




SWEETRUSH®

The AI Blueprint

Actionable Strategies for L&D Leaders
to Reduce Risk, Add Value, and
Unlock the Benefits of AI at Scale



Introduction

Welcome, learning innovator!

It's 2025, and AI is no longer knocking at our doors: It's already in the building.

Our people are using AI for a wide range of daily tasks, the demand for candidates with AI skills is growing, and **AI is itself an important tool in the recruiting process.**

Meanwhile, L&D and talent leaders have moved beyond outsourcing rote tasks to AI and are now engaging it as a strategic and creative partner. As we continue to dream bigger, our AI ambitions unfold within the larger context of our organization's AI strategy and policies.

Read on to see how a few different stages of adoption might look.



Scenario

1

The organizations that recognized AI's transformational power early in the game have built an AI ecosystem where AI and human knowledge live and thrive together—and are [already seeing the ROI](#) of their efforts.

These successful adopters, characterized as "Frontier Firms" by the latest [Microsoft Work Trends Report](#), "blend machine intelligence with human judgment, building systems that are AI-operated but human-led."

Every employee embodies their organization's AI strategy, thanks to enterprise-wide training on AI guidelines, ethics, and role-based AI skills and tools. In-house AI tools for every business function draw upon a shared knowledge base of specialized terms, branding language, and institutional wisdom.

Finally, the organization has clear processes for human oversight and review to catch any "hallucinations," inaccuracies, and bias in AI outputs—*before* they have a chance to harm the organization's relationships or reputation.

These organizations aren't unicorns: This level of digital transformation is within reach for every organization. We'll share some steps for creating a [solid foundation](#) you can build upon for years to come.



Scenario 2



Other organizations are leveraging AI on a more ad hoc basis:

Certain teams leverage AI tools for specialized use cases; for example, to screen job applicants or field customer inquiries. Other teams might still be looking for a compelling use case.



It's on the individual level that AI is flourishing: Team members are tapping into a virtual smorgasbord of AI tools to help them take meeting notes, summarize lengthy reports, and generate code, images, text, and other content.

We applaud their initiative and curiosity!

Caveat: Without organization-wide AI usage guidelines, security standards, vetted tools, and procedures for human oversight and review, individual AI users are inadvertently opening the door to substantial risk.

These risks range from publishing inaccurate or biased AI-generated content to sharing sensitive information with an AI tool that “learns” from it...and later passes it on to other users.

If this scenario feels a little too close to home, read on to transform your ad hoc approach into a safe, productive AI strategy that empowers your people and extends your (already mighty) in-house wisdom.



Scenario

3



Still other organizations have banned AI outright. Their (understandable!) goal: Protect sensitive information, intellectual property, or even personal safety.

Financial or government institutions, as well as developers of leading-edge technology, might adopt this stance simply because they have so much to protect.

By keeping AI off-limits, these organizations are losing out on the potential benefits of AI—and keeping their people behind the curve. There may be AI power users on staff but, because AI is treated as a liability, they're not allowed to share their superpower.

Does that mean everyone checks their AI tools at the door? Unlikely.

Whether their answer to AI adoption is “not yet” or a hard “never,” these organizations face a hard truth: **AI is already in the building.**

Three out of four employees currently bring their own AI tools to work to meet the ever-mounting demands to do more, better, and faster—and over half are reluctant to admit that they use them. And by using unvetted AI tools, without consistent human oversight, they're risking their organizations' data security, trade secrets, and reputation.

If you're looking to replace resistance with realism, **our guide to responsible AI strategy will help you embed security and responsibility from the ground up.**



The Common Theme: AI

AI is in every organization described above, but **only those in Scenario 1** have met it with a proactive governance, implementation, and usage strategy. Without this foundation, organizations face the following risks:

IP and Data Security

Without clear guidelines, the use of AI tools can inadvertently compromise sensitive organizational and client data.

Quality Assurance (QA) and Inconsistency

When individuals with varying levels of AI expertise use different AI tools to complete job-related tasks, the outputs can be inconsistent, of varying quality, and based on incomplete knowledge and context.

Duplicated Efforts

Without centralized AI knowledge bases that start everyone on the same page, multiple teams or individuals tackle similar AI-driven tasks independently, leading to inefficient, duplicated efforts and outputs of varying quality (see QA, left).

Ethics and Accuracy

Lacking a clear review process, teams and individuals are at risk of using, and even publishing, AI outputs that are inaccurate, biased, or otherwise misaligned with organizational values.

Reputational and Legal Harm

Inaccurate or biased AI outputs that are published or leaked to the public can jeopardize an organization's reputation and/or legal standing.

Missed Opportunities

[Research](#) by Randstad found that most workers view their companies as lagging in AI integration. Meanwhile, [the World Economic Forum](#) reports that a full 50% of global executives cite lack of skills as their biggest barrier to AI adoption. **In short, the L&D opportunity is immense.**



Without guidance or clearance from the top, employees are taking things into their own hands and keeping AI use under wraps...

This approach means missing out on the benefits that come from strategic AI use at scale. It also puts company data at risk in an environment where leaders' #1 concern for the year ahead is cybersecurity and data privacy.



[—Microsoft Work Trend Index Annual Report, 2024](#)



Addressing these risks and capturing the AI opportunity requires a proactive and strategic approach...and a timely one.

We've talked before about how, as L&D leaders, we're in the perfect position to host the conversation about AI, and that's still true! But now that the conversation is already underway—and many of us face ad hoc AI implementation efforts or outright bans—we need to approach it differently.

We hope the tools in this eBook will help you shape a compelling story about AI's potential value and risks with your stakeholders, leaders, and team.

We're optimists, so we believe that AI can be a force for good that helps us connect, create widespread abundance, and generally do more good in the world.

We're on the cusp of an era of transformation on par with the invention of the printing press, the Industrial Revolution, and the dawn of the internet.

Imagine being able to look back at this extraordinary moment in human history and say not only, "*I was there*," but **"I helped make it happen."**

Let's do this.



Warmly,

Tiffany Vojnovski



We are at this moment where we could actually play a difference in shifting the zeitgeist of this country, potentially the world. And it's not out of the question that people in 50, 100, 1,000 years from now, they're going to look back and say, 'Ah, people in that period, they figured out AI. They finally got over to AI, and the AI that we're still using centuries in the future'.



—Peter Leyden, [The Great Progression Explained](#)

The SweetRush AI Team



Danielle Silver
Solution Architect



Adrián Soto
Director of
AI and Immersive
Technologies



Andrés Villalobos
AI Engineer
Strategist



From Sepia to Technicolor:

The Current State of AI in L&D



For years, we've built highly effective L&D programs and offered our learners a rich blend of modalities.

In the age of AI, we're facing a major shift: Our role as L&D leaders now entails curating dynamic content that creates even deeper, more impactful learning.

With AI, we can amplify our impact by offering learners agency, active exploration, and truly personalized experiences.

The benefits are also business-facing: Integrating AI into the powerful work of learning and skilling helps us add even more value. Read on to discover the four key areas.

The AI Value-Add: Four Key Areas



Personalization at Scale:

Adapting learning to individual needs while seamlessly connecting vast organizational knowledge, often through the power of **custom knowledge bases**—a topic we'll explore in the [Human-Centric Knowledge Design for AI](#) section.



Content Curation:

Using AI-driven insights to curate relevant and impactful learning journeys.




Immersion and Agency:

Engaging the entire learner in free-form exploration and interaction with AI-powered, 3D, immersive digital experiences in virtual reality (VR), augmented reality (AR), and mixed reality (MR).



Business Impact and ROI:

Where the other three values lead: measuring learning effectiveness and aligning L&D with your strategic goals.




Beneath it all, AI enables **Deeper Organizational Insights**, uncovering patterns and knowledge gaps that inform crucial business decisions.

Discover how to unlock all of the above with our [human-centered strategic framework for AI in L&D](#). It's designed to help you systematically map your AI ecosystem, amplify human potential, and cultivate ethical, nuanced AI engagement across your organization.

Introduction

Strategic Foundation

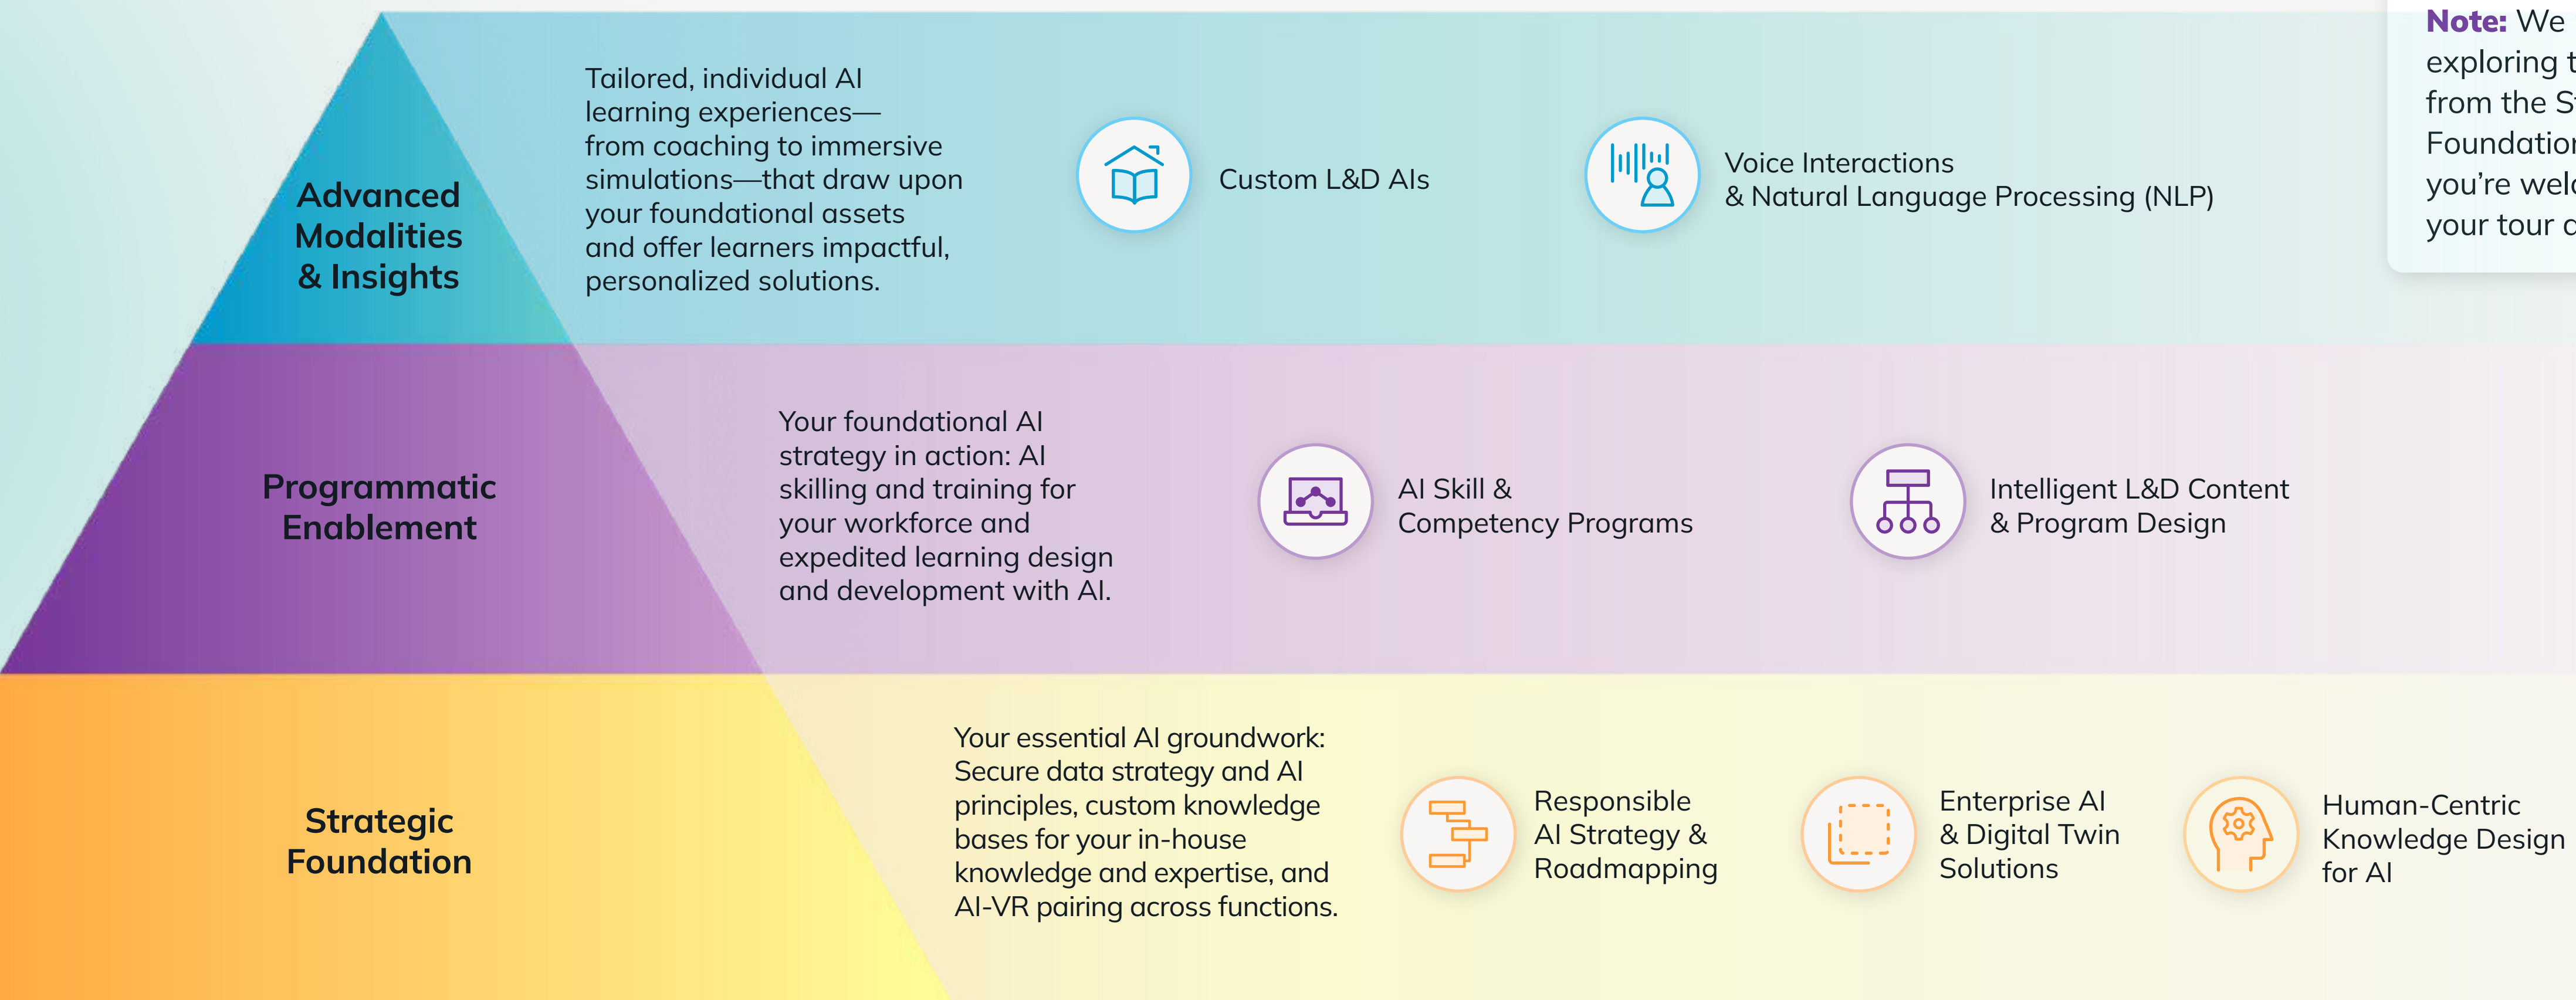
Programmatic Enablement



Conclusion

Advanced Modalities & Insights

AI for L&D: A Human-Centered Strategic Framework



Note: We recommend exploring the pyramid from the Strategic Foundation upward, but you're welcome to start your tour anywhere!

