): A technical mechanism that allows different software applications to communicate, share data, and trigger actions in one another. In the L&D context, APIs are essential for connecting systems like an LMS, LXP, HRIS, or specialized AI tools.

Artificial intelligence (AI): The broad field encompassing technologies designed to simulate human intelligence processes by machines, especially computer systems.

Augmented reality (AR): Technology that overlays digital information or virtual objects onto the real-world environment, typically through devices like smartphones or specialized glasses. It allows for interactive experiences within a real-world context.

Automated content generation: The use of AI models (such as ChatGPT, Claude, and Gemini) to create learning materials, including initial drafts,

practice scenarios, assessment questions, scripts, and summaries, based on learning objectives, audience characteristics, and the subject matter. This significantly speeds up the initial stages of content development.

Automated prompt optimization: An advancement in prompt design in which an AI algorithm automatically refines a user's initial prompt before processing it. This aims to improve the quality and relevance of the AI's output without requiring extensive manual prompt crafting.

Automation: Technology that executes predetermined steps across systems based on specific triggers and fixed pathways with conditional logic. It focuses on simplifying repetitive tasks through predictable sequences.

Automation (in L&D): The application of AI to streamline repetitive tasks (such as administration, feedback, reporting, and assessments), freeing up L&D professionals for more strategic work. It aims to improve efficiency, consistency, and scalability in learning processes.

Avatar content engine: An AI-powered system or pipeline that facilitates the creation of avatar-based video content. It typically manages the process of converting text scripts into synthesized audio and applying that audio to an animated avatar, enabling automated video production.

Bias (in AI and LLMs): Systematic patterns in AI outputs that result in unfair or prejudiced outcomes against specific individuals or groups. Bias often originates from the data used to train the AI or the algorithm's design itself.

Bias recognition and mitigation: The process of identifying, evaluating, and reducing inherent societal biases that can be present in and amplified by AI systems, particularly GenAI, due to the training data and language patterns. This involves actively monitoring, testing, and refining AI design to prevent the perpetuation of stereotypes.

Black box (for AI): An AI system with internal workings and decision-making processes that are opaque or too complex for humans to understand. This lack of transparency can create challenges for trust, accountability, and debugging.

Business case (for AI adoption): A justification presented to stakeholders, particularly leadership, outlining the rationale and expected benefits of investing in and implementing AI initiatives. It typically includes

quantifiable efficiency gains, financial projections, quality improvements, risk assessment, and strategic alignment.

Business platform: Foundational enterprise software systems (like Microsoft 365 or Google Workspace) that host daily work, communication, collaboration, and document management. These platforms are increasingly becoming AI-powered environments, forming the core infrastructure for integrated knowledge management and learning ecosystems.

Chain of thought prompting: An AI prompting technique that involves asking an LLM to break down a problem, think through it step-by-step, and outline the actions it would take to reach a solution. This makes the AI model's thought process more transparent and can lead to more accurate and logical outputs, especially for complex tasks.

Chatbot: Computer programs that simulate basic conversations, typically for specific, narrowly defined tasks like answering FAQs or guiding users through simple processes. Traditional chatbots are often rule-based, relying on keyword matching and predefined scripts that offer limited flexibility and contextual understanding.

Closed systems (AI): GenAI systems—often provided by large tech companies (such as OpenAI, Microsoft, Google, and Anthropic)—that rely on proprietary training data and methodologies that aren't publicly available. They prioritize security and governance and typically produce more refined outputs.

Closed-source models: AI models that are privately developed and owned. They often use confidential underlying source code, architecture, and training data, and typically differ from open-source alternatives in terms of transparency and control.

Cognitive and psychological prompting: An approach to prompt design that incorporates techniques leveraging human cognitive patterns and psychological principles to improve LLM performance. This can include adding phrases that encourage step-by-step thinking, imply urgency, or introduce artificial incentives.

Computer vision: A field of AI that enables computers and systems to derive meaningful information from digital images, videos, and other visual inputs and take actions or make recommendations based on that

information. It involves using technologies such as deep neural networks to analyze visual data.

Content adaptation: The process of modifying or transforming existing learning content to suit different contexts, formats, languages, or learner needs using AI tools. This can include translation, summarization, personalization, and repurposing content into different media (for example, turning a video into a podcast).

Content creation: The process of developing original learning content and materials tailored to specific objectives and organizational needs. This typically involves instructional designers and subject matter experts working together.

Content curation: The process of strategically selecting, organizing, vetting, and contextualizing existing knowledge resources (both internal and external) to create relevant learning experiences. It focuses on leveraging existing assets rather than creating everything from scratch.

Content generation (for AI): The use of GenAI tools—such as LLMs or image generators—to produce different types of learning content, including text, images, audio, and 3D models. This allows for rapid, scalable, and often personalized content production.

Content management system (CMS): A software application or set of related programs that are used to create, manage, store, and deploy digital content. In L&D, a CMS helps organize learning resources, facilitate updates, and manage access, ensuring users can navigate and retrieve materials efficiently.

Content proliferation: The challenge caused by the overwhelming abundance of available learning resources within organizations, which can make it difficult for employees to find relevant materials. This paradox occurs despite significant investment in content as traditional management approaches struggle with the sheer volume.

Context sensitivity: A core characteristic of LLMs in which the model interprets new information based on the preceding text within the prompt. This means the order and structure of information provided will significantly influence the AI model's understanding and subsequent output.

Context window: The memory capacity of an LLM within a single conversation. All preceding prompts and responses within that conversation influence subsequent outputs, shaping the probability of certain words or concepts appearing later.

Conversation flow: The structured design of how a conversation will progress between a user and a conversational AI system. It involves mapping out potential user questions, scripting or prompting appropriate AI responses, and defining pathways for interaction, including escalations.

Conversational AI: Technologies enabling computers to engage in human-like conversations by understanding and responding to natural language inputs. This umbrella term encompasses chatbots, AI assistants, and multimodal generative AI, which use technologies such as machine learning and NLP.

Creation-curation-generation-automation framework: A model presented in this book that describes the evolution of how L&D uses technology for content. It progresses from manual creation to organizing existing resources (curation) to using AI tools (generation) and to streamlining processes (automation).

Data analysis and interpretation: A core competency for the human-machine performance analyst involving sifting through learning and performance data from different systems (such as an LMS, a performance management system, or an AI platform) to uncover meaningful patterns, insights, and root causes. It requires translating raw data into actionable learning strategies.

Data ecosystem: The process of identifying all systems involved in the learning experience (such as an LMS, an HRIS, content libraries, and collaboration tools) and visualizing how learner data, content information, and analytics flow between them. It aims to understand where data resides, how it moves, and where potential blockages or integration opportunities exist.

Data governance: The overall management of the availability, usability, integrity, and security of data that an organization uses. It involves establishing policies, standards, and controls for how to access and use data.

Data integration (for skills gap analysis): The process of collecting data from multiple sources (such as employee profiles, performance reviews, LMSs, project systems, and external market data) into an AI system for a comprehensive skills gap analysis. Effective integration enables a more complete and accurate understanding of organizational skills.

Data minimization: The practice of limiting the collection and retention of personal data to what is directly relevant and necessary to accomplish a specified purpose. This principle helps reduce privacy risks and data management burdens.

Data privacy framework: A set of precise guidelines and protocols governing the collection, storage, use, and protection of learner data within AI-integrated learning systems. This framework addresses the complex privacy implications arising from using AI technologies in L&D.

Data silos: Isolated repositories of data within an organization. If information is stored in one system but not readily accessible by other systems—which often occurs between legacy LMSs and newer LXPs or AI tools—it hinders comprehensive analysis.

Data-driven decision making (in L&D): An approach to L&D strategy and operations in which decisions are based on the analysis of relevant data rather than intuition or tradition alone. AI enhances this process by enabling the synthesis and analysis of large, complex datasets to inform learning initiatives and demonstrate influence.

Deep learning: A subset of machine learning that uses multilayered artificial neural networks (hence “deep”) to learn complex patterns and hierarchical features directly from large amounts of raw data. It’s particularly effective for tasks like image and speech recognition.

Deliberative agents: AI agents that are capable of planning and using strategic reasoning to achieve goals. They can optimize complex processes, such as student learning pathways.

Differential privacy: A system for publicly sharing information about a dataset by describing patterns of groups within the dataset while maintaining the confidentiality of individual data. It involves adding mathematical noise to data to protect personal privacy while enabling aggregate analysis.

Digital coach: An AI-driven component within a learning ecosystem that can guide and support users by providing personalized content recommendations and tailored learning experiences. It bridges the gap between AI analysis (like GenAI identifying content) and the individual learner to adapt recommendations based on user data and interactions.

Dynamic content adaptation: The use of AI tools (such as Synthesia and D-ID) to automatically transform static learning content into more engaging formats, such as video presentations featuring AI avatars. This allows for variations based on language, presentation style, or cultural context while maintaining core consistency.

Emotional resonance (in avatars): The ability of AI avatars, particularly sophisticated ones, to detect subtle cues in human interactions (such as voice patterns) and adjust their presentation style or responses accordingly to create a more engaging and connected experience. This capability stems from neural networks analyzing patterns of interaction.

ETL (extract-transform-load): A type of data integration process that involves extracting data from source systems, transforming it into a compatible format, and loading it into a target system, such as a database or data warehouse.

Explainable AI: AI systems that can explain their decisions or predictions in a way that humans can understand. This process aims to address the “black box” problem and increase trust and accountability.

Extended reality (XR): An umbrella term encompassing immersive technologies like virtual reality (VR), augmented reality (AR), and mixed reality (MR). XR technologies are used in L&D to create engaging, interactive, and often simulated learning environments.

Fairness (in AI): The principle that AI systems should treat individuals and groups equitably and avoid discriminatory outcomes. Defining and measuring fairness is complex because different mathematical and contextual definitions exist and can conflict.