Strategic Foundation

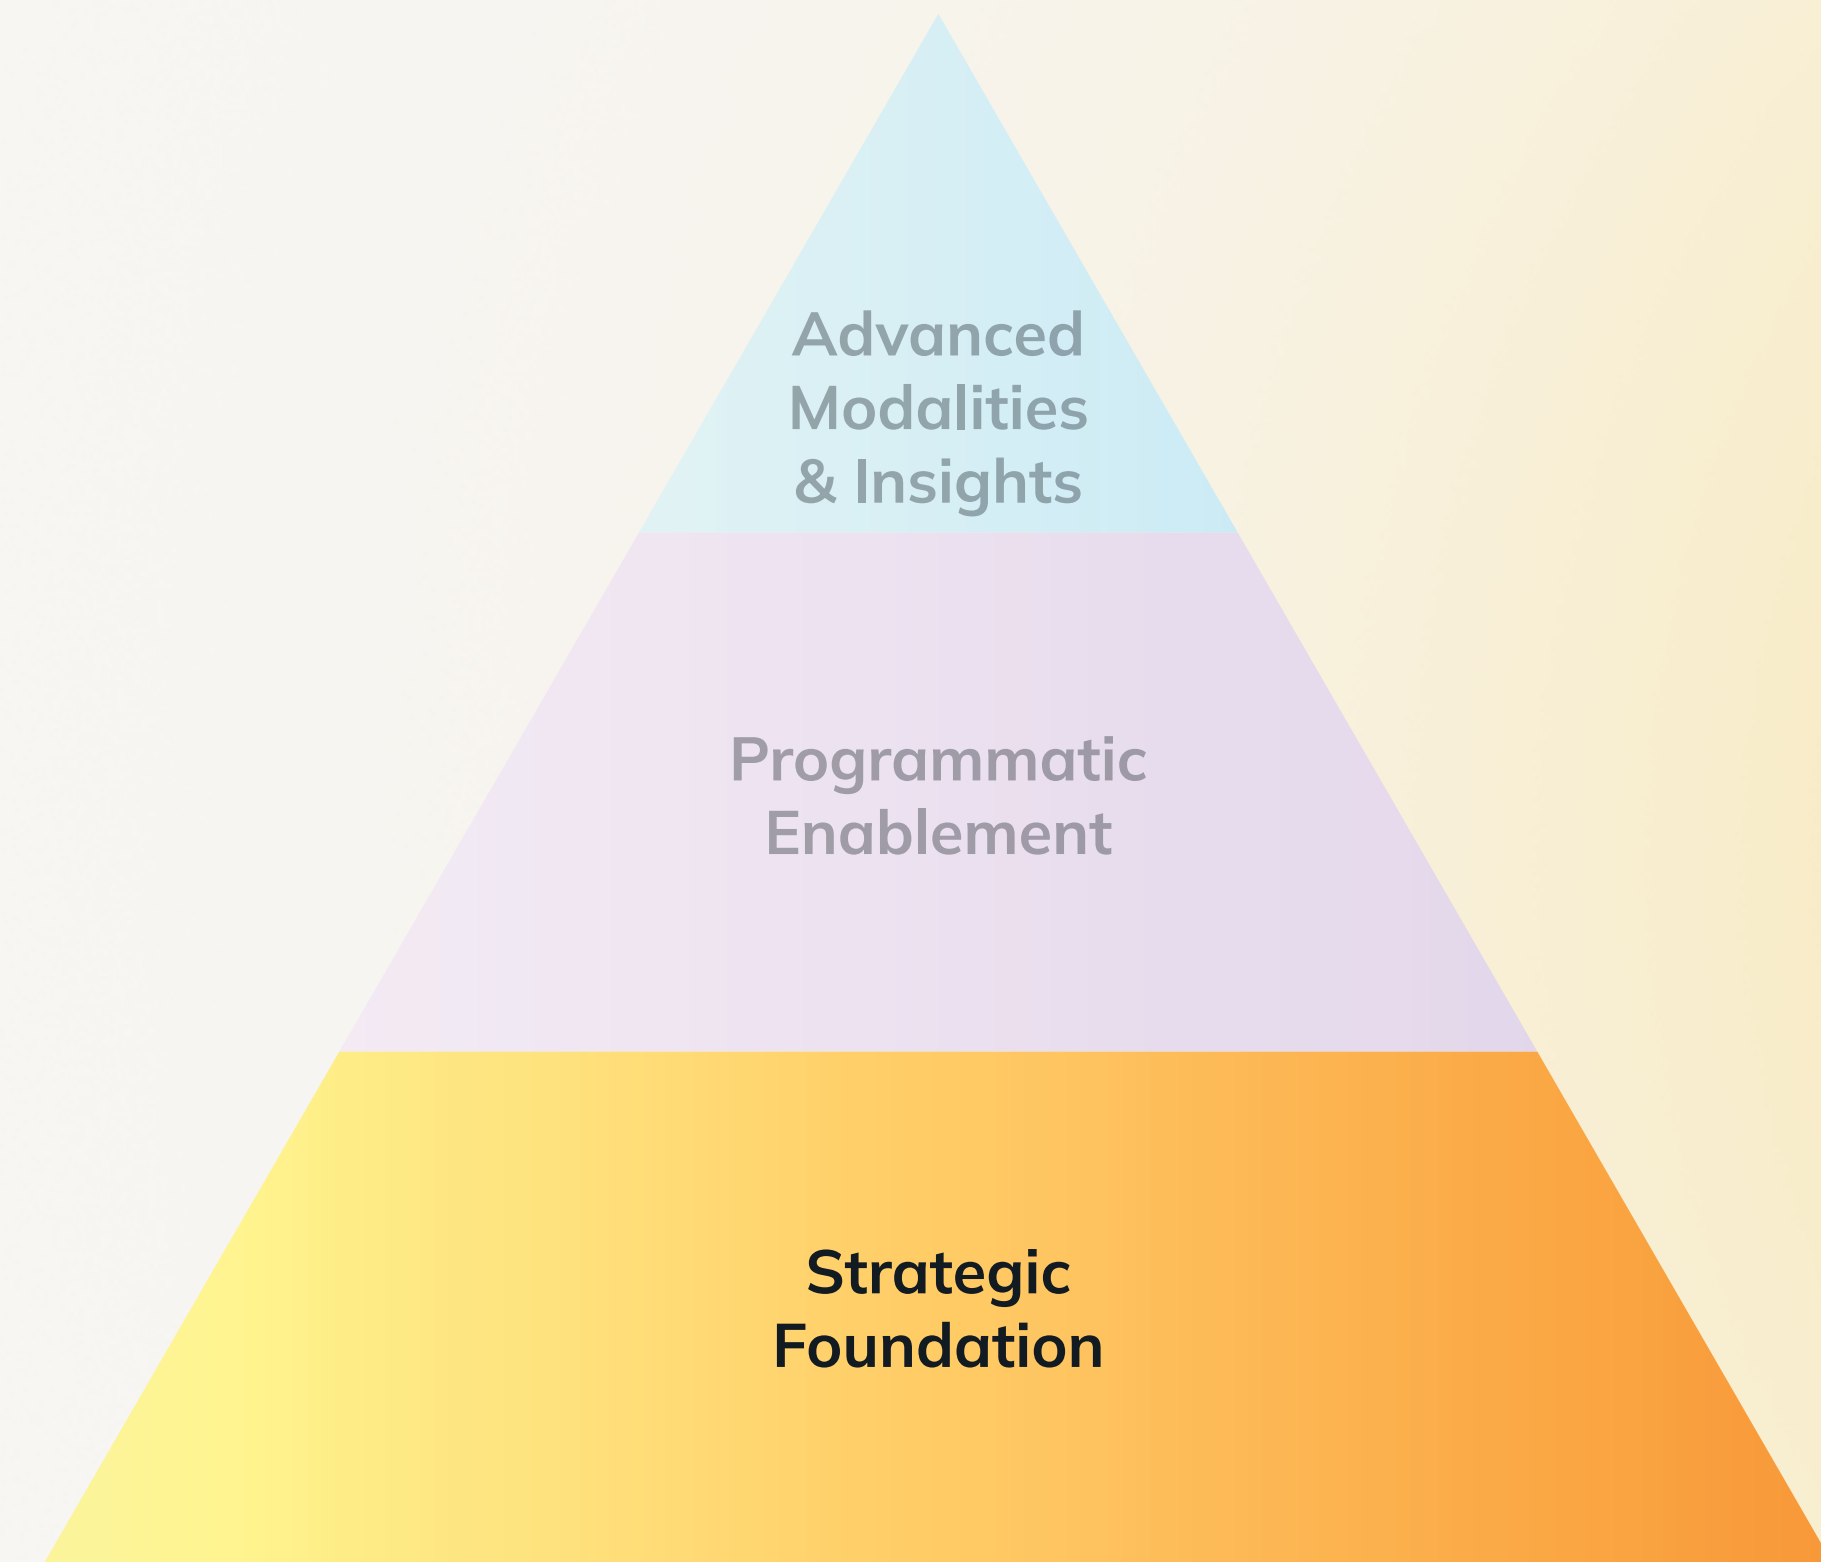
Earlier, we looked in on three types of organizations, each at a different level of AI adoption. In Scenario 1, we described the early adopters, where AI is living, thriving, and actively managed by humans.

What do they have that the others don't? A strategic AI foundation.

Like anything worth building, AI implementation needs an unshakable base: firm enough to protect its inhabitants, but flexible enough to build upon and reinforce over time.

When we *don't* have this, our AI implementation is not only ad hoc; it's incredibly risky. Scenarios 2 and 3 are riddled with unvetted AI tools and unreviewed, untested outputs of varying quality and accuracy.

Let's examine how to mitigate the risks—and also get inspired with some real-life AI implementations by forward-thinking L&D teams that dared to dream bigger.



Responsible AI Strategy and Roadmapping



Embedding responsibility in your AI strategy from the ground up isn't just about compliance—it's a strategic imperative for long-term success.



We also need to consider the human impact of AI: **AI is part of a rapidly evolving skill set that will drastically reshape workers' roles and opportunities in the near future** (See below).




If the world's workforce was made up of 100 people:

59 would need training by 2030.
Of these, employers foresee that:

29 could be upskilled in their current roles.

19 could be upskilled and redeployed elsewhere within their organization.

11 would be unlikely to receive the reskilling or upskilling needed, leaving their employment prospects increasingly at risk.



—World Economic Forum [Future of Jobs Report 2025](#)



On average, workers can expect that

two-fifths (39%)

of their **existing skill sets will be transformed or become outdated** over the 2025–2030 period.



—World Economic Forum [Future of Jobs Report 2025](#)

Organizations have an urgent need to develop ethical strategies and roadmaps to manage AI's human impact. That includes creating [guardrails](#) for responsible use and human oversight. These guidelines and guardrails encourage AI adoption: **Team members who understand AI and its use are more likely to use it and less likely to fear or mistrust it.**

Here's how the process of roadmapping looks from the ground up.



AI Principles


To build a responsible AI strategy, **a focus on values-driven AI principles and needed mental model shifts (MMS) is critical.**

How do we land on these? We leverage information from your internal communication channels, values, existing AI guidelines, and much more (see next page). We also conduct qualitative analyses and empathy interviews with your people.


We then use our findings about what's important to your organization **to co-create unique AI principles** such as **"transparency" or "mindfulness of cost."**



AI Ethics



Transparency



Continuous
Improvement



Cost
Awareness



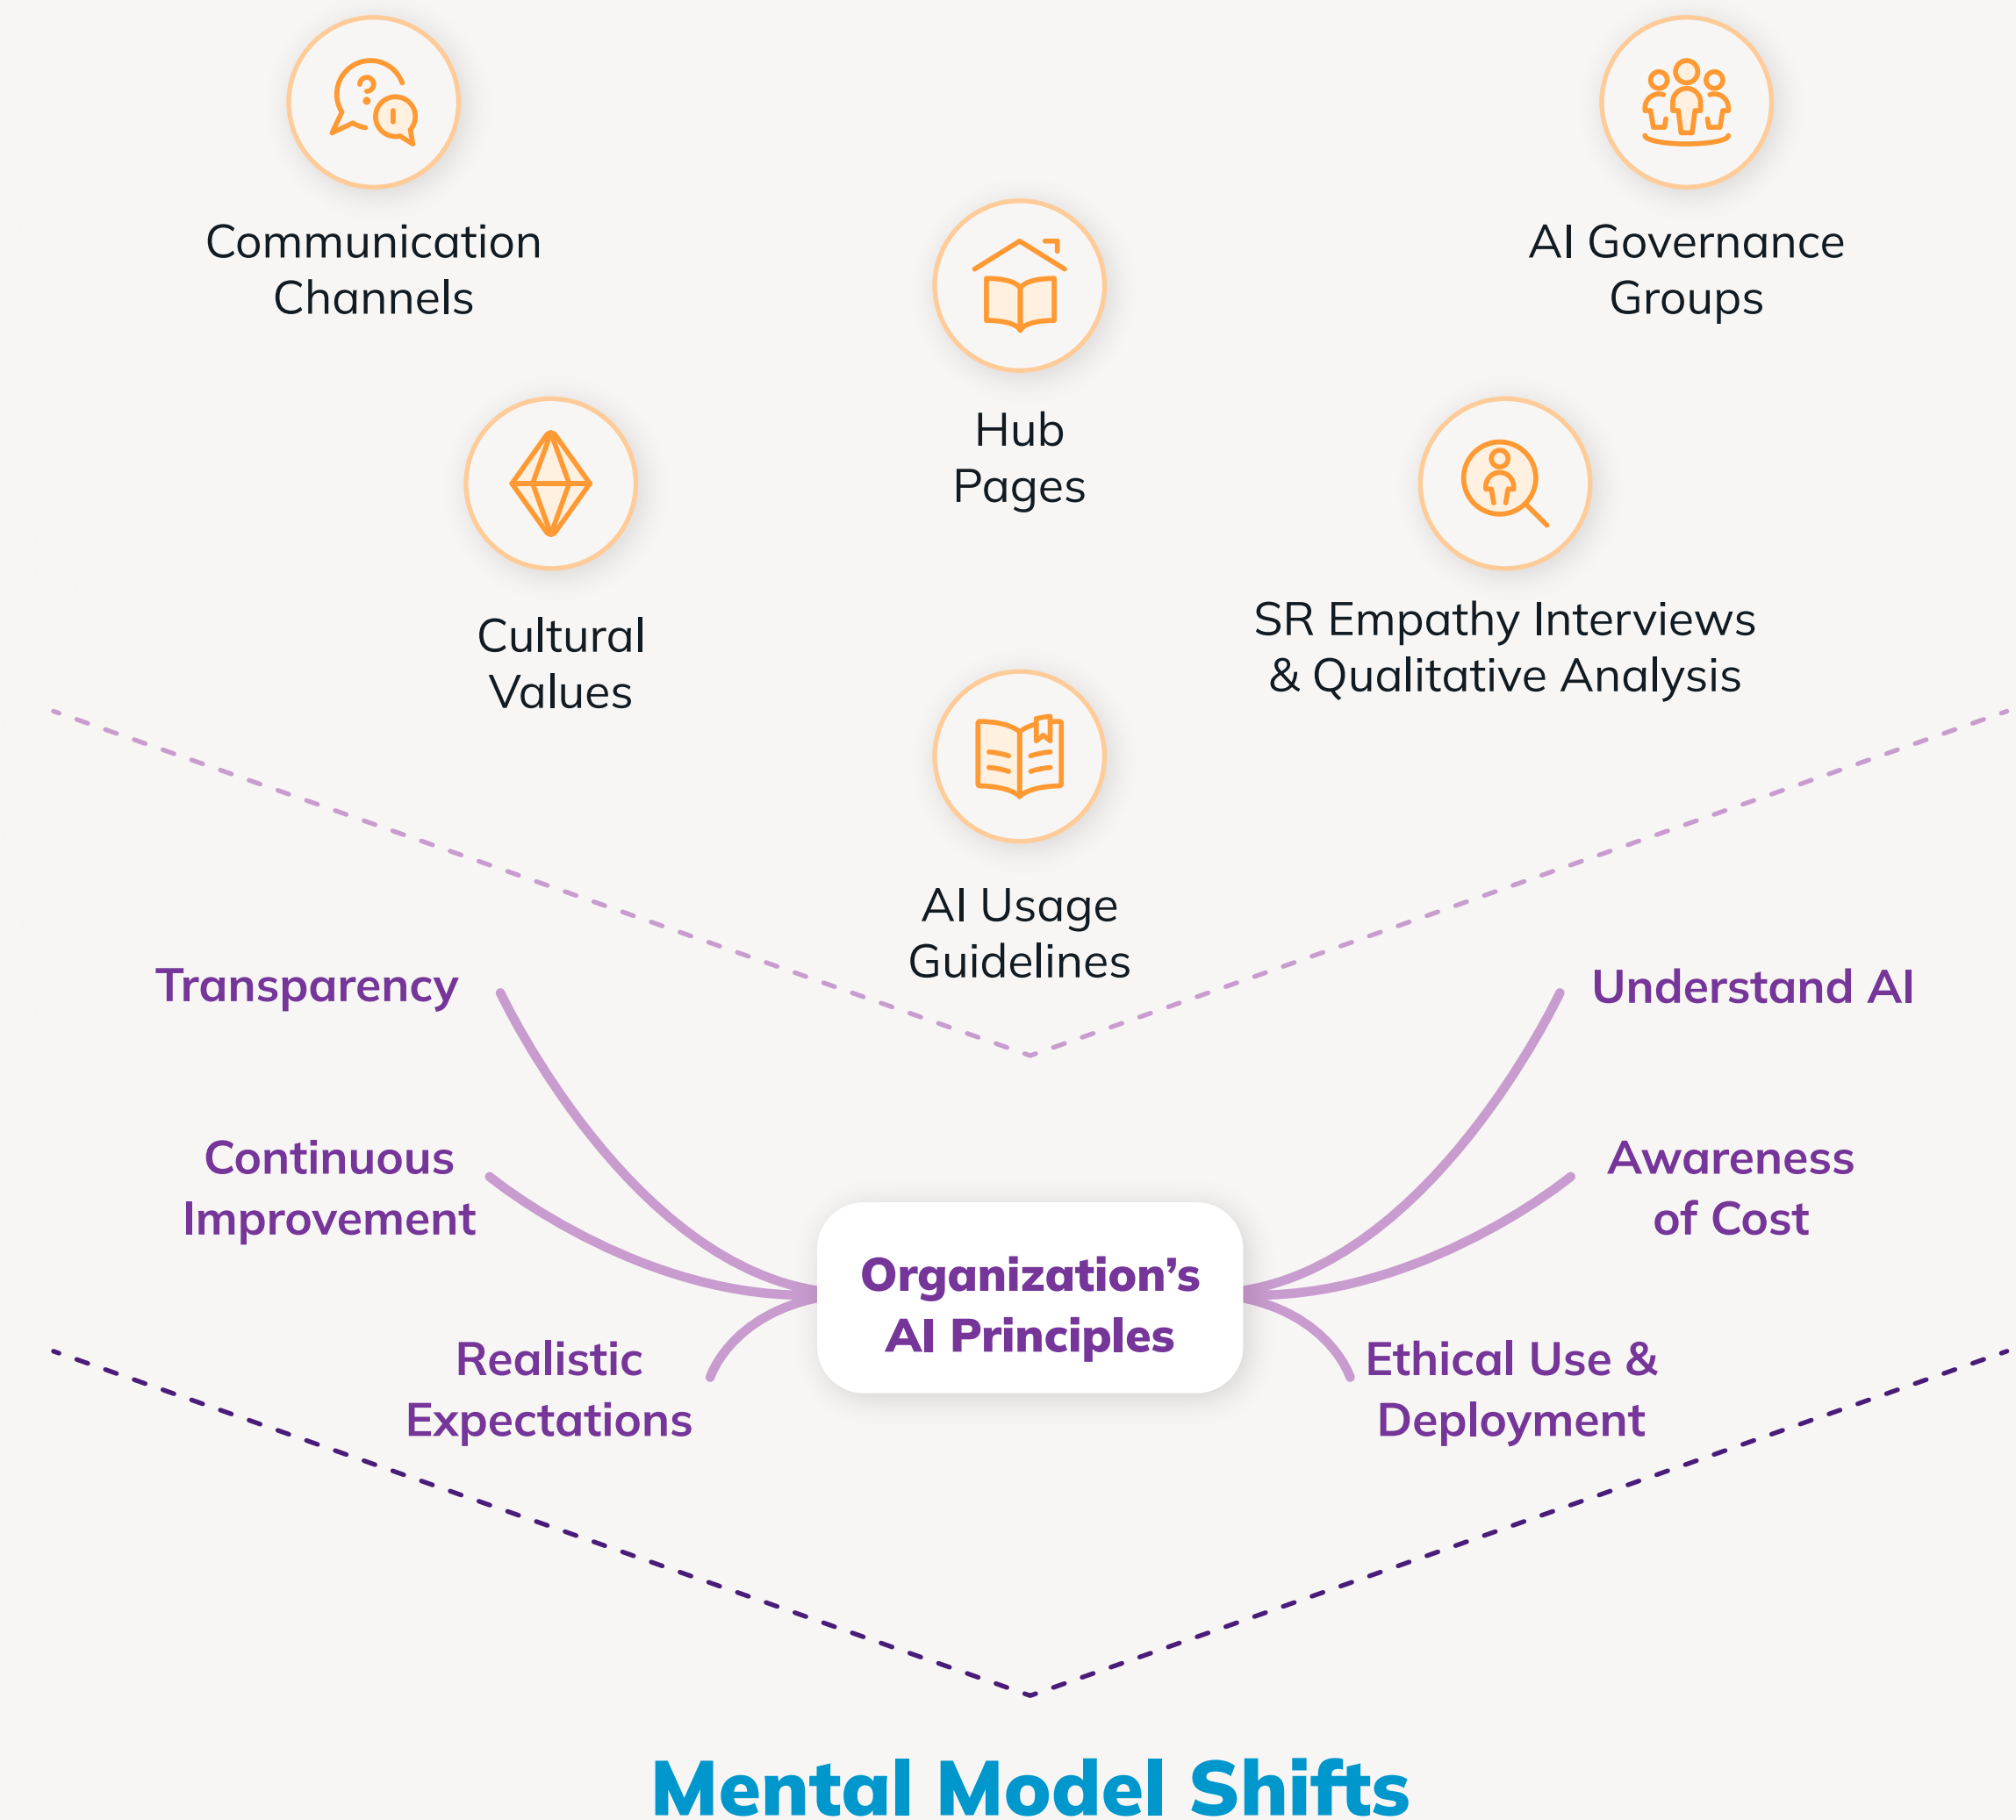
Bias
Mitigation

Mental Model Shifts: Our Sources

Of course, it's not enough just to *tell* people about new tools or policies. To effect real change and foster new behaviors, we must understand and address underlying attitudes, for example, "I'm apprehensive about AI," or "I have no use for AI tools in my role."

Therein lies the shift: We then define the vital mindset and behavioral changes your people will need to embrace AI and contribute to a cohesive AI culture in your organization. By aligning the MMS with your unique AI principles, we ensure your teams enact and embody your principles.

See right for some concrete examples.



Case Study: Mental Model Shifts

We worked with a global construction organization to co-create unique AI principles based on its values and culture. After conducting empathy interviews to *learn* team members’ *current* mindsets about AI, we mapped the new mindsets and behaviors we wanted to cultivate. This approach is how we guarantee that people don’t simply learn about AI but truly **embody** their organization’s vision of responsible and impactful AI use.

From

To

I can’t use AI tools specific to my role because they aren’t available.



I can adopt existing AI tools to save time and add value to my day-to-day work as I keep abreast of new and evolving AI technology relevant to my role.

I’m apprehensive that AI will make my role redundant.



I am confident that AI technology is a tool I can use to enhance my role and support my contributions to my organization.

It’s OK to experiment with AI without boundaries.



It’s great to experiment with AI, but I need to follow my organization’s AI guidelines, policies, and best practices to make sure I’m using AI safely, ethically, and effectively.

AI Meets Humanity:

Safeguarding Human Workflows

Turing tests are another step we use in building a responsible AI strategy. Proposed by Alan Turing in 1950, these tests were intended to answer the question, “Can machines think?”

If a human judge couldn’t tell whether they were conversing with a human or a machine, the machine was said to demonstrate human-level intelligence.

At SweetRush, **we conduct internal Turing tests to evaluate when and how AI amplifies human potential**—and when it doesn’t, or shouldn’t.

Here’s how that works: By reviewing AI outputs, such as language or storytelling, against those of human experts, human judges identify the origin (human or AI), where AI excels, where it needs development, and where human insight is critical.

This process helps us gain clarity on which tasks to automate and identify new opportunities to enhance human expertise.

The Turing Test

During the test, the **interrogator asks questions** to both players to **identify which is human** and **which is a computer**.



AI Integrity: Practicing What We Preach

Quality and accuracy of results are *the* No. 1 concern our client-partners express about AI, and we feel the same.

As L&D professionals, [misinformation and bias](#) are anathema to us—and can cause lasting harm to our organizations.

That's why **we've committed to an ethical alignment strategy that prioritizes accuracy, transparency, and human validation in *all* processes**—without exception.

We practice what we preach by fact-checking all AI outputs and employing AI detection, using tools like Perplexity and GPTZero, respectively, to maintain accuracy and support human oversight.


These practices ensure AI initiatives are inherently built on trust and ethical principles, which are a priori for a robust strategic foundation.



Human-Centric Knowledge Design for AI



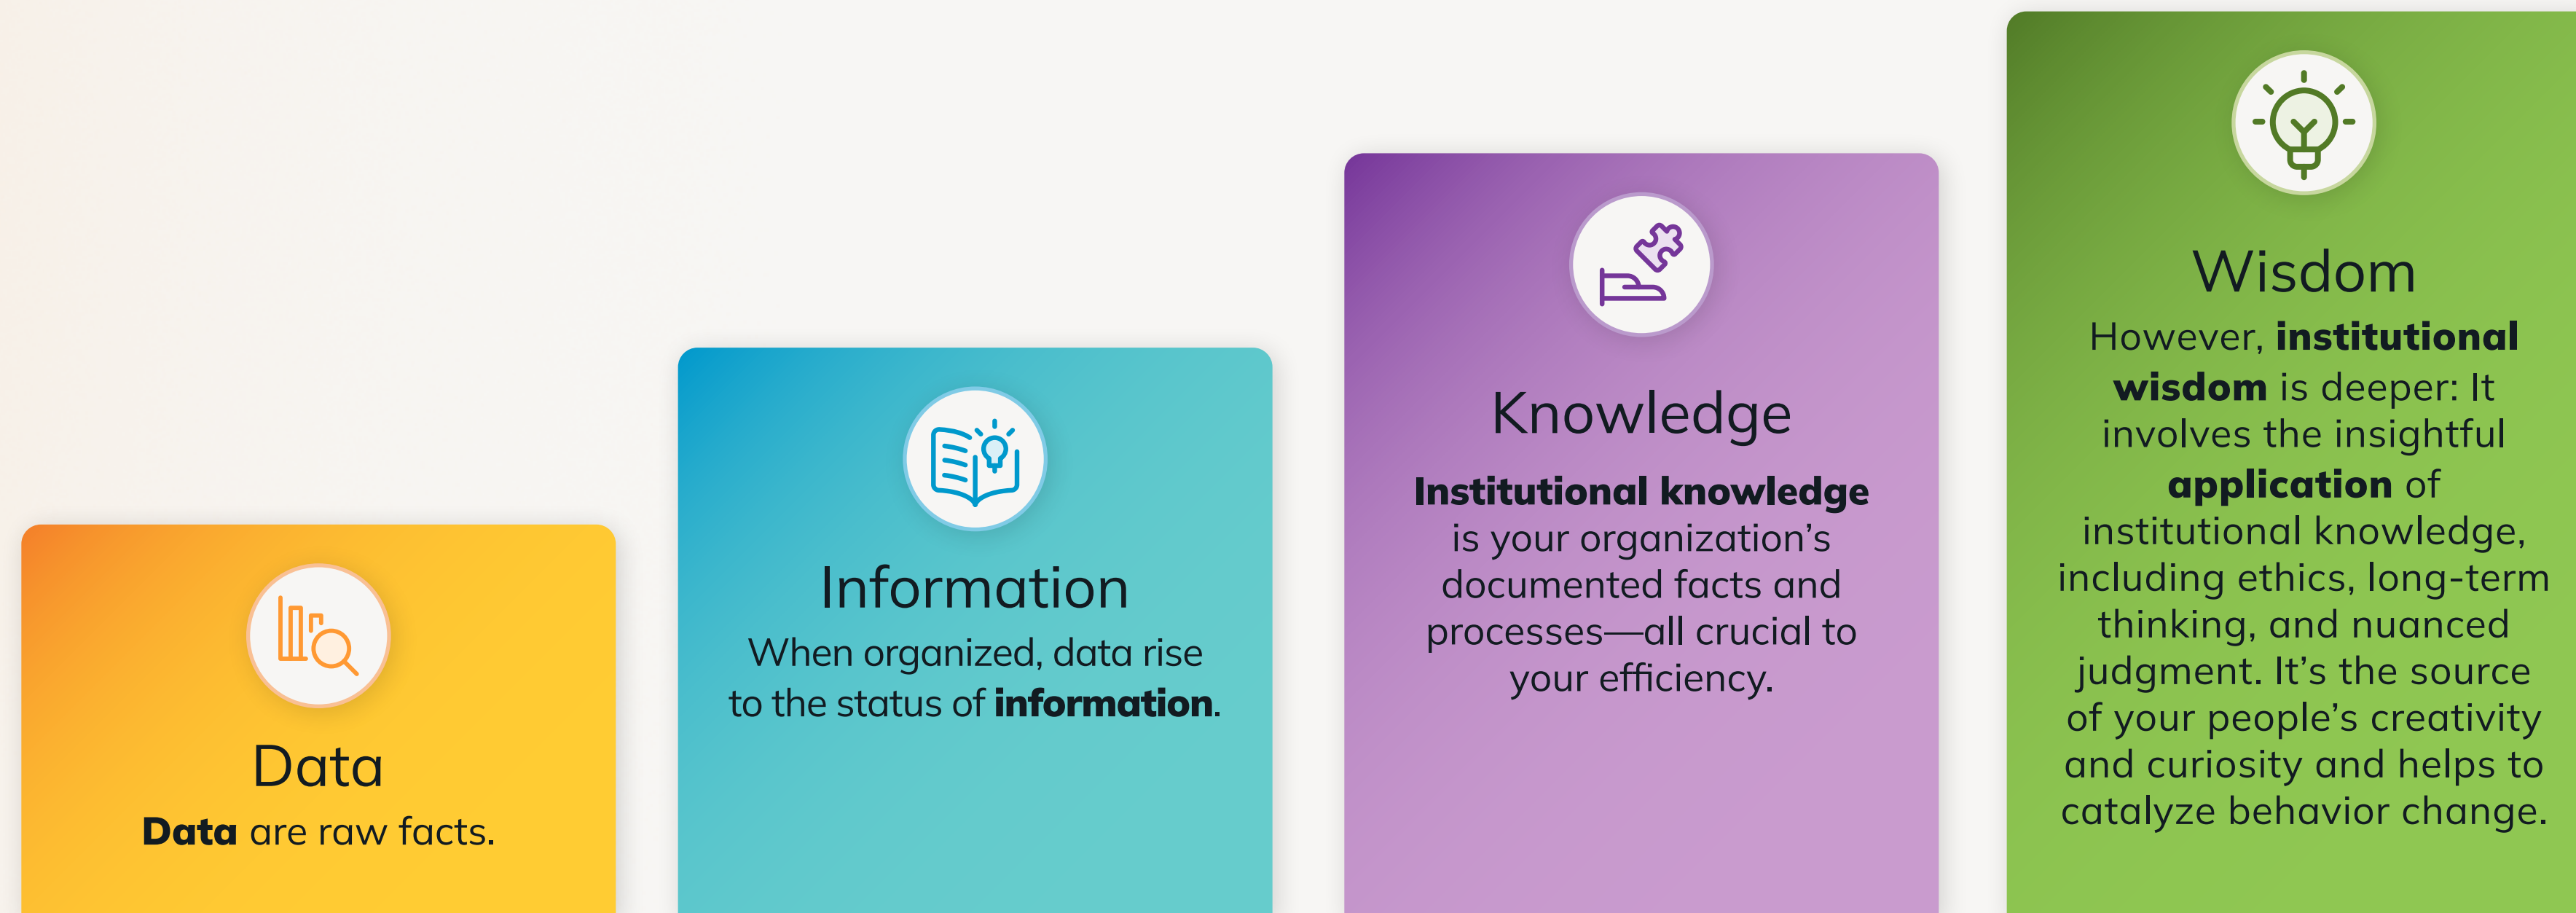
As designers of [transformational learning](#) programs, we're always thinking about how to pair human insight with AI. One way we bring the two together is by creating custom **knowledgebases**, or meaningful, multi-functional repositories of institutional knowledge and wisdom. These are part of the foundation of your entire AI ecosystem.



Cultivating Institutional Wisdom in the AI Era

With over [400 million terabytes](#) of data generated every day, it's challenging to filter out what's relevant, what's accurate, *and* how it all fits together. That's true for us *and* our learners.

But as the graph below shows, we can add value by organizing and indexing information.



Human-Centric Knowledge Design for AI

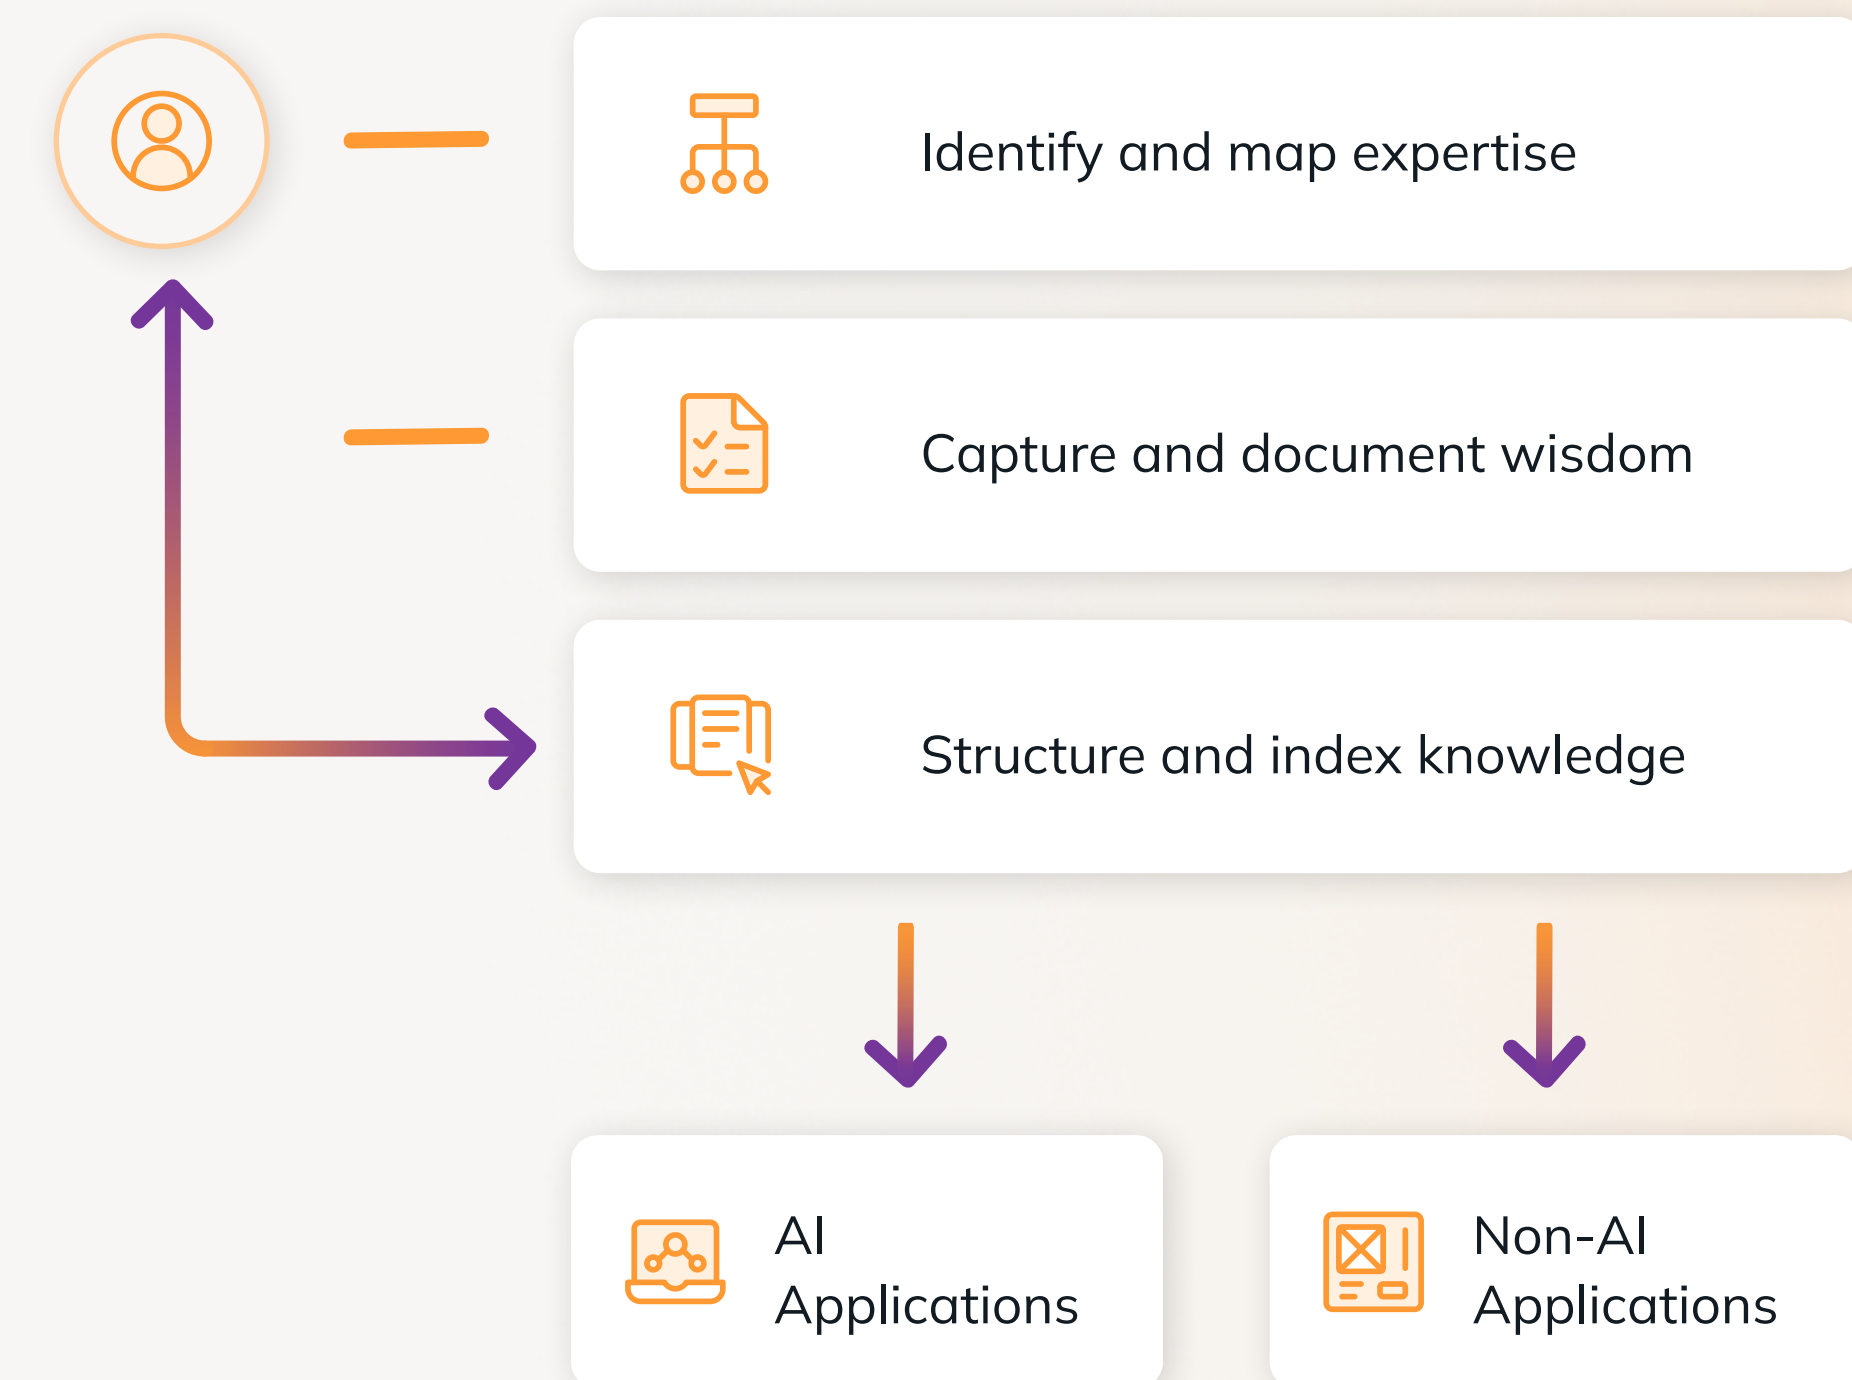
Our knowledge design process (see right) outlines a holistic, human-centered approach to transform wisdom that lives in the heads of in-house experts into powerful, versatile content. Once captured, this content can fuel AI applications, enrich your websites, enhance your learning programs, and much more.