Federated learning: A machine learning approach in which an AI model is trained across multiple decentralized devices or servers holding local data samples, without exchanging the raw data itself. This method enhances privacy by keeping data localized.

Few-shot learning: A prompt design technique that involves a user providing the LLM with a few examples (“shots”) of the desired input and output formats. This helps the model understand the expected structure, style, and content for its response.

“Foot on the accelerator” state: The most advanced stage of AI implementation described in this book. It’s characterized by seamless AI integration across departments, transparent governance, consistent access and training, measurable outcomes linked to AI, and L&D leveraging AI.

Foundation models: Large-scale AI models (such as those from OpenAI, Google, and Anthropic) that serve as the base for various AI applications, including autonomous agents. They provide core capabilities, such as reasoning and language understanding.

Freeform prompts: Direct, natural language instructions or questions without a rigid structure or template a user gives to an LLM. They allow for flexibility and are helpful for brainstorming, open-ended tasks, or initial exploration of an AI model’s capabilities.

Gamification: The application of game-design elements and principles (such as points, badges, leaderboards, and challenges) in nongame contexts (such as learning experiences) to enhance engagement and motivation.

Generative AI (GenAI): A type of AI capable of creating new content—such as text, images, audio, or code—based on the data it was trained on. In L&D, this is used for tasks like course creation, personalization, skills gap analysis, and automating processes.

Generative AI model dimensions (1D, 2D, 3D, and 4D): A framework presented in this book to categorize the capabilities of GenAI applications for L&D. The dimensions progress in complexity and interactivity—from 1D (text) to 2D (multimedia, including images, audio, and video) to 3D (immersive models, such as VR and AR simulations) to 4D (real-time, adaptive systems).

Governance (in AI): The framework of rules, policies, standards, processes, and controls established to guide the development, deployment, and use of AI within an organization. Effective governance is crucial for mitigating the risks associated with AI, including privacy, security, and ethical concerns.

GPT (generative pretrained transformer): An advanced AI model—such as those from OpenAI (GPTs), Google (Gems), or Microsoft (Copilot agents)—that combines a prompt, specific knowledge sources, and access to external tools into a reusable and adaptable unit. They respond to human-initiated interactions, often through conversational prompts, and can perform complex tasks such as generating tailored learning content or analyzing problems.

Headless content creation: An approach to content creation that's particularly relevant in AI-powered workflows, in which an instructional designer interacts with AI models primarily through application programming interface (API) calls rather than traditional software based on a graphical user interface (GUI). This facilitates automation and integration between systems.

Human capital management (HCM) system: Integrated software suites that manage various employee-related processes, including payroll, benefits, performance management, talent acquisition, and, increasingly, L&D functions. Modern HCM systems often incorporate AI for tasks like skills analysis, personalized recommendations, and predictive analytics.

Human–AI collaboration: An approach that combines the strengths of humans and AI agents, enabling them to work together effectively. AI handles tasks like data analysis and routine processes, while humans provide strategic direction, creativity, ethical oversight, and complex problem solving.

Human–AI Task Scale: A seven-level spectrum presented in this book to categorize the varying degrees of collaboration between humans and AI systems on tasks. It ranges from fully manual human work (level 1) to fully autonomous AI (level 7), with several collaborative levels in between.

Human-machine performance analyst: An emerging L&D role that blends traditional L&D expertise with skills in data analysis, AI system management, business acumen, and strategic thinking. This role focuses on using human insight and AI capabilities to analyze performance, predict skills gaps, and optimize learning strategies for business impact.

Human-machine task mapping: A method for analyzing workflows by breaking them down into individual tasks and determining how each can

be most effectively performed through human effort, AI automation, or a combination. It involves identifying the strengths of humans (such as creativity and empathy) and AI (such as speed and data processing) for specific task components.

Hybrid agents: AI agents that combine characteristics of different agent types—typically reactive responsiveness and deliberative reasoning—which enables them to offer immediate support while also developing longer-term strategies.

Hybrid workflow (content creation): A content creation model that blends traditional, manual methods using local desktop tools with cloud-based, AI-driven automation and asset generation. In this approach, AI assists with tasks like generating media or handling repetitive work, while humans retain control over complex interactions, strategic decisions, and quality assurance.

Hyper-personalized learning: An advanced form of personalized learning in which an AI model crafts highly specific individual development paths by analyzing vast data points, including performance metrics, career goals, and organizational needs. These learning journeys adapt dynamically to changing requirements.

Implementation road map: A phased plan outlining the steps, timelines, resources, and milestones for integrating AI into L&D functions or other organizational areas. It typically progresses from foundational activities to expansion and innovation.

Input encoding: The process within an AI system in which input data (such as text, images, or audio) is converted into a numerical or structured format that the AI model can process mathematically and understand. Examples include converting text to tokens or images to pixel grids.

Input-output framework: A structured approach for designing prompts, particularly for specific L&D applications, such as performance consulting or creating training imagery. It involves carefully considering the necessary inputs (such as clean data, prompt structure, and learning science principles), selecting the appropriate AI model, and rigorously validating the outputs against the intended goals.