Here's what each stage of the process involves:

- **Identify and map expertise**, or uncover who holds the critical knowledge and insights and where your collective genius truly resides.
- **Capture and document** the wisdom directly from your experts and learners, focusing on the tacit knowledge, invaluable stories, and nuanced advice that forms your institutional “truth.”
- **Structure and index** this knowledge using robust frameworks, making it reusable and accessible.

This collection of knowledge then powers two critical pathways:

- **Powerful AI applications:** For example, AI coaches that provide contextual feedback, or just-in-time knowledge, bringing a high-touch, one-on-one coaching experience to a wider audience and extending the reach of your human experts.
- **Non-AI applications**, including insights for strategic planning, optimizing operations, and enriching HR initiatives using the knowledge you’ve “downloaded” from your most seasoned and successful in-house experts.



Case Study:

Institutional Wisdom Brought to Life with Founder AI

Here's an exciting use of an AI application we've explored with some of our most innovative client-partners—and are currently building in-house for SweetRush.

The goal: immortalize Cofounder Arturo Schwartzberg in a conversational AI.

See right for how the SweetRush 3D team captures Arturo's movements, mannerisms, and speech as he reads from a script using a range of emotional expressions and vocal tones.

To build our Founder AI's knowledge base, we're interviewing Arturo himself. Other sources of Founder wisdom might include interviews with colleagues and family members; analysis of the founder's writings, interviews, and biographies; and wisdom and stories about the founder shared by past and present colleagues.

By “bottling” the founder's unique personality, expressions, and spirit in digital form, the interactive Founder AI experience makes accessing institutional knowledge fun and inspiring.



Enterprise & Digital Twin Solutions

Digital twins are dynamic virtual replicas of places, objects, and even people. They're often built with immersive technologies such as augmented reality (AR), virtual reality (VR), or mixed reality (MR).

Headsets are a great way to visit, but they're not required: Digital twins can be built with WebXR, a web-based immersive technology that can be accessed via a standard browser window on any device.

AI-powered digital twins can offer profound business insights and help you maximize the value of 3D digital assets across your organization, from internal functions like L&D to external-facing functions like recruitment or customer engagement.



Virtual Worlds, Real Impact:

Enterprise AI and Digital Twins

Whether they represent people, places, or things, **digital twins are powering new insights and efficiencies across a wide range of industries and sectors.**

Benefits of digital twins include:

- **Deeper understanding via spatial insights and dynamic feedback**
- **Increased efficiency**
- **Enhanced collaboration**
- **Transformative experiences in learning, marketing, and strategy**

Places

Building, Site,
Infrastructure




People

Conversational
Avatar or
Physiological Model



Objects

CAD 3D Model with
Sensor Readings



Digital Twins: Multi-Platform Strategy

The true power of digital twins lies in their multifunctional value. Below are just a few of the ways they can be leveraged to bring different areas of the organization to life for learners, customers, and key stakeholders alike.

Digital twins can be accessed via a range of devices, including trade show “caves,” laptops or tablets, headsets, and even 3D models embedded in videos.



Case Study: Places

Experiential AI with a Building Information Model (BIM) Demonstrator

For a large global construction client, **we conceptualized a highly visual, interactive BIM Demonstrator** (right) to make abstract AI principles more tangible.

For example, **learners might explore how data impact project sustainability in real-world architecture, engineering, and construction (AEC) applications**, thus leveraging the powerful storytelling and visualization capabilities of this tool.

If you have existing digital twins—perhaps for supply chain or manufacturing simulations—we can import them into our game engines and leverage them for immersive learning.

If you're looking to build for the first time, we can create foundational digital twins, then work together to map other strategic enterprise-wide uses to ensure scalability and customization to meet your organization's needs.



Visual and Interactive Learning of AI Principles

Real-World Application of AI in AEC

Storytelling and Visualization


Scalability and Customization for Diverse Learning Needs

Case Study: People

SweetRush's Founder AI

Founder AI, we meet again! In the previous section on human-centric knowledge design, we discussed this exciting application of institutional knowledge. (You can find a more detailed description [here](#).)

This conversational AI experience channels your organization's founder and can be used to serve numerous business units and functions. Imagine the immediacy, engagement, and emotional realism a Founder AI experience would bring to L&D programs, recruitment events, and marketing demonstrations.



Case Study: Objects

VR from the Heart: A Digital (Almost) Twin for ECG Mastery

Here's a learning experience (literally) from the heart: We created this model of the human heart in partnership with a medical nonprofit organization.

Via dynamic visualizations of electrical pathways and interactive navigation, this experience challenges learners to deepen their understanding of cardiac anatomy and practice interpreting electrocardiograms (ECGs).

Though a true digital twin would integrate with physical medical equipment or provide a real-time, two-way connection to a physical heart, this digital almost-twin still profoundly immerses learners.

Imagine how a foundational model like this could offer multifunctional value beyond skills practice: for example, in marketing, recruiting, or industry partnerships.



Ready for the Next Level?



We hope you're feeling inspired to break ground on—or expand—your strategic foundation. Read on to travel to the middle level of the pyramid and discover how to translate your foundational AI strategy and roadmap into action.



Programmatic Enablement

Now that we've talked about creating a foundational AI strategy, we'll dive into how it translates into actionable capabilities across your workforce.

To avoid a chaotic, ad hoc, and—above all—*risky* AI implementation like the one in Scenario 2, **your entire workforce needs to be onboarded to your AI strategy and vetted collection of AI tools.**

It's also great for your people: Research shows that folks who [understand AI](#) are more likely to engage with AI tools and less likely to view AI with fear or mistrust. And given that [82% of leaders](#) planning to expand the capacity of their workforce with digital labor, there's no time like the present.

Sharing your AI strategy, [guardrails](#), and policies—along with the WIIFMs of increased [efficiency and productivity](#)—with your people helps them understand AI and approach it as a long-term partner.

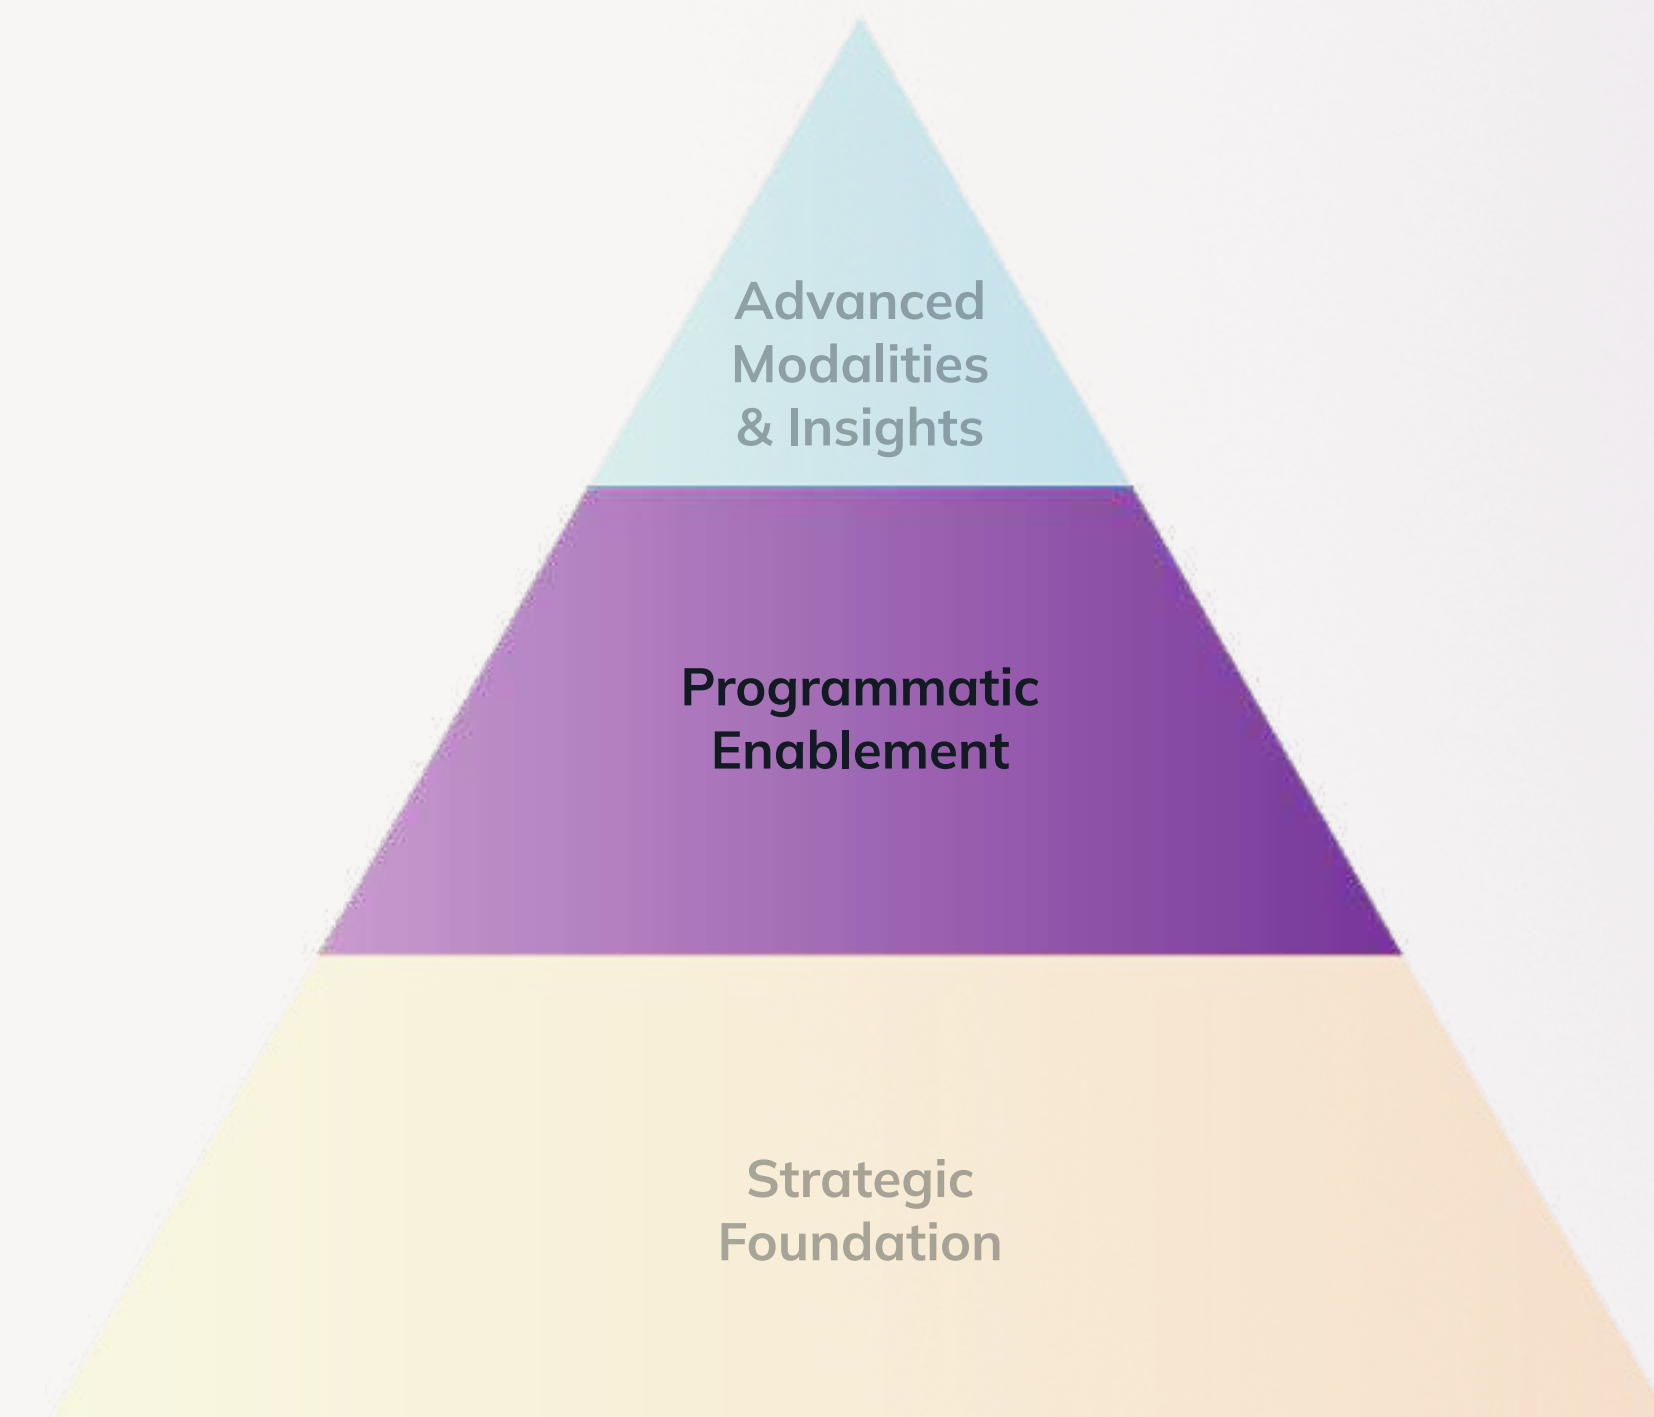
And now, a zoom into the L&D value-add.

Naturally, L&D leaders are motivated by those same WIIFMs: Lack of time is our top obstacle in creating new content; meanwhile, the demand for custom learning is only ramping up.

We'll explore how *you* can partner with AI to expedite the production of L&D and programs; curate content and internal wisdom and recommend personalized

learning paths; to help your team scale high-value learning experiences to your entire learner audience—all the while tracking progress, business impact, and ROI.

With our solid foundation in place, it's time to start building skyward.



AI Skill & Competency Programs

Why is workforce AI skilling so important? The numbers paint a vivid picture of the need:

76% [of working professionals](#) say they need AI skills to remain competitive in the job market.

57% reported feeling [some or a great deal of impact](#) from AI on their jobs, with overall experiences skewing positive.

56% report that they do not feel [prepared](#) to use AI at work.

55% reported that their employer provides [no clear guidance](#) on the use of AI tools.

41% feel they have a [good understanding](#) of AI.

35% are using AI at work, but only [only 31%](#) report having received workplace AI training.

16% [of respondents](#) reported that they had access to employer- or school-provided AI tools.

As the data demonstrate, frequent AI usage among team members does not equal organizational or individual AI readiness.

Without structured, practical training and clear guidance on usage, AI adoption risks being superficial, spotty, and incredibly risky, just as it was in Scenarios 2 and 3.

Real AI capability involves shifting how people think, act, and integrate AI into their daily workflows. It's about achieving genuine, lasting behavioral change and tool mastery, ensuring that team members not only understand AI but also **embody** responsible and effective AI use.



What if AI isn't just a butler, automating the mundane, but a **muse and **co-creator** that sparks innovation and helps us achieve what was once unimaginable?**



—The SweetRush AI Team



AI Training Pathways

Every organization's AI skill needs are unique. In general, though, training pathways tend to fall into three themes:

Theme 01

The **AI Fundamentals and Responsible Innovation** pathway addresses the foundational understanding of AI, including:

- Defining and embodying your unique AI principles
- Fostering psychological safety
- Mitigating bias
- Managing data securely and responsibly
- Ensuring copyright traceability

Theme 02

The **Practical AI for Enhanced Productivity** pathway focuses on immediate, tangible skills to equip your team members to leverage AI for daily tasks. This pathway leverages insights from your **custom knowledge bases that reflect nuanced institutional knowledge and wisdom**—so that this valuable source of human expertise doesn't get lost in your digital transformation. Skills include:

- Using vetted AI tools and platforms
- Prompt engineering
- Human-AI collaboration

Theme 03

The **Strategic AI Integration and Transformation** pathway is designed for organizational leaders and internal change agents. This path aims at forging a cohesive, strategic AI roadmap across the organization to transform operations and culture, and ensure initiatives are strategically aligned, continually calibrated, and evaluated to achieve measurable impact. This pathway delves into:

- Researching and developing new AI initiatives, projects, and tools
- Managing AI projects
- Integrating AI into specific industry contexts
- Mastering change management for AI adoption

As we see in all three pathways, **building AI capability involves far more than just tool training**. Each is built upon the values-driven AI principles we established together in the AI strategy mapping phase and the mental model shifts we identified. (Return to the foundation of our pyramid for more on [values-driven AI strategy](#) and [mental model shifts](#), including a [case study](#).)



AI Skill and Competency Programs: Building Upon the Foundational AI Strategy

Here's a quick recap of how we work with you to create these fundamental building blocks for your AI skilling program:

We begin by gathering information through internal communication channels, qualitative analyses, and empathy interviews with your people to co-create unique AI principles specific to your culture—such as “continuous improvement” or “ethical use and deployment.” These principles address common AI fears and provide a North Star for your teams.

They also form the foundation for defining mental model shifts (MMS), or the critical attitudes and behavioral changes your AI training program needs to foster to build a truly cohesive AI culture.

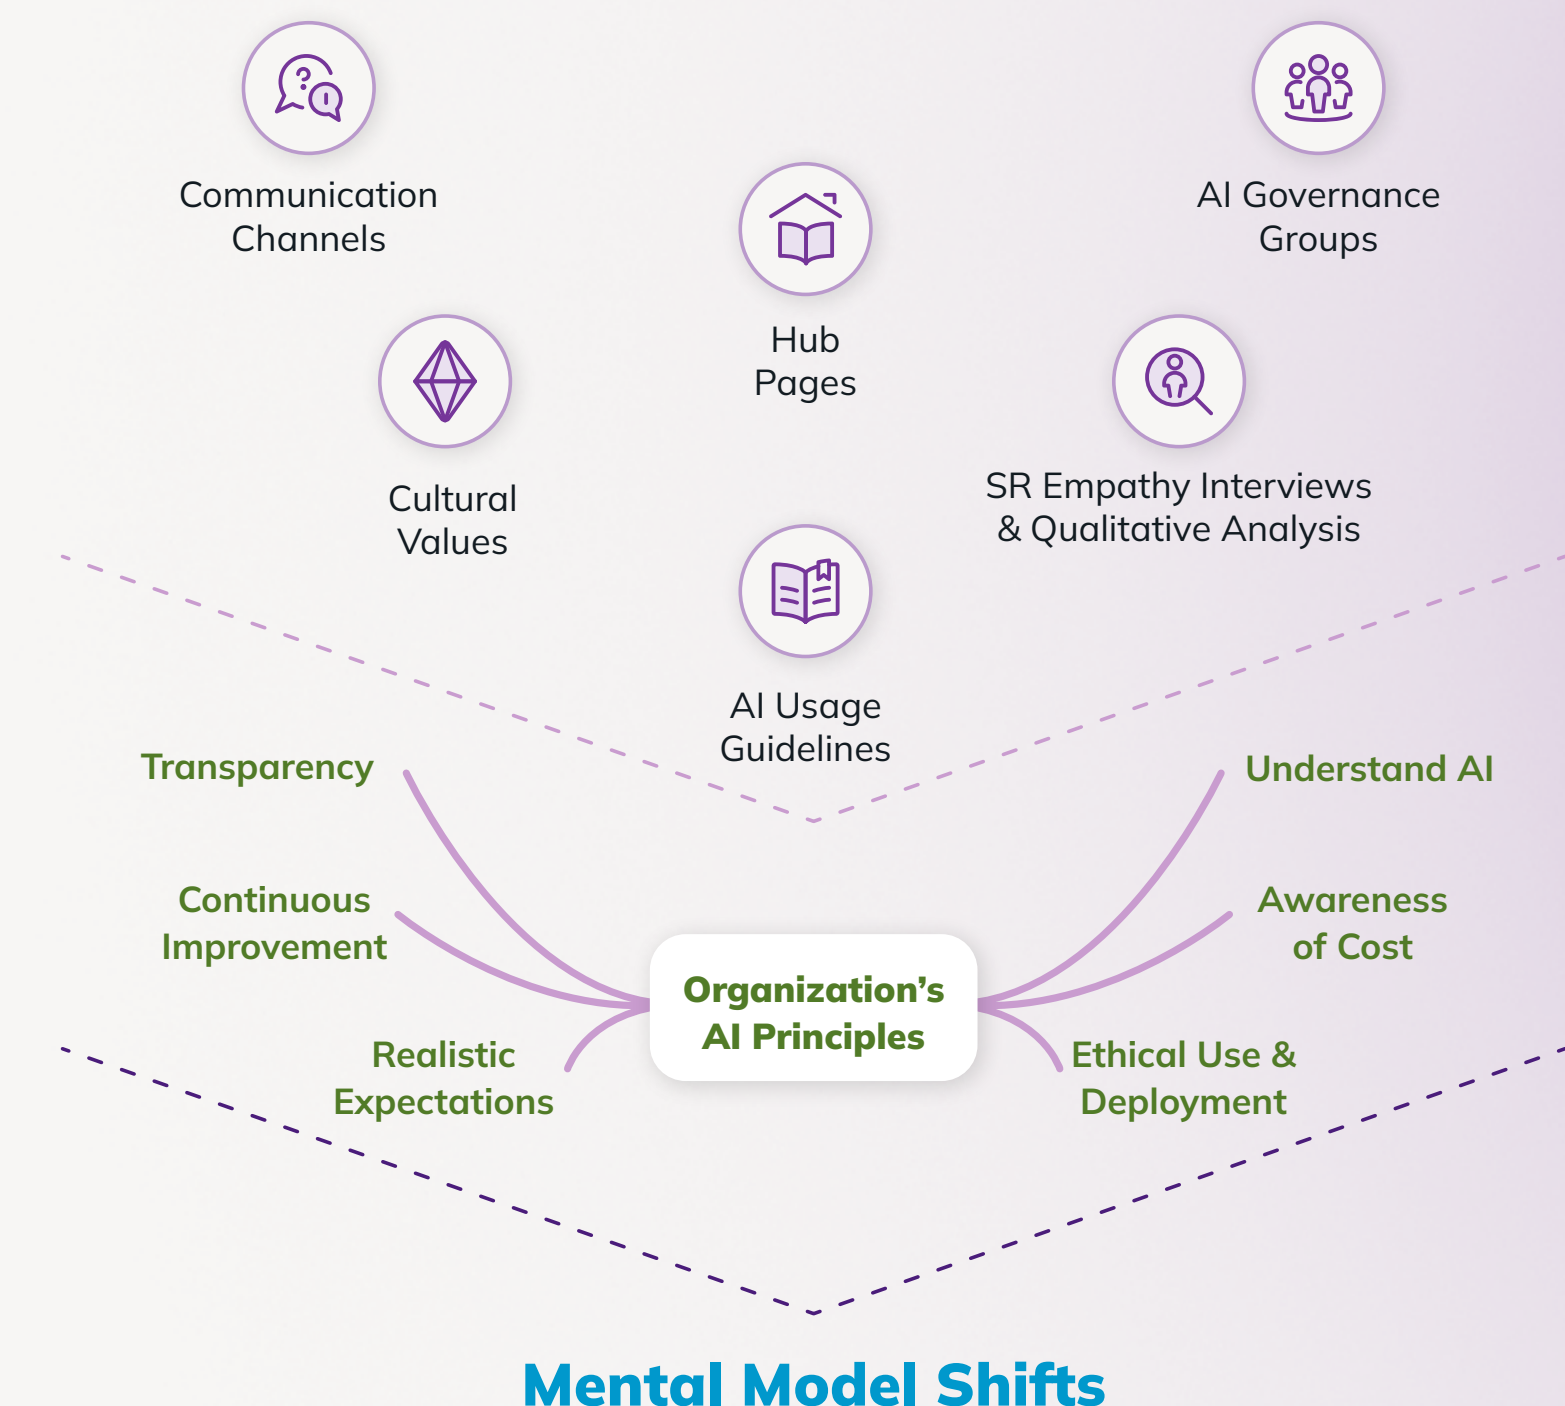
For example, your goal might be to shift your people's mindset from “It's OK to experiment with AI without boundaries” to “It's great to experiment with AI, but I need to follow my organization's AI guidelines, policies, and best practices to make sure I'm using AI safely, ethically, and effectively.”

(See [examples](#) of the MMS we defined for one innovative client-partner's comprehensive AI training program.)

By aligning your programs with these custom MMS, we ensure that your program drives sustained behavioral change among your people and builds the desired AI competencies in every area of your organization.

Curious about how an AI skill and competency program might look in action?

We've got a few examples to share!



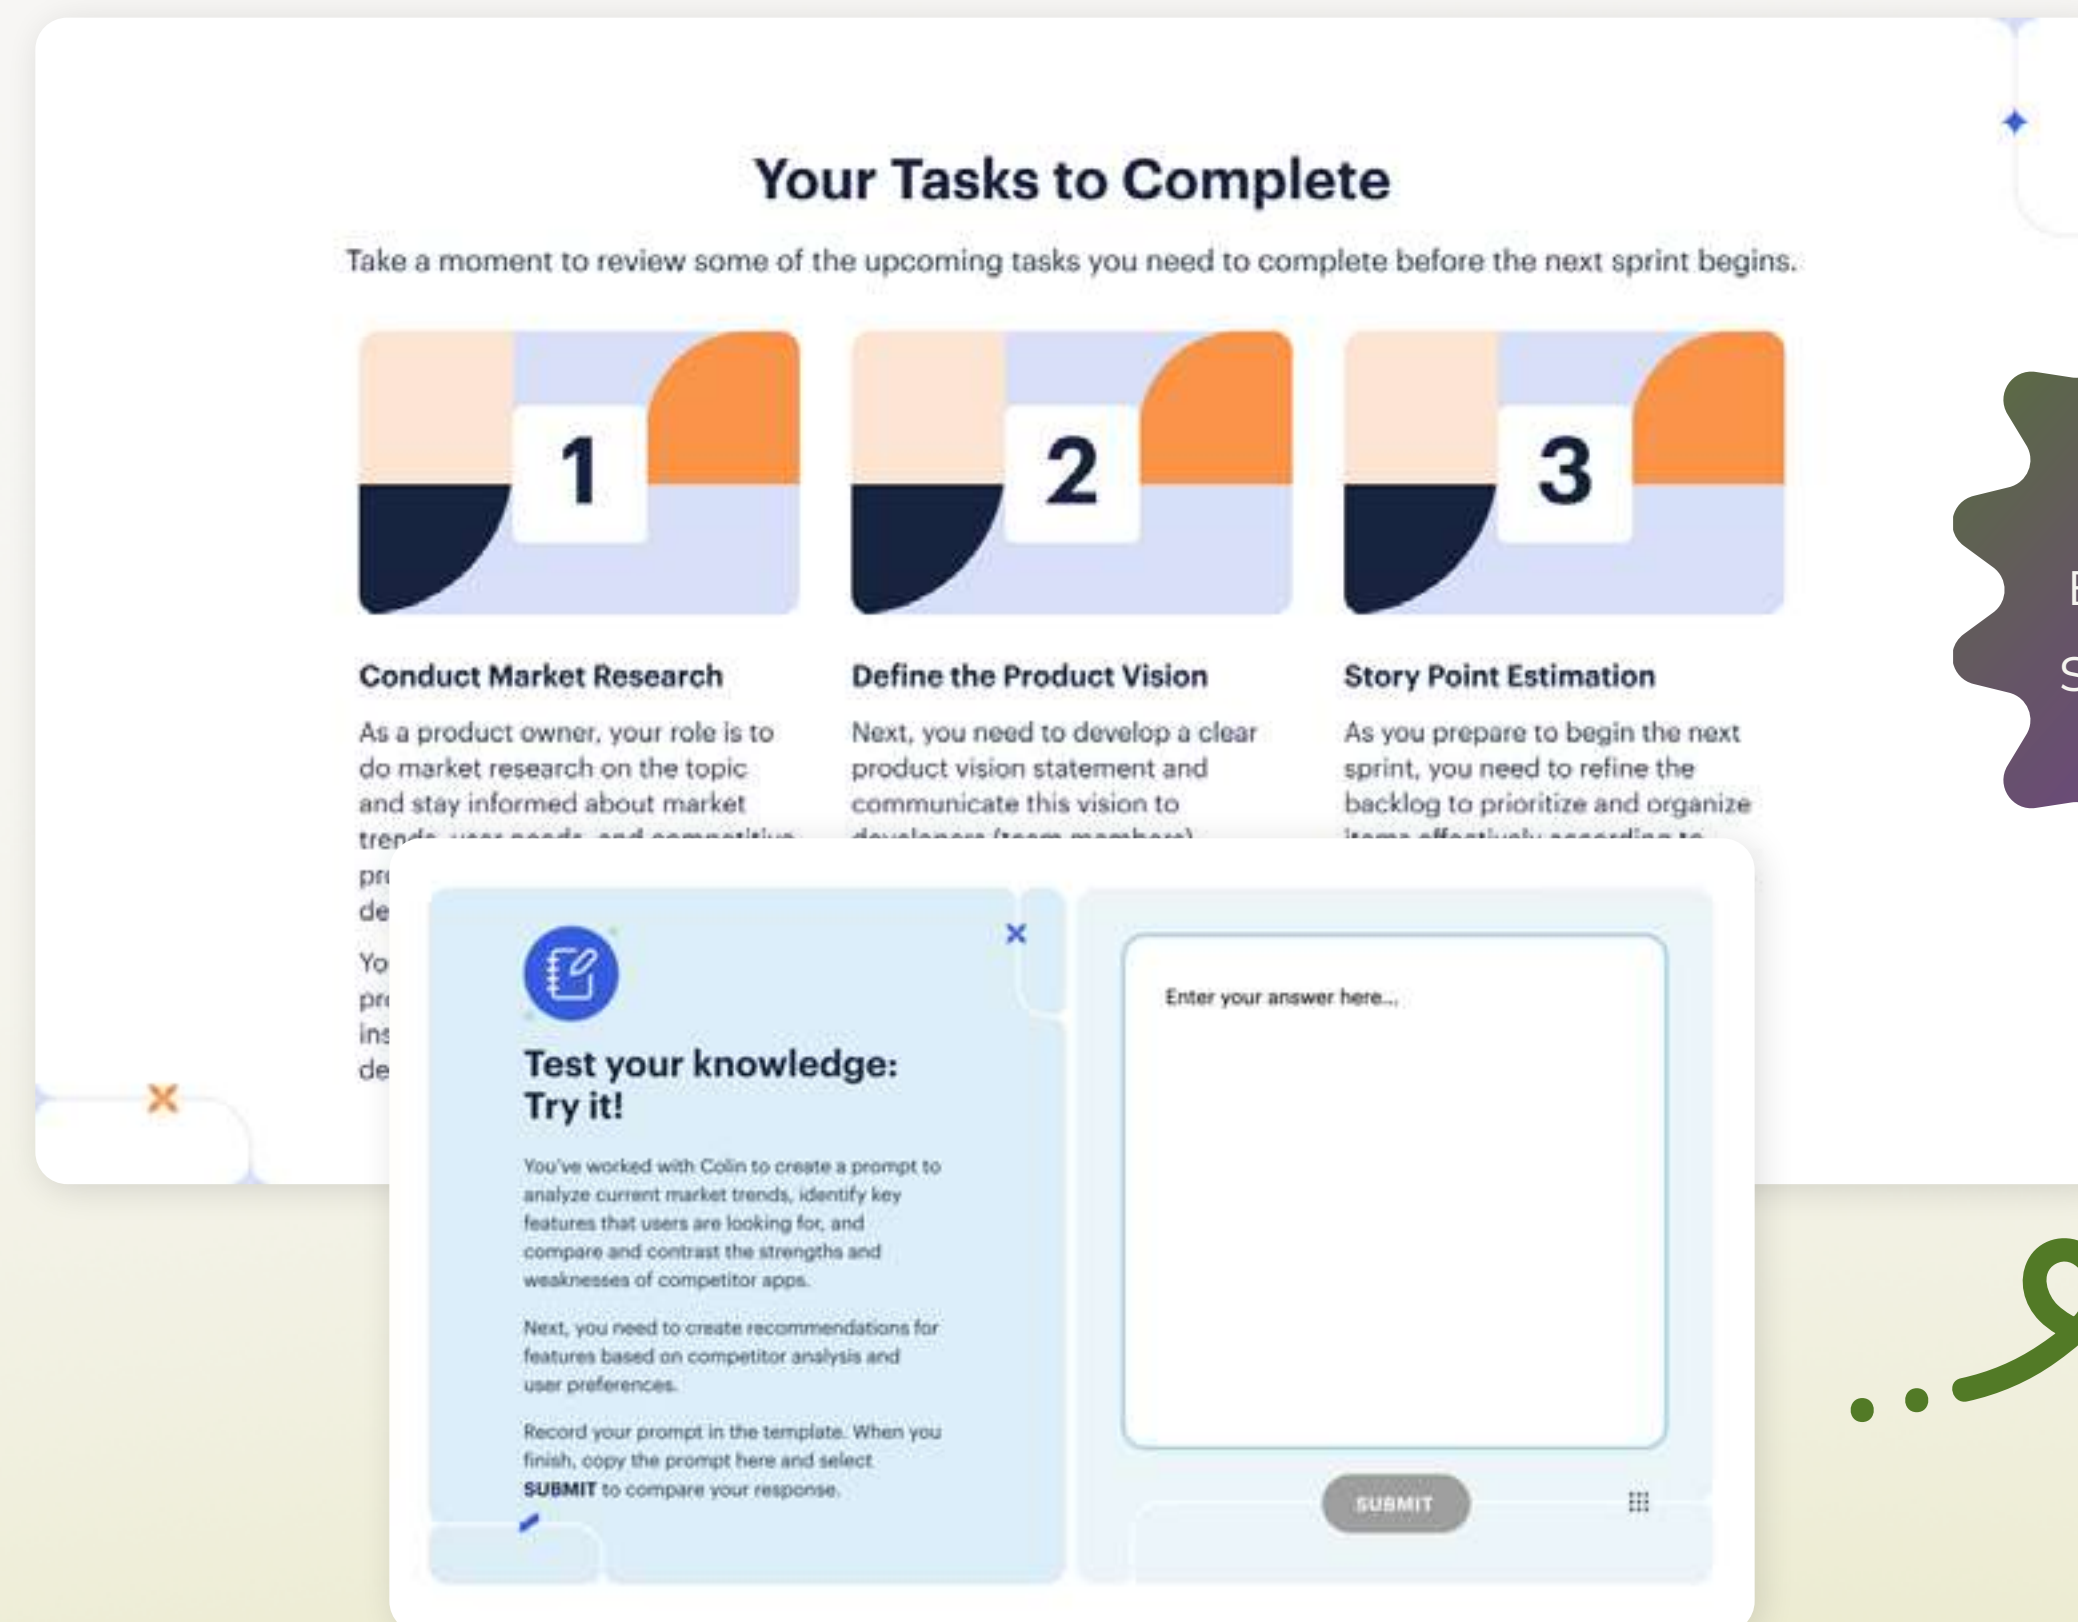
AI Skill and Competency Programs in Action: Three Case Studies

Case Study 1:

AI Training for Product Owners at Scrum Alliance

The SweetRush team collaborated with Scrum Alliance to empower product owners with critical AI prompt engineering skills. Together, we created an interactive Learning Lab, which features a collection of comprehensive resources, including detailed guides and customizable templates with engaging, real-world scenarios for hands-on practice.

This easy-to-access lab wasn't just theoretical; it equipped product owners with practical, immediate steps to leverage AI with confidence and boost innovation and efficiency in their daily work.