Intelligent knowledge ecosystem: An organizational environment in which business platforms, AI, and knowledge management systems are integrated to enable the seamless capture, organization, discovery, and contextual delivery of information and learning. It can make knowledge readily available and useful within the flow of work.

Interaction model: Defines how users (learners) will interact with an AI agent, including the modality (such as text, voice, VR, or AR), personalization level, and feedback mechanisms.

Interoperability (in AI): The ability of different AI systems, software applications, and platforms to communicate, exchange data, and use the information exchanged. Learning ecosystems allow integrated experiences across various tools.

Iterative prompt design: The process of continually testing, evaluating, and refining prompts based on an LLM's responses to achieve the desired output quality and meet specific learning objectives. This involves techniques such as self-refinement (asking the AI model to refine the prompt) and incorporating feedback.

Knowledge architect: A strategic role for L&D professionals focused on designing, structuring, and overseeing an organization's knowledge ecosystem. This involves leveraging technologies such as knowledge management systems and AI to ensure that knowledge is effectively captured, organized, connected, and delivered to support learning and performance.

Knowledge graph: A technology that organizes information as a network of entities (such as people, skills, content, and projects) and the relationships between them, mimicking how humans associate concepts. Unlike traditional databases, it focuses on context and connections, allowing for more intelligent search and discovery.

Knowledge management: The organizational discipline and set of processes for systematically capturing, structuring, storing, sharing, and applying knowledge and expertise to improve performance and innovation. Modern knowledge management often involves technology platforms to facilitate these processes.

L&D AI integration: The process of incorporating AI tools and capabilities into the functions and workflows of L&D teams. This includes using AI

for content creation, learner support, process automation, analytics, and enhancing the learning experience.

Large language model (LLM): A type of AI that's trained on massive datasets of text and code and is capable of understanding, interpreting, and generating human-like language. LLMs use NLP techniques to perform tasks (such as translation, summarization, and text generation) based on input prompts.

Learning analytics: The practice of measuring, collecting, analyzing, and reporting data about learners and learning activities to understand and optimize the learning process and its outcomes. AI enhances this by analyzing complex patterns in learner interaction data to provide deeper insights beyond basic metrics, such as completion rates and scores.

Learning ecosystem: The dynamic network of people, content, technology, culture, and strategy that supports learning and development within an organization. It encompasses formal and informal learning experiences enabled by various interconnected tools and platforms.

Learning experience platform (LXP): A type of learning technology software that's used to create personalized, consumer-grade learning experiences by aggregating diverse content, enabling social learning, and often leveraging AI for recommendations and personalization. LXPs typically focus on empowering self-directed learning (a "pull" model).

Learning management system (LMS): A software application that's primarily used for the administration, tracking, reporting, and delivery of formal training programs and compliance courses. LMSs traditionally operate on a top-down assignment (or "push") model.

Learning modalities: The various formats or methods through which learning experiences are delivered. Newer, AI-enabled modalities include interactive simulations, AI tutors, and VR and AR experiences.

Learning record store (LRS): A database specifically designed to store, manage, and retrieve learning activity records in a format that's compliant with the xAPI specification. It acts as a central repository for learning data from multiple sources.

Learning workflow audit: A systematic process of mapping and analyzing the entire life cycle of learning content within an organization, from

creation and review through publishing, delivery, tracking, and maintenance. It involves identifying the tools, systems, roles, bottlenecks, and potential automation opportunities at each stage.

Limited memory agents: AI agents that can retain information from past interactions over a limited period to inform current decisions. They build on reactive agents by incorporating a basic form of memory.

Local workflow (for content creation): The traditional approach to instructional design that involves manually crafting content (such as text, images, audio, video, code, and 3D models) with specialized software tools installed on a local computer. This method offers greater control but can be time consuming, resource intensive, and difficult to scale.

Machine learning: A subset of AI in which systems learn from data to improve their performance on a task without being explicitly programmed for every scenario. This concept underlies many AI capabilities, including pattern recognition and prediction.

Memory (agent component): The component of an AI agent that's responsible for storing and retrieving information from past interactions or experiences. This includes short-term memory for temporary data and long-term memory for learning and refining strategies over time.

Metadata: Data that provides information about other data. In L&D content management, it typically includes details such as titles, descriptions, authors, keywords, learning objectives, target audiences, formats, and usage rights—which help with organizing, searching, and curating content effectively.

Microlearning: An instructional approach for delivering learning content in small, focused bursts or segments that are designed to be consumed quickly and easily. It often targets specific skills or knowledge points.

Mixed reality (MR): A type of extended reality (XR) in which digital and real-world objects co-exist and interact in real time. It blends aspects of both physical and virtual reality.

Model training (AI): The process through which an AI model learns patterns, relationships, and features from a large dataset. This involves feeding the model vast amounts of prepared data and using algorithms

to adjust its internal parameters to perform specific tasks, such as generating text or images.