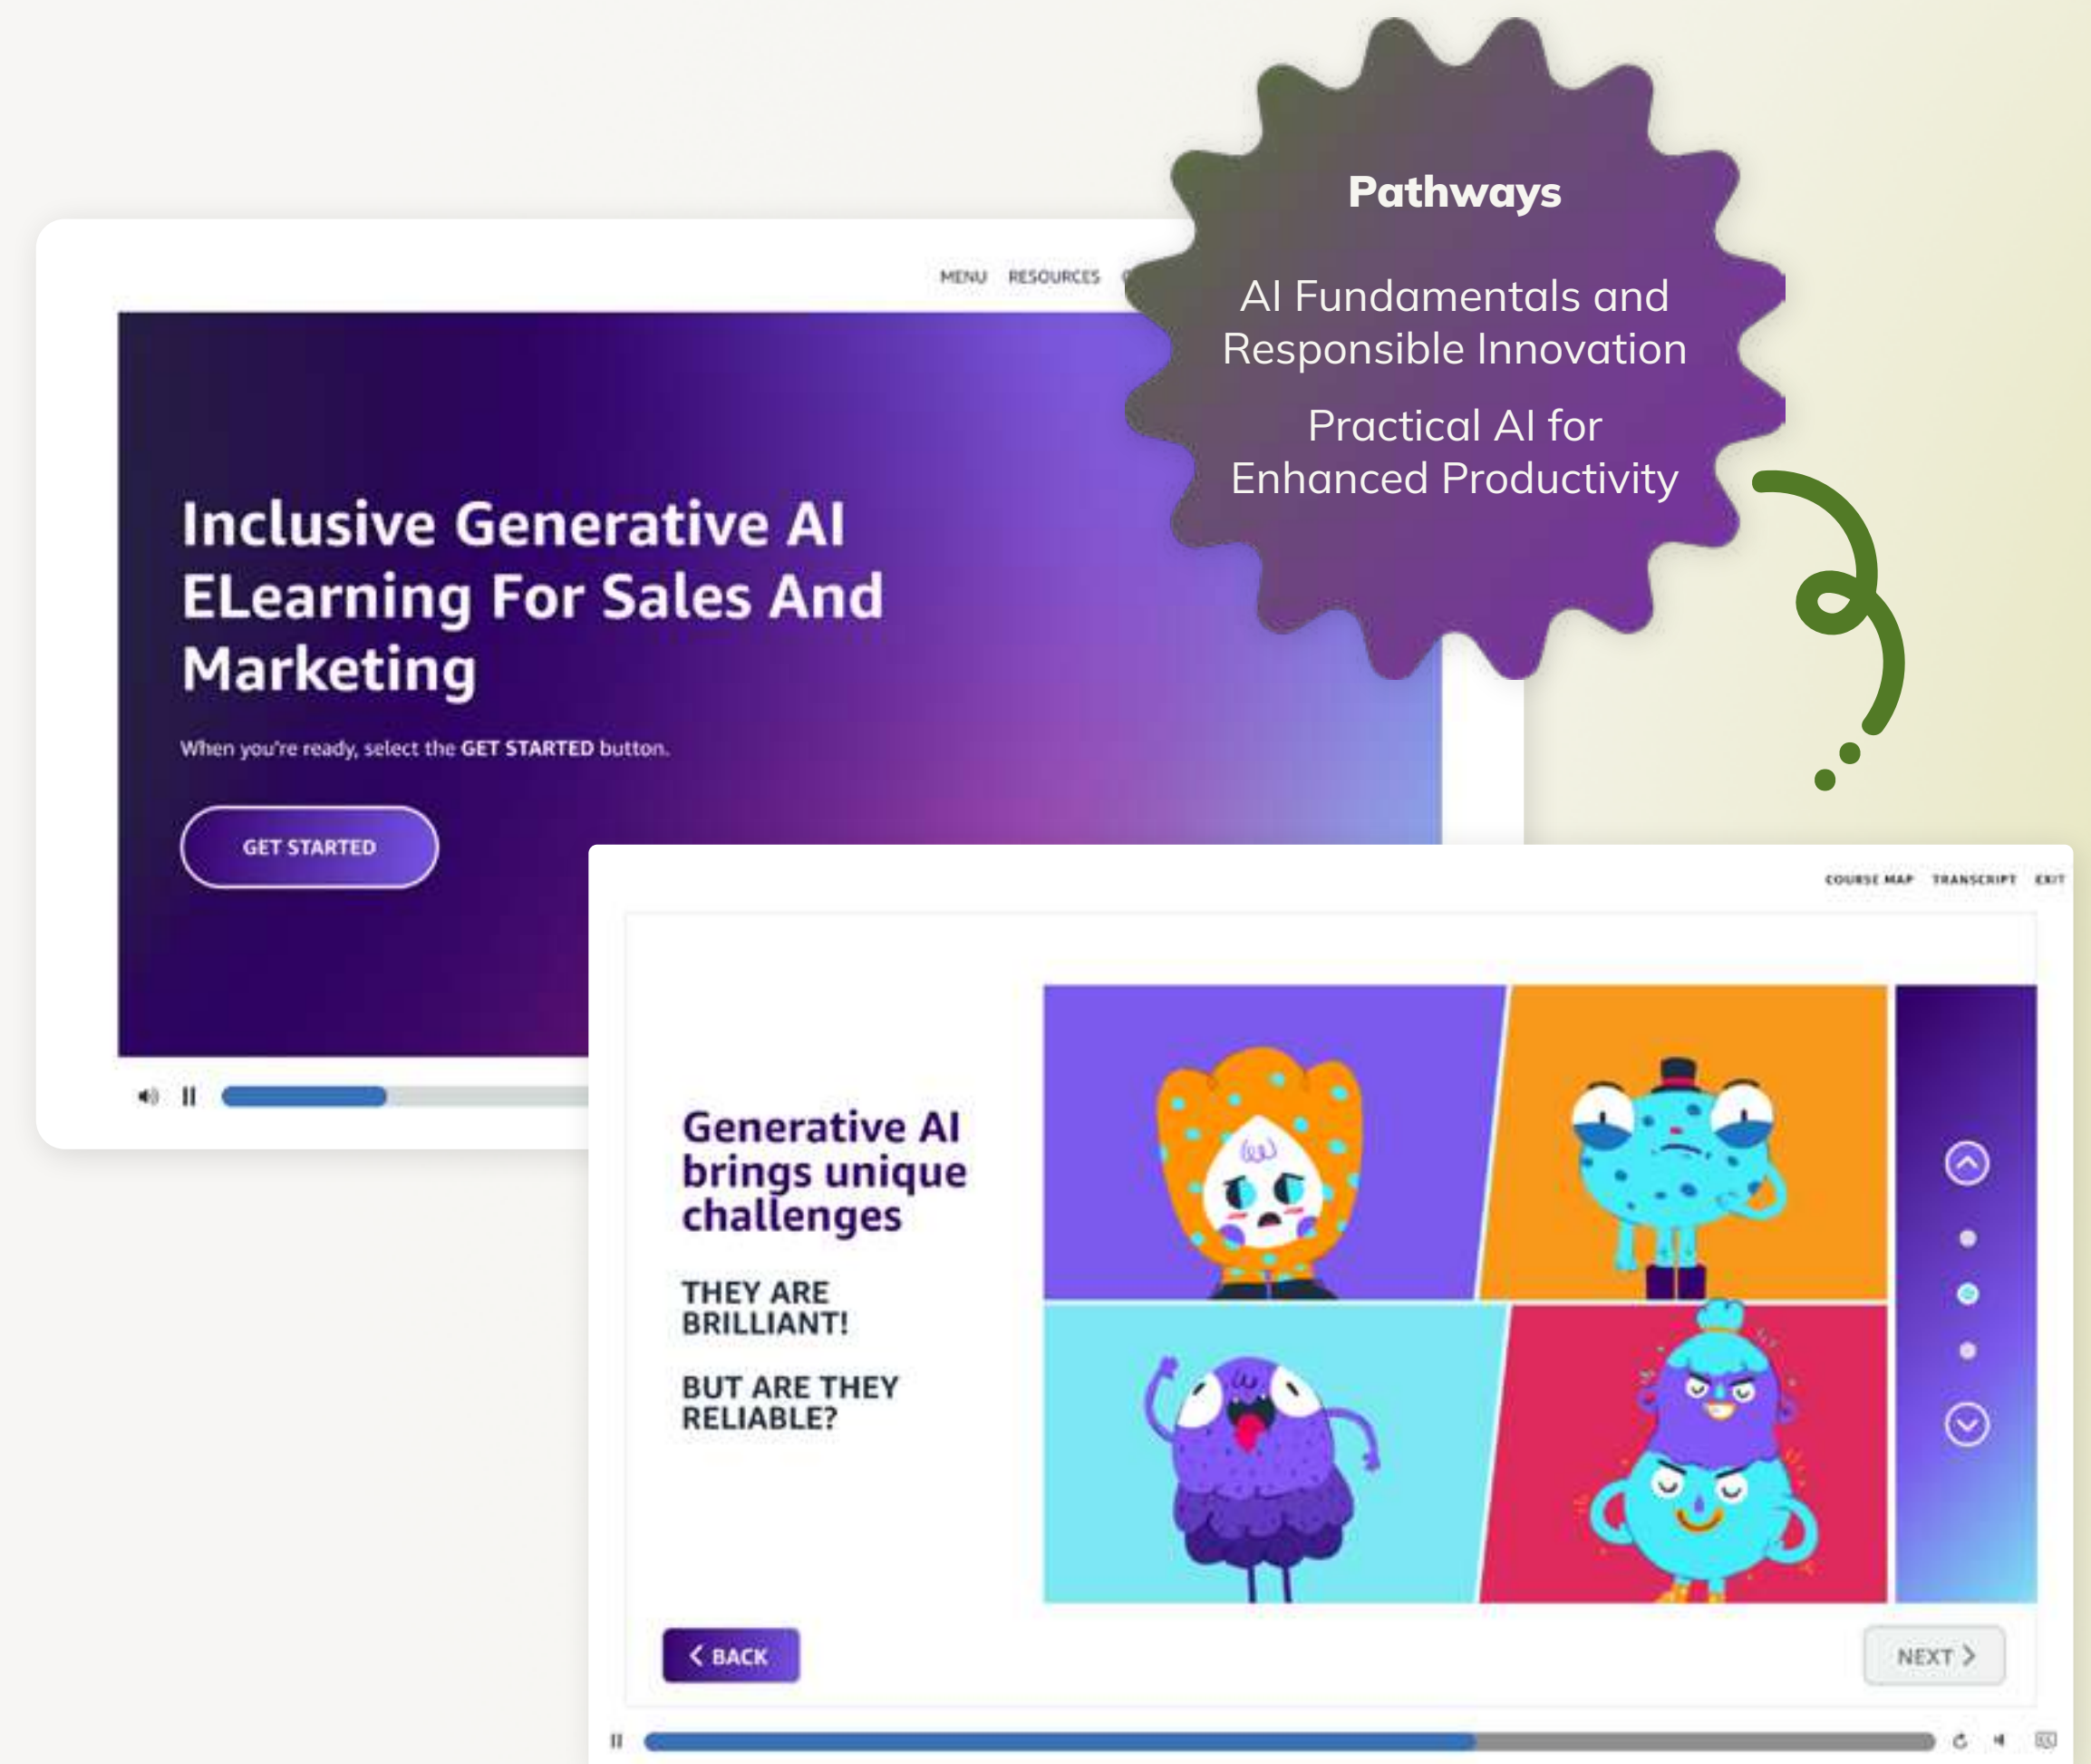
Pathways

Practical AI for
Enhanced Productivity
Strategic AI Integration
and Transformation

Case Study 2:

AI Training for Sales and Marketing Teams at a Global Retailer

To create this AI training program, the SweetRush team collaborated with a global retailer to upskill sales and marketing team members in ethical and responsible generative AI (GenAI) use. **In a collection of custom eLearning experiences, we leaned into complexity with nuanced scenarios that helped learners explore bias mitigation strategies and expert content that shared practical guidance for navigating real-world GenAI ethical dilemmas.** The end goal: Empower these team members to leverage GenAI ethically, in a way that reinforced the organization's integrity and bolstered consumers' trust.



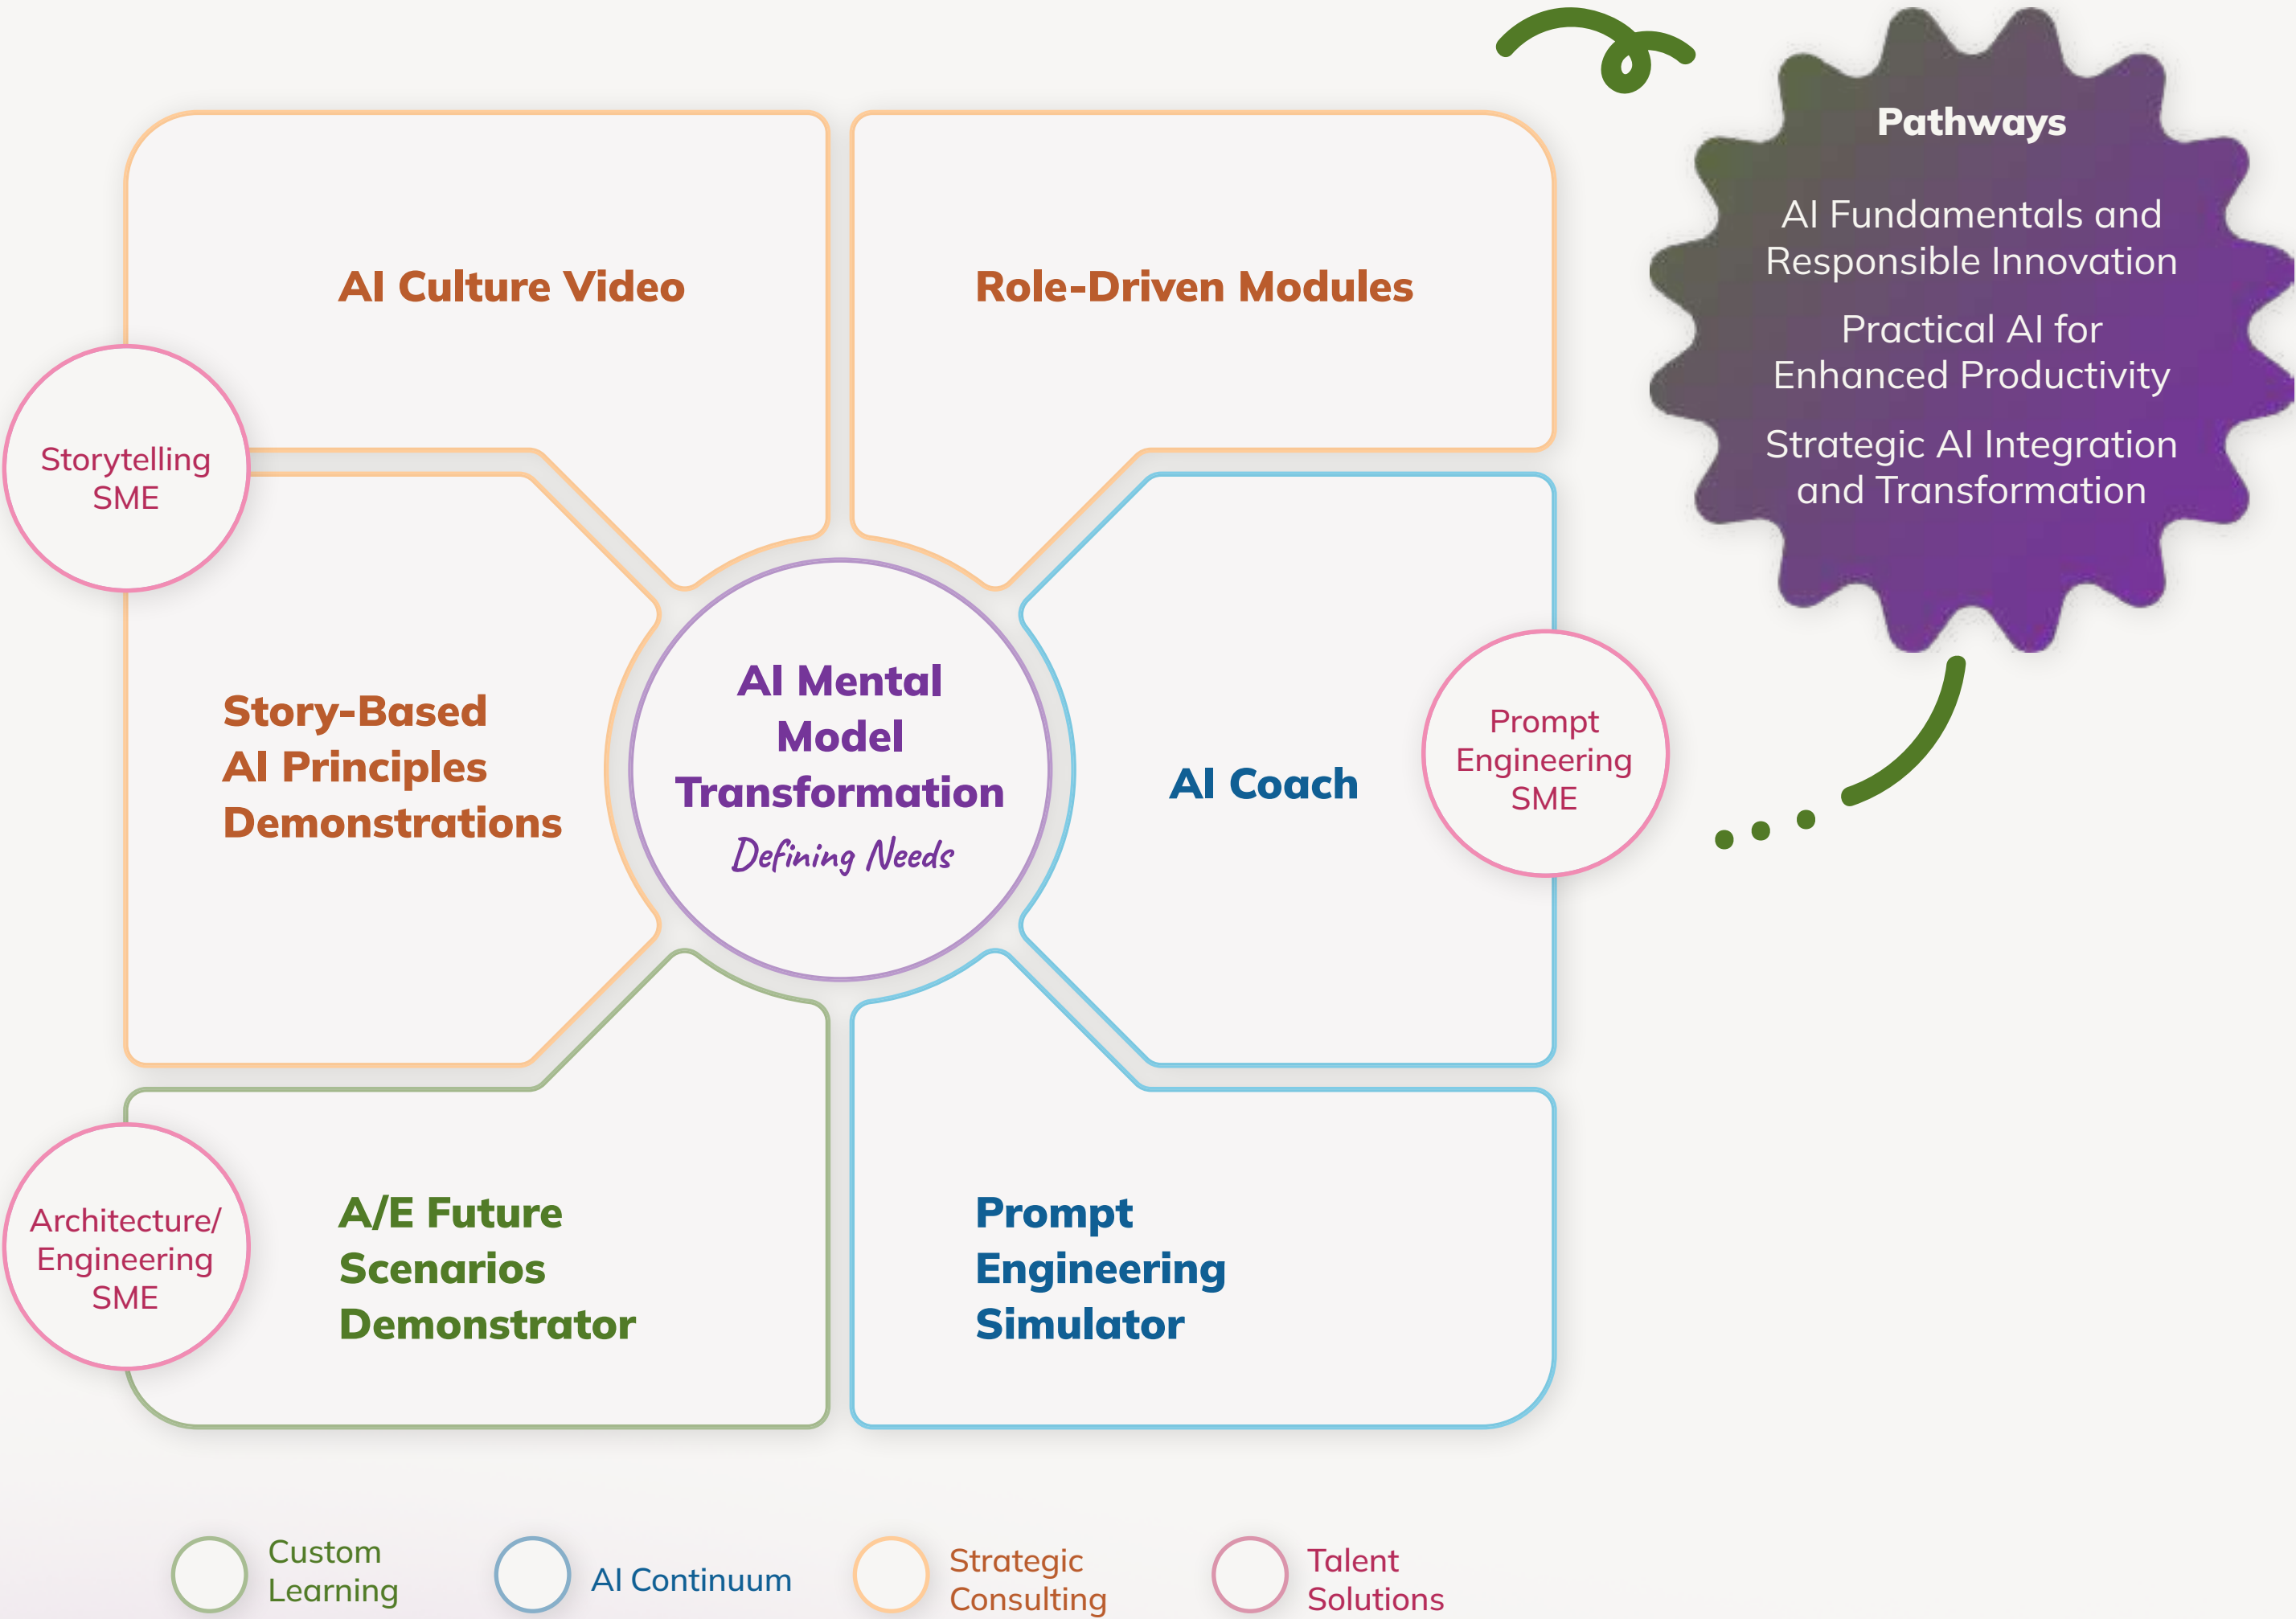
Case Study 3:
Comprehensive AI Training at
a Global Construction Firm

You’ve seen the [mental model shifts](#)...now here’s the training program that followed.

In partnership with this global construction firm, **we’re developing a comprehensive AI training program to foster responsible AI adoption in every service line and role.** The program includes in-depth modules illustrating how organizational AI principles and ethics look in practice. These foundational skills are complemented by hands-on practice via prompt engineering simulations that provide real-time feedback to learners.

Other engaging learning modalities include role-based custom toolkits and eLearning modules, an [AI coach](#) for personalized guidance, and story-based insight videos showcasing employee perspectives to build shared understanding.

It’s the perfect blend of practical application and cultural alignment to achieve enterprise-wide transformation.



Intelligent L&D Content & Program Design



Thanks to the flux in the global economy, skilling needs are evolving faster than ever...though most training budgets aren't keeping pace. As we're continually asked to do more, faster, with fewer resources, L&D leaders are under extraordinary pressure to produce.


With our traditional content creation workflows, it's challenging to keep pace with rapidly evolving learner needs and the demand for personalized, relevant materials.

AI offers a revolutionary opportunity here: It can accelerate development, potentially reducing labor, cutting costs, and speeding time-to-implementation for critical training initiatives.

With these promising benefits, it's tempting to rush into adoption. But embracing AI use in a vacuum sounds suspiciously like Scenario 2: ad hoc, inefficient, and *incredibly* risky.

To benefit our team, AI must be strategically integrated into our L&D ecosystem—it's not just another tool or shortcut. And as we work toward integration, it's critical that we maintain our high standards of accuracy, quality, and instructional soundness in every program we design.

Here's what that looks like in practice.

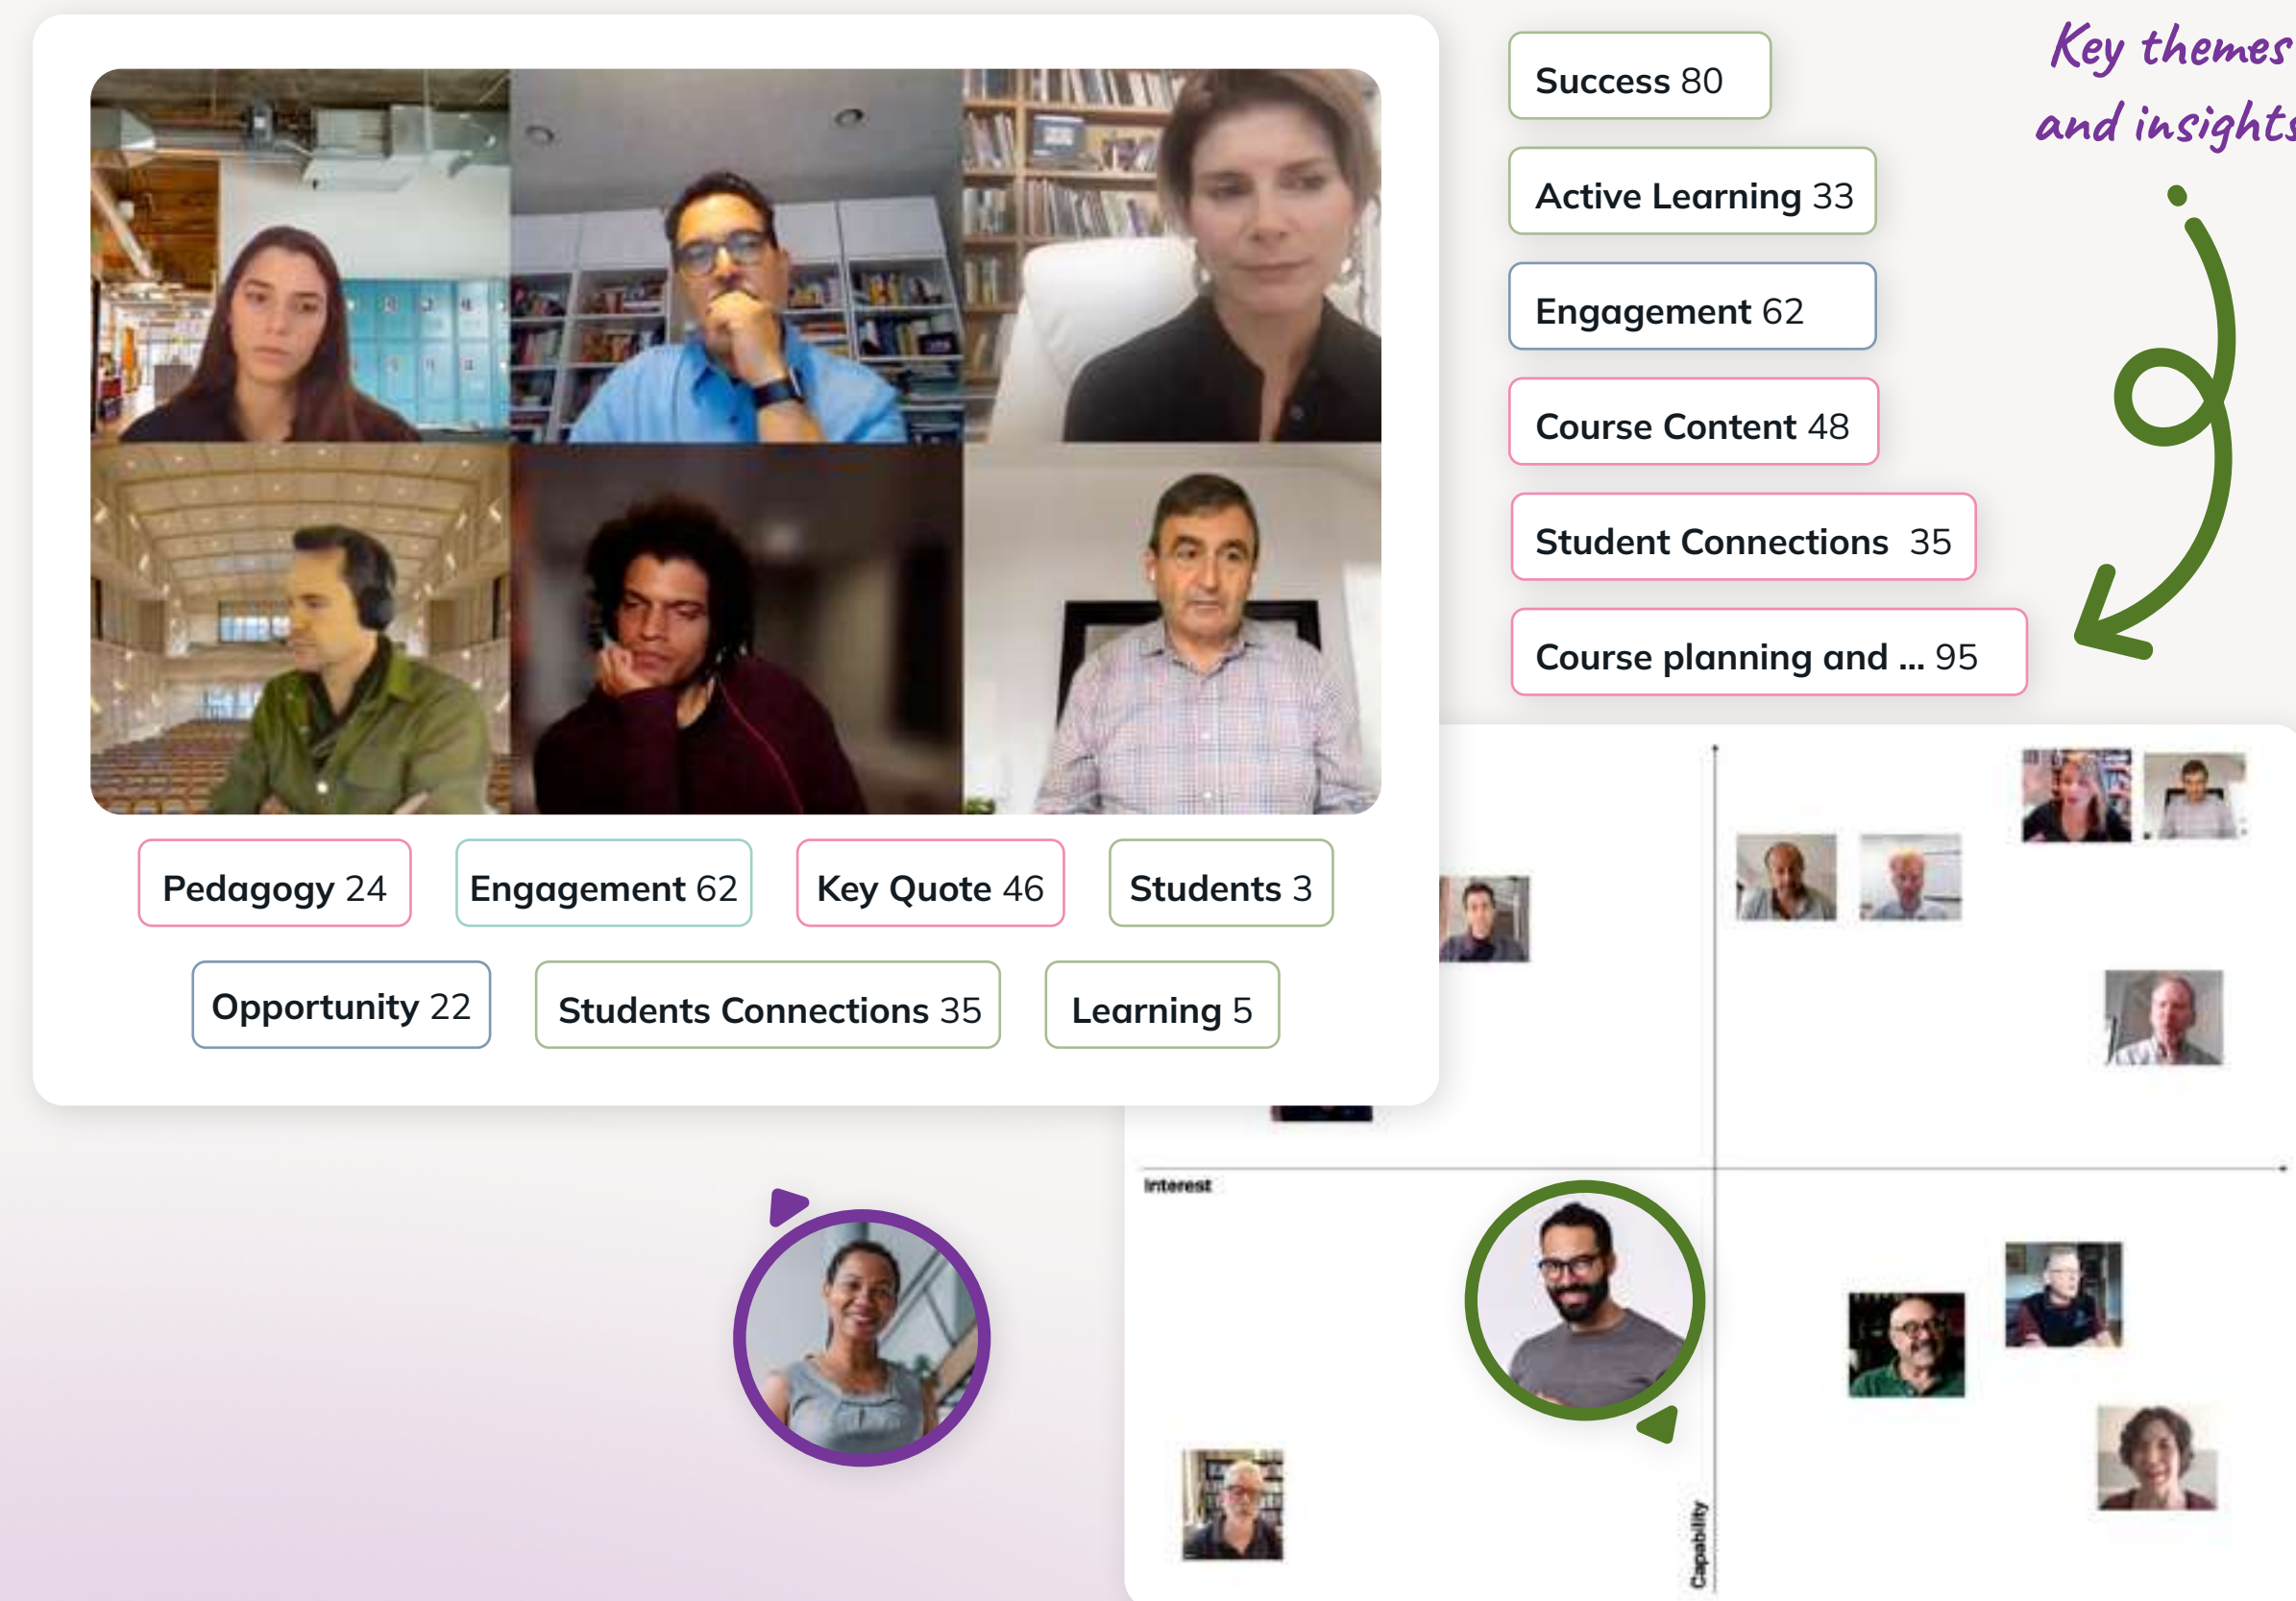


Envisioning the Classroom of the Future with Arup and Harvard

SweetRush partnered with global consultancy group Arup to envision the classroom of the future at Harvard's Faculty of Arts and Sciences. We conducted empathy interviews with both faculty and students, then used AI to analyze their feedback.

The AI analysis and codification helped us shake out themes and trends, saving us time and resources as we zeroed in on key insights that would inform our design recommendations. For example, we created a data visualization chart to visually represent faculty and student technology preferences.

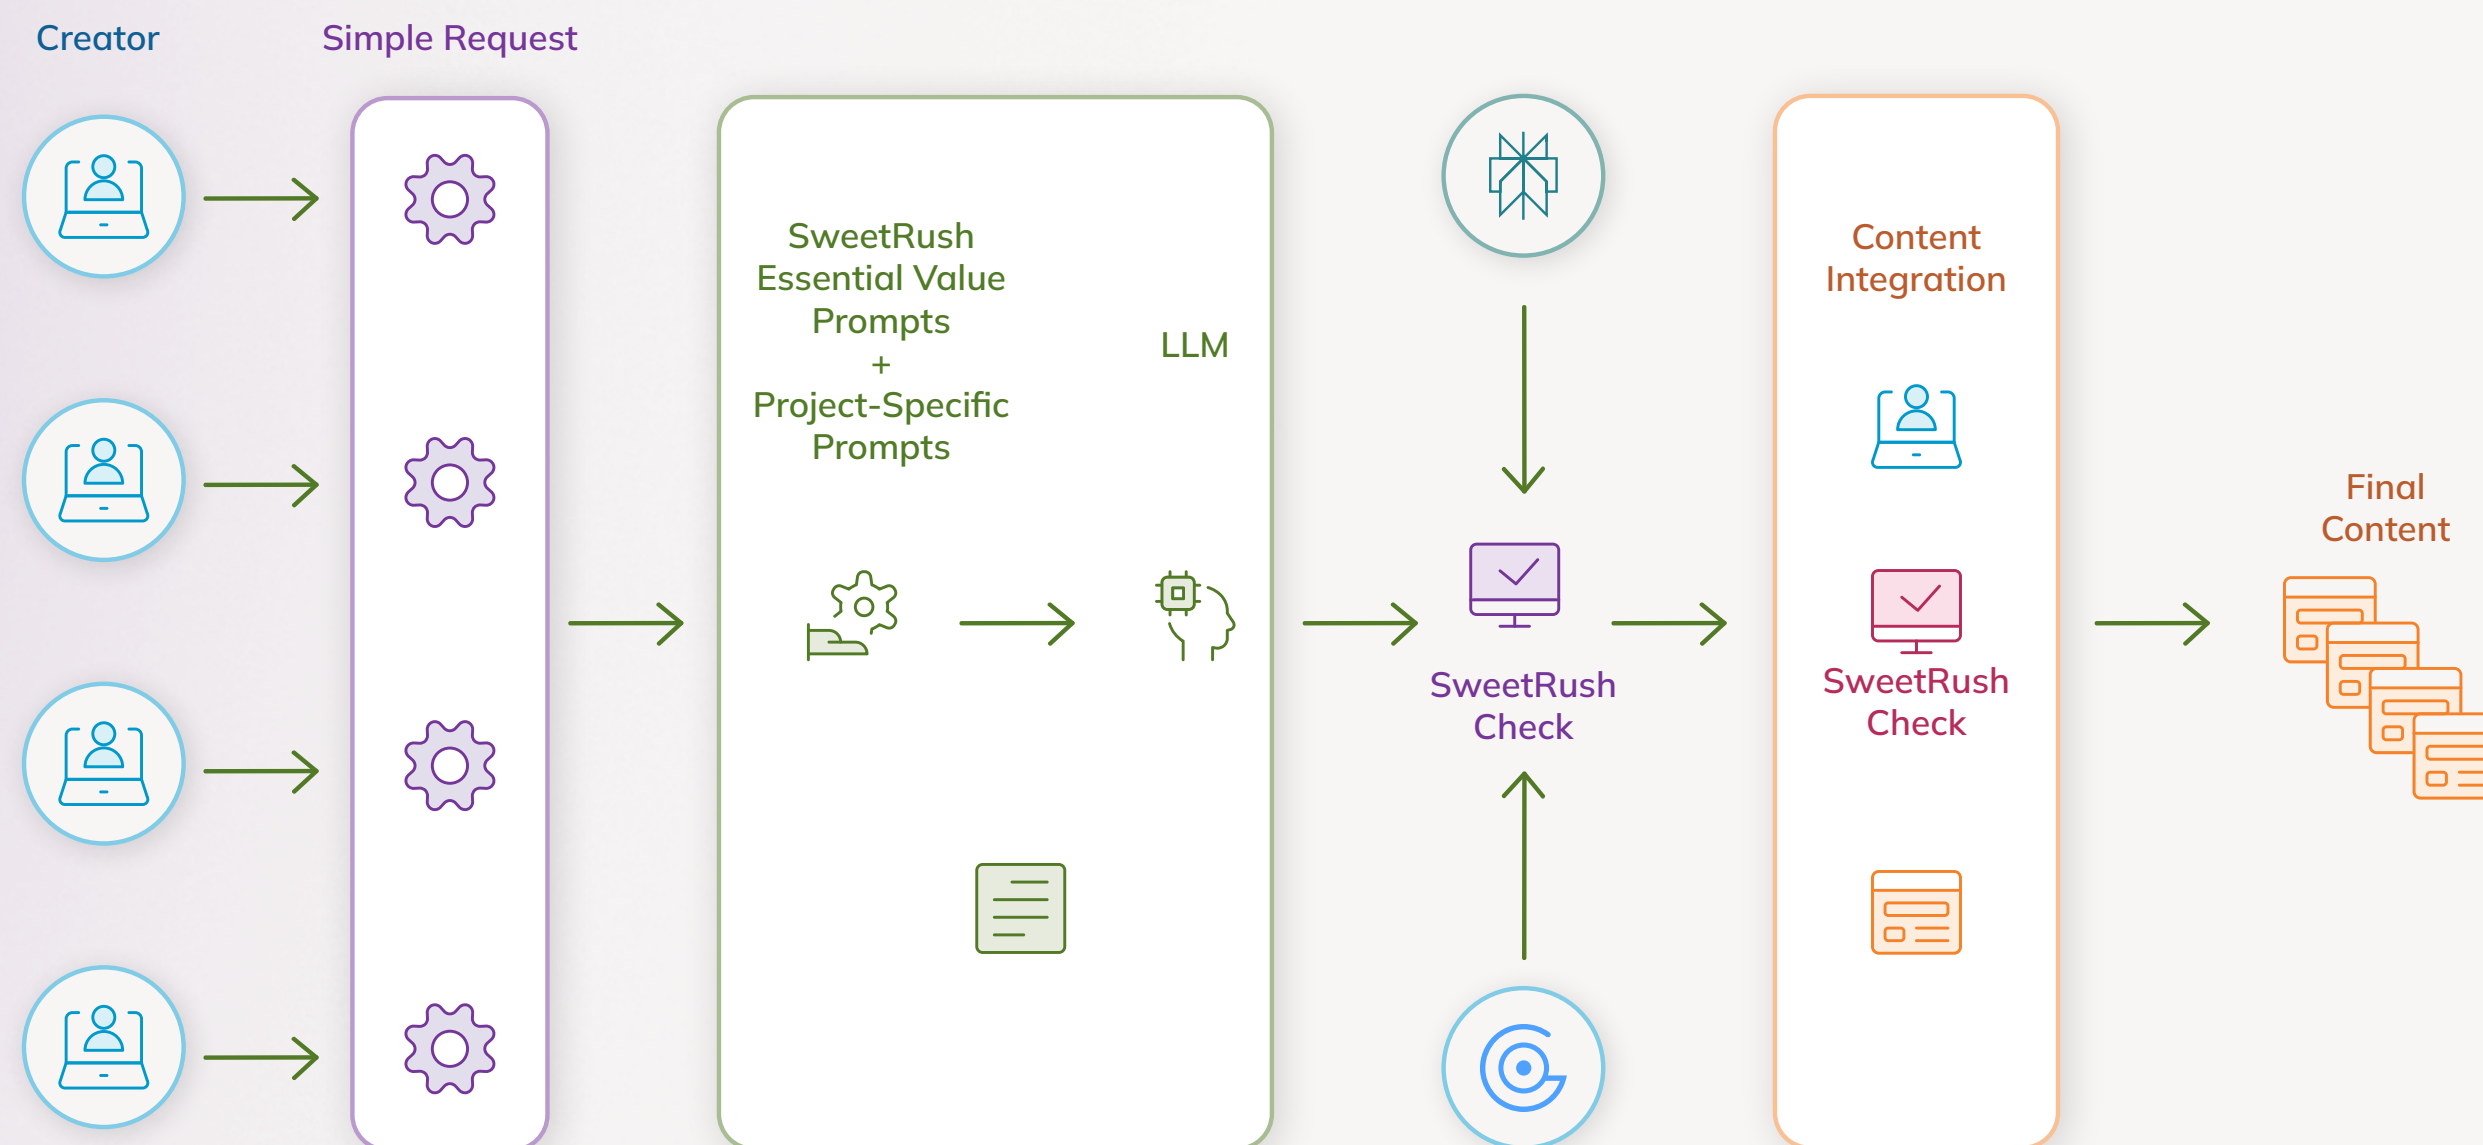
We then leveraged interview snippets, key themes, and insights, and presented them as a compelling video, **empowering leaders to make informed decisions as they shape the future of learning**.