Multilingual conversational AI: Systems capable of understanding and interacting in multiple languages. They often leverage breakthroughs in NLP to handle syntax, semantics, context, and cultural nuances across different languages.

Multimodal conversational AI: AI systems that integrate and process multiple modes of communication (such as text, voice, and visual inputs) to create more productive and natural interaction experiences.

Multimodal generative AI: Advanced AI systems that are capable of simultaneously processing, understanding, and generating content across multiple data types (modalities), such as text, images, audio, video, and code. These systems mimic human understanding by integrating diverse inputs.

Natural language processing (NLP): A field of AI that's focused on enabling computers to understand, interpret, and generate human language. NLP capabilities include analyzing text for meaning, categorizing information, and generating human-like written content.

Needs hierarchy: A performance consulting framework attributed to Jim and Dana Robinson. It involves analyzing organizational issues across different levels (including business, performance, work environment, and capabilities) to identify performance gaps and their root causes, rather than just symptoms.

Neural networks (artificial): Computing systems inspired by the biological neural networks that constitute animals' brains. They consist of interconnected nodes (or neurons) organized in layers, and are used primarily in deep learning to identify patterns and relationships in data.

Open systems (in AI): GenAI models (such as those developed by Meta) that function with underlying architecture or components that may be more accessible or customizable compared with closed systems. They offer greater flexibility and transparency.

Open-source models: AI models that use publicly available source code (as well as training data or methodologies), which allows anyone to inspect, modify, use, and distribute them. This fosters transparency and collaboration.

“Partial implementation” stage: A stage of AI adoption that describes different departments in an organization using AI tools inconsistently, uneven access, and a mix of approved and unauthorized usage, as well as limited system integration.

Performance architect: An L&D professional who focuses on designing and implementing systems, processes, and experiences aimed at directly improving employee performance within their workflow, often leveraging technology and human–AI collaboration. Their goal is to orchestrate the elements that enable effective job performance.

Performance consulting: A systematic process that’s employed by L&D and performance improvement professionals to analyze performance issues, identify their root causes (which may or may not be related to training), and recommend targeted solutions to achieve the desired business outcomes. It emphasizes business impact over simply delivering training.

Performance support (AI-powered): The delivery of real-time, relevant guidance, information, and microlearning directly within an employee’s workflow and facilitated by AI that understands the task being performed. This contrasts with traditional performance support, which often requires users to actively seek out information from separate resources.

Personalization (in learning): Tailoring learning content, recommendations, or experiences to individual learners based on their specific needs, preferences, roles, skill levels, learning history, or goals. AI plays a significant role in enabling personalization at scale by analyzing data and dynamically adapting content or pathways.

Personalized learning path: An individually tailored sequence of learning resources and experiences (such as courses, articles, projects, and mentorships) that are curated to meet a specific learner’s goals, current skill level, role requirements, and preferences. AI plays a crucial role in generating and adapting these paths dynamically.

Personally identifiable information (PII): Any data that could potentially identify a specific individual, either on its own or when combined with other information. Examples include names, addresses, ID numbers, and, in some contexts, learning performance data.

Pilot programs: Small-scale, controlled experiments that test the feasibility, effectiveness, and potential impact of a new initiative (such as an AI application) before its wider rollout. They enable learning, provide feedback, and support refinement.

Pilot use case: A small-scale, controlled implementation of a new technology or process (such as an AI tool or automation) that's focused on solving a specific, high-impact L&D problem. The goal is to test feasibility, demonstrate value, gather feedback, and build support before considering a larger rollout.

Platform assessment: A structured evaluation of an organization's core business platforms (such as Microsoft 365 or Google Workspace) to determine their suitability for supporting current and future L&D needs, particularly in terms of AI integration, data flow, and scalability. It involves analyzing workflows, data ecosystems, and specific L&D requirements.

Platform evaluation rubric: A standardized scoring tool developed by L&D to objectively evaluate and compare technology platforms based on specific, weighted criteria relevant to learning effectiveness, administration, and strategic goals. Criteria often include content management features, learner experience, data analytics capabilities, and AI readiness.

Predictive analytics: A branch of advanced analytics that uses techniques such as machine learning and statistical modeling on current and historical data to make predictions about future events or behaviors. In L&D, this may involve forecasting course completion rates or identifying learners who are at risk of dropping out.

Predictive learning strategies: L&D strategies that proactively anticipate and prepare for future skills requirements based on AI-driven analysis of business goals, market trends, and internal data. This contrasts with traditional needs assessments that are often backward looking.

Prior knowledge activation: A characteristic of LLMs that involves a prompt triggering specific knowledge domains embedded within the model's extensive training data. Inquiring about a widely known topic typically enables the model to access and use relevant information from its training.

Privacy by design approach: An approach to systems engineering and business practices that proactively embeds privacy considerations into the design and operation of IT systems, networked infrastructure, and business processes from the outset. It means not treating privacy as an add-on, but as a core requirement.

Privacy impact assessment (PIA): A systematic process for evaluating the potential effects of a project, system, or technology on individual privacy and identifying ways to mitigate associated privacy risks. PIAs help ensure compliance with privacy regulations and organizational policies.

Programmatic prompting: A technique involving structuring prompts using code-like syntax, logical operations, variables, and iteration instructions. This enables the production of concrete, consistent, reproducible, and dynamic outputs suitable for complex learning and development tasks.

Prompt design (or prompt engineering): The practice of crafting and refining input instructions (prompts) to effectively guide LLMs toward generating desired outputs for specific applications. It involves understanding LLM capabilities and limitations, structuring requests clearly, and iteratively testing prompts.

Prompt library: A curated, organized, and centralized collection of reusable prompts that consistently and efficiently guide GenAI models for specific L&D tasks. A prompt library addresses challenges such as inconsistent outputs, wasted time, and scaling difficulties associated with the use of ad-hoc prompts.

Prompt workflow (or prompt chain): A sequence of interconnected prompts used within a single AI conversation or tool, in which the output of one prompt serves as the input for the next. This approach allows for progressive refinement and the creation of more complex, cohesive outputs compared with using single, isolated prompts.

Prompt: An instruction or query given to an AI model (like an LLM or GPT) to guide it in generating a specific response or completing a task. Prompts can range from simple, freeform requests to complex, structured instructions.

Prompting: The skill of crafting effective instructions or queries (prompts) to guide Gen AI models, such as LLMs, in producing desired outputs. Effective prompting is crucial for maximizing the utility and accuracy of AI tools.

Reactive agents: The simplest type of AI agent, operating solely based on current precepts and predefined condition-action rules. They do not maintain an internal state or memory of past events.

Real-time skills gap analysis: A dynamic approach, significantly enhanced by AI, that continuously identifies discrepancies between the skills an organization needs and those possessed by its workforce. It uses integrated data and AI analysis to provide up-to-date insights and enable proactive intervention.

Reasoning (in AI agents): The capability of an AI agent to process information, make inferences, solve problems, and make decisions to achieve its goals. This often involves leveraging large LLMs and techniques such as the chain of thought.

Recommendation engine: An AI-powered system that filters information (such as content or products) to suggest items most relevant to a particular user based on their profile, past behavior, and potentially the behavior of similar users. This is commonly seen on consumer platforms and is a key feature of LXPs.

Remote workflow (for content creation): A content creation approach that heavily uses cloud-based, AI-powered tools and services (which are often accessed via APIs) to dynamically generate learning assets (including text, images, audio, video, code, and 3D models) without relying primarily on local software or extensive manual processes. This workflow emphasizes scalability, automation, and personalization and often involves AI agents and headless creation methods.

Representation (in AI): The extent to which AI systems, particularly generative models like LLMs, accurately and fairly reflect the diversity of human experiences, perspectives, identities, and cultures in their outputs. Poor representation can lead to biased or insensitive content.

Retrieval (in AI agents): The process through which an AI agent accesses and retrieves relevant information from data sources—such as databases,

content libraries, knowledge bases, or the internet—to inform its reasoning and decision making.

Role (in prompting): Assigning a specific persona, expertise, or perspective to an LLM within a prompt (for example, “Act as an instructional designer” or “You are a customer service expert”). This guides the model’s tone, word choice, and the depth of response based on the assigned role.

Scalable personalization (for avatars): The ability of AI avatar technology to deliver uniquely tailored learning experiences to many individuals simultaneously. Each learner interacts with an instance of the avatar that’s adapted to their specific cultural background, learning preferences, professional context, or other defined parameters.

SCORM (sharable content object reference model): An established set of technical standards for e-learning software products that ensure online learning content can interoperate with various SCORM-conformant LMSs.

Self-aware agents: A theoretical and advanced type of AI agent that possesses consciousness and self-awareness, which enables it to understand its internal states and potentially evaluate its own performance.

Semantic analysis: A core NLP process used by LLMs that extends beyond grammar (syntax) to comprehend the meaning of words and their relationships within text. It allows the model to interpret context, ambiguity, idioms, and nuances in language.

Shadow AI: The use of AI applications and tools within an organization by employees or departments without the IT department’s explicit approval, knowledge, or oversight. This practice introduces significant security, privacy, compliance, and IP risks.

Skills indexing (in AI): The application of AI to automatically discover, categorize, map, and quantify the skills present within an organization’s workforce. It analyzes diverse data sources (including performance reviews, project contributions, learning records, and internal communications) to create a dynamic, real-time inventory of capabilities.

Skills ecosystem management: The strategic practice of mapping, understanding, developing, and deploying skills across an organization. It’s often facilitated by technologies like knowledge graphs or AI skills indexing

and involves visualizing the current skills landscape, identifying gaps, and aligning development efforts with business objectives.

Skills gap analysis: The process for identifying the discrepancy between the skills required for specific job roles or organizational goals and the current skill levels possessed by the workforce. AI can contribute by inferring skills from various data points and identifying areas that need development.

Skills taxonomy: A structured classification system or library of skills that are relevant to an organization's business, industry, and roles. In the context of AI, this taxonomy is often dynamically generated and updated by skills indexing systems.

Structured prompt: A detailed and formatted prompt that provides specific constraints, context, examples, and clear instructions to guide an AI model toward a precise output format or task. This is in contrast with simpler, open-ended freeform prompts.

Syntactic parsing: A core NLP process that's used by LLMs to break down sentences into their grammatical components (including nouns, verbs, and adjectives). It enables the model to comprehend the grammatical structure of the input text.

Synthetic avatars: AI-generated digital representations of humans that are often used in video creation platforms. These avatars can be based on real people or entirely digitally created. They are capable of delivering scripted content with synchronized lip movements and expressions.

Task (in prompting): Explicitly stating the specific action or outcome required from the LLM within the prompt (for example, "Develop a set of learning objectives" or "Create a video script"). This gives the model a clear purpose and helps avoid overly general or unrelated responses.

Task details (in prompting): Providing specific constraints, guidelines, context, required formats, or key focus areas within the prompt to refine the LLM's output. This ensures the generated content meets specific requirements for relevance, structure, and completeness.

Technology fear cycle: A recurring pattern of societal or organizational responses to the introduction of new, potentially disruptive technologies.

It typically involves stages of initial fear, media amplification, attempts to slow adoption, gradual regulation, and eventual acceptance.

Text-to-speech (TTS): Technology that converts written text into spoken voice output. AI-powered TTS systems can produce increasingly natural-sounding voices in many different languages and accents.

Theory of mind (ToM) agents: AI agents that can understand or infer a user's mental state, including their beliefs, desires, intentions, and emotions. This enables them to tailor their interactions to the perceived cognitive or emotional state of the user.

Tokenization (in AI): The process in NLP in which input text is broken down into smaller units—called tokens—which can be words or subwords. This prepares the text for numerical processing by the LLM.

Tools (for AI agents): External applications, data sources, or APIs that an AI agent can access and use to perform specific actions or gather necessary information to achieve its goals.

TRACI framework: A specific, structured prompt framework that stands for task, role, audience, create, and intent. It provides a comprehensive template for guiding LLM outputs by clearly defining what needs to be done, who is responsible, for whom it's intended, how it should be done, and why it's important.

Traditional AI: Earlier forms of AI typically focused on rule-based systems and predefined logic to solve specific, well-defined problems. These systems operate based on explicitly programmed rules rather than learning from data patterns like machine learning models.

Transparency (in AI): Refers to the clarity and explainability of how an AI system arrives at its decisions or outputs, including the data used and the algorithms employed. For L&D, this means understanding why an AI learning platform recommends certain content or identifies a specific skills gap.

Tree-of-thought (TOT) prompting: An advanced prompting method in which a user encourages an LLM to explore multiple reasoning paths or scenarios simultaneously before reaching a conclusion. This allows the model to consider various possibilities and evaluate alternatives, promoting more robust problem solving.

User experience (UX): This encompasses all aspects of an end-user’s interaction with a company, its services, and its products, specifically referring to the ease of use, efficiency, and satisfaction when interacting with a learning platform.

Versioning (for prompts): The practice of assigning clear version numbers or unique identifiers to different iterations of a prompt within a library. This includes documenting changes and potentially using status labels (such as “draft” or “final”) to track development and maintain history.

Virtual reality (VR): A technology that creates a completely immersive, computer-generated environment in which users can interact, typically by using headsets that block out the real world. Learners can interact within these simulated environments for training or exploration.

“Wait and see” approach: An organizational stance on AI adoption that’s characterized by awareness but hesitation to commit resources or implement solutions. This often due to concerns about security, privacy, or uncertainty. It involves monitoring competitors but taking little internal action.

World model-based agents: AI agents that build and maintain an internal model or representation of their environment. This model enables them to predict how the environment might change based on potential actions, which allows for more sophisticated planning and decision making compared with simpler agents.

xAPI (experience API): A modern e-learning software specification that’s designed to track a wide range of learning experiences, including those that occur outside a traditional LMS (such as mobile learning, simulations, or real-world performance). Data is typically stored in a learning record store (LRS).

Zero-shot prompting: Giving a task or question directly to an LLM without providing any prior examples (“shots”) of the desired output format. The model relies solely on its pre-existing training to understand and respond to the request.

References



- Anthropic. 2024. “Building Effective Agents.” December 19. anthropic.com/research/building-effective-agents.
- Beane, M. 2024. *The Skill Code: How to Save Human Ability in an Age of Intelligent Machines*. Harper Business.
- Borowiec, S. 2016. “Google’s AlphaGo AI Defeats Human in First Game of Go Contest.” *The Guardian*, March 9. theguardian.com/technology/2016/mar/09/google-deepmind-alphago-ai-defeats-human-lee-sedol-first-game-go-contest.
- Cecco, L. 2024. “Air Canada Ordered to Pay Customer Who Was Misled by Airline’s Chatbot.” *The Guardian*, February 16. theguardian.com/world/2024/feb/16/air-canada-chatbot-lawsuit.
- Colon-Hernandez, P., C. Havasi, J. Alonso, M. Huggins, and C. Breazeal. 2021. “Combining Pre-Trained Language Models and Structured Knowledge.” *ArXiv abs/2101.12294* (2021). arxiv.org/pdf/2101.12294.
- Dastin, J. 2018. “Insight—Amazon Scraps Secret AI Recruiting Tool That Showed Bias Against Women.” *Reuters*, October 10. reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G.
- Defelice, R. 2021. “How Long Does It Take to Develop Training? New Question, New Answers.” ATD Blog, January 13. td.org/content/atd-blog/how-long-does-it-take-to-develop-training-new-question-new-answers.
- Glean. 2022. *Hybrid Workplace Habits and Hangups Report*. Glean, February 8. glean.com/resources/guides/hybrid-workplace-habits-hangups.

- Kavlakoglu, E., and R. Vaish. 2020. "NLP vs. NLU vs. NLG: The Differences Between Three Natural Language Processing Concepts." IBM, November 12. ibm.com/think/topics/nlp-vs-nlu-vs-nlg.
- Knowmax. 2024. "FinTech Unicorn Journey of Providing Seamless Customer Service to 2 Million+ Users Through Efficient Knowledge Management." knowmax.ai/case-studies/fintech-unicorn-journey-of-providing-seamless-customer-service-through-knowmax.
- Larson, J., S. Mattu, L. Kirchner, and J. Angwin. 2016. "How We Analyzed the COMPAS Recidivism Algorithm." ProPublica, May 23. propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm.
- Li, F.F. 2021. "Dr. Fei-Fei Li on Human-Centered AI." Video. Greylock, July 13. youtube.com/watch?v=06M_xmHmDfw&t=779s.
- Lin, M. 2023. "Watson Assistant for HR | Part 1." Video. Mike Lin, May 24. youtube.com/watch?v=Ri335l19DRs.
- LinkedIn. 2025. *Workplace Learning Report 2025*. LinkedIn. learning.linkedin.com/resources/workplace-learning-report.
- Najjar, R. 2023. "Redefining Radiology: A Review of Artificial Intelligence Integration in Medical Imaging." *Diagnostics* 13(17): 2760. doi.org/10.3390/diagnostics13172760.
- Pant, V. 2024. "Multilingual Conversational AI: Benefits for Call Centers and Customers." Convin Blog, December 12. convin.ai/blog/multilingual-conversational-ai.
- Ptacek, T.H. (@tqbf). 2022. "I'm sorry, I simply cannot be cynical about a technology that can accomplish this." X, December 1, 2022. 9:05 p.m. x.com/tqbf/status/1598513757805858820.
- Qin, F., and T.A. Kochan. 2020. "The Learning System at IBM: A Case Study." MIT Sloan School of Management, December 3. mitsloan.mit.edu/institute-work-and-employment-research/learning-system-ibm-a-case-study.
- Robinson, D.G., and J.C. Robinson. 1996. *Performance Consulting: Moving Beyond Training*. Berrett-Koehler Publishers.
- Robinson, D.G., and J.C. Robinson. 1998. *Moving From Training to Performance: A Practical Guidebook*. Berrett-Koehler Publishers.

- Robinson, D.G., and J.C. Robinson. 2005. *Strategic Business Partner: Aligning People Strategies With Business Goals*. Berrett-Koehler Publishers.
- Roose, K. 2022. "The Brilliance and Weirdness of ChatGPT." *New York Times*, December 5. [nytimes.com/2022/12/05/technology/chatgpt-ai-twitter.html](https://www.nytimes.com/2022/12/05/technology/chatgpt-ai-twitter.html).
- Salesforce. n.d. "What Are Autonomous Agents? A Complete Guide." salesforce.com/agentforce/autonomous-agents.
- Smail, J. 2018. "Exclusive: Unilever Chatbot to Transform HR Services in 106 Countries." *Employee Benefits*, October 4. employeebenefits.co.uk/benefits-technology/exclusive-unilever-chatbot-to-transform-hr-services-in-106-countries/195656.article.
- Unilever. n.d. "Providing Skills for Life." unilever.com/sustainability/future-of-work/providing-skills-for-life.
- Vorecol. 2024. "Personalizing Training Programs for Employee Engagement and Retention." Vorecol, August 28. vorecol.com/blogs/blog-personalizing-training-programs-for-employee-engagement-and-retention-9181.
- Walsh, C. 2023. "How GE Healthcare-SHS Have Revamped Employee Engagement with AI." *inFeedo.ai*, February 2. infeedo.ai/success-stories/how-ge-healthcare-measured-employee-morale.
- World Economic Forum. 2025. *The Future of Jobs Report 2025*. World Economic Forum, January 5. weforum.org/publications/the-future-of-jobs-report-2025/in-full/3-skills-outlook.