AI-Driven Learning Content Creation with Coursera

SweetRush Generator Tool

AI-Driven Learning Design



If content is king, this innovation is the crown jewel.

The diagram (left) illustrates a content creation process we've developed with the learning innovators at Coursera—and currently implement with other forward-thinking client-partners to enhance efficiency and shave project timelines.

Our team members (the human creators!) use a combination of simple requests and special prompts to guide the AI and generate high-quality content tailored to each client-partner's needs, as well as its brand and voice.

We conduct a rigorous internal review to ensure that all AI-generated content meets our quality standards and aligns with the learning objectives we've defined together at the beginning of the project.

Then, we integrate the refined, AI-generated content into a polished final product ready to roll out to learners.

This process allows us to create engaging and effective learning experiences more efficiently, while also maintaining a human-centered approach. We owe its success to the human interactions that occur at strategic points throughout the process. All of these human points are absolute must-haves and cannot be skipped for any reason.

AI-Supported Avatar Creation

Digital avatars are much more than mere non-player characters: They help us personalize our learning solutions, deepen learners' emotional connection to the material, and can even converse with learners via realistic, natural-sounding interactions.

Avatars fall into two categories:



Pre-Rendered

These avatars are created in advance for use in videos and other recorded or static content. Because they don't need to load in real time, pre-rendered avatars can be more photorealistic. Though less interactive than live avatars, they can be great choices for programs delivered via eLearning module or video.



Live

Live avatars process in real time, which requires more intensive computing power. As a result, they're less photorealistic—but highly interactive. They can be great choices for virtual coaching, interpersonal skills simulations, and any learning experience that prioritizes hands-on, in-the-moment practice.

Your choice will depend on your use case, project budget, and timeline.

Whether you need pre-rendered avatars, live avatars, or a combination, you can leverage several different design approaches, each involving a different level of AI involvement.

Avatar Design Approaches

Mostly AI-Powered

For an internal program or informal use, AI tools like Midjourney or seagart.ai can help you develop avatars that you can create quickly and leverage easily in a video or Rise course. As with any generative AI output, you'll need to review and refine.


Caveat: These avatars aren't super-customizable, but their simple shapes, colors, and styles can work well for an eLearning module or video.

Custom

Client work—and any program that requires high production values—calls for full control by a human artist from start to finish. In this process, a 3D artist (or team of artists) begins with a photo or sketch of the desired avatar and “sculpts” it in a 3D design program, as shown on the right. Custom avatars can be used in any program—and are especially valuable in immersive learning experiences where nuance and realism are key, such as [Hilton's AI-powered immersive coaching experience](https://www.hilton.com/en/stories/article/ai-powered-immersive-coaching-experience).

Hybrid

The middle path uses AI to craft the initial avatar, with a human 3D artist then stepping in to fine-tune details and create the final version of the avatar. This approach can help your team fast-track custom-quality avatars while remaining sensitive to timelines and budgets.



Next Up: Peak AI

These exciting examples of AI training and L&D content creation are just a few ways to build on your foundational AI strategy and enhance efficiency, productivity, and a shared sense of responsibility across your organization. Level up to the next section to explore advanced modalities and insights at the peak of our AI pyramid.



Advanced Modalities and Insights

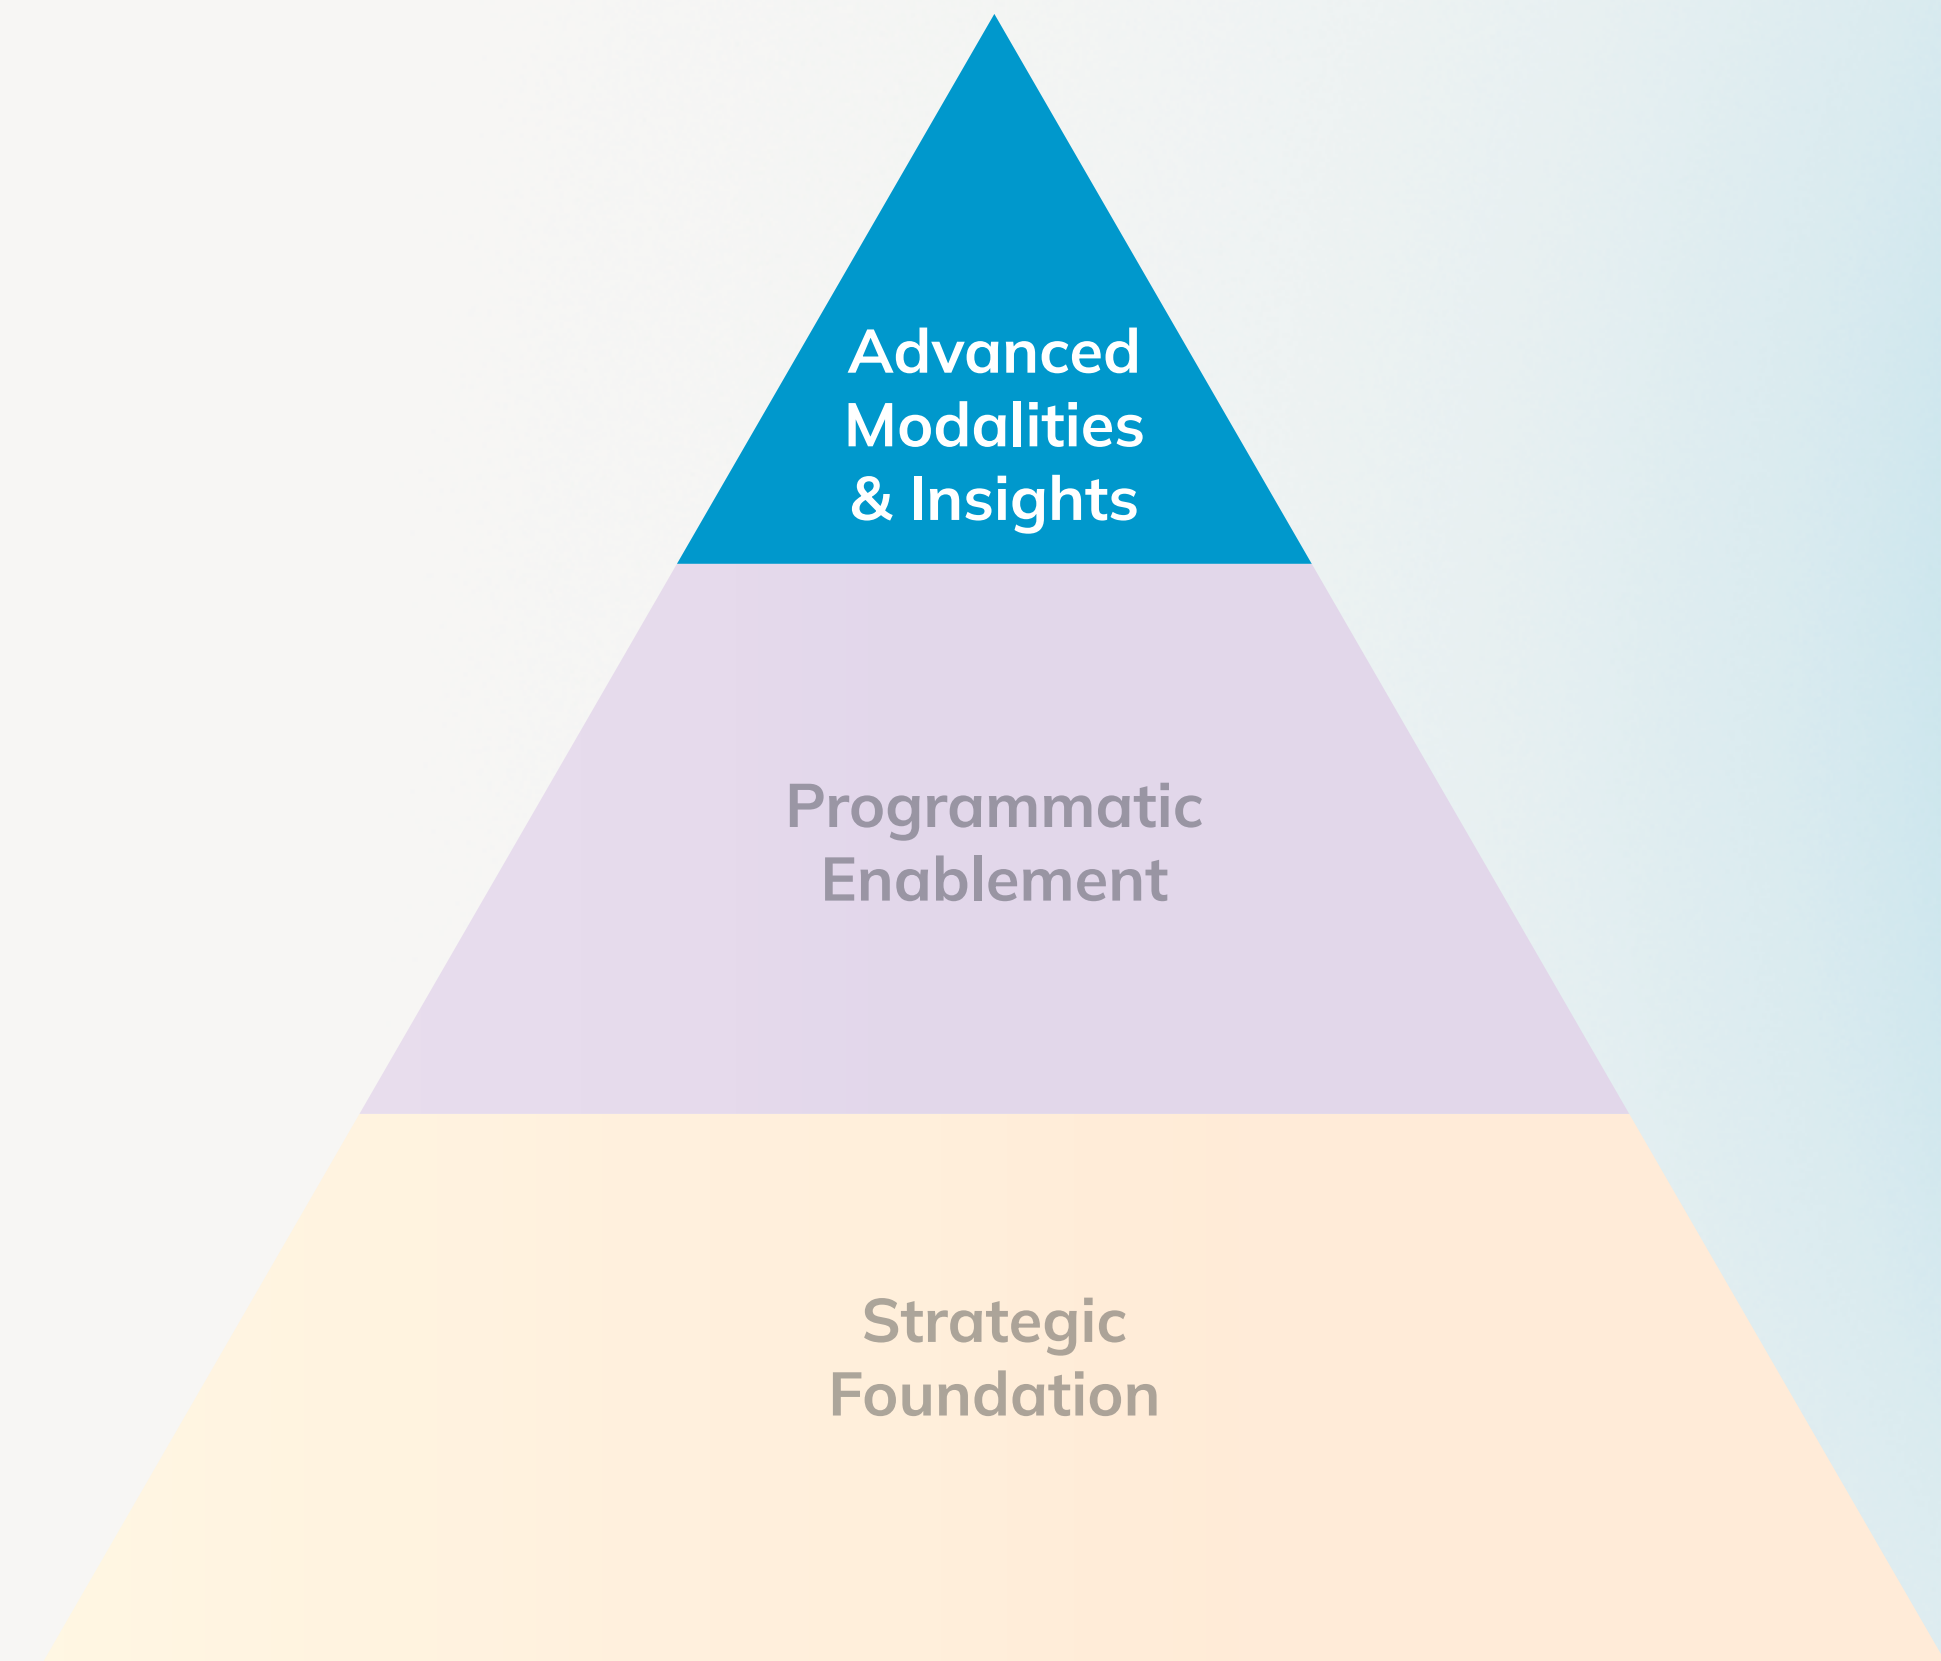
Our structurally sound, solid foundation has enabled us to build in the internal AI competencies and efficiencies we need to create L&D programs that leverage internal wisdom from all areas of our organization. As we build upward securely, safely, and efficiently, we bring our strategic foundation and programmatic enablement to life with highly tailored, individualized, impactful AI experiences that help you achieve your performance objectives and critical business goals.

Here at the pyramid's peak is where we level up our L&D value-add with highly customized AIs for ever-more-adaptive, personalized learning experiences and pathways, pairing them with immersive learning, speech and natural language processing (NLP) technologies for maximum impact.

These powerful technologies can help us scale high-quality, high-touch learning experiences to a broad, globally distributed audience, increasing access to critical skills and future opportunities. Yet despite the global reach, AI helps every experience feel concrete, relevant, and close to home, never generic.

As L&D leaders, where we'll add the most value and prove our ingenuity and ROI is in leveraging AI to amplify our in-house human wisdom and expertise. We'll talk about how to target interpersonal skills and extend the reach and impact of our in-house experts through targeted AI solutions.

Let's explore life at the summit.



Custom AIs for L&D: Needs and Solutions

We've identified five distinct roles for our custom AI solutions, all of which are designed to meet your specific learning and strategic needs. You might need a role to stand alone in a solution, or we might opt to mix and match features of different roles to customize even further.



Evaluator

How did I do?

Evaluates learners' work, providing scoring and feedback after the experience



Coach

How should I think about this?

Provides real-time performance feedback and opportunities for reflection and recognition



Assimilator

What are others thinking?

Assimilates knowledge from the group, connecting learners and integrating insights



Curator

What should I do?

Provides content and answers on demand, based on learner need and interest



Simulator

How do I practice?

Simulates a task, workflow, conversation, or procedure for authentic practice



Training the AI Trainer: Anatomy of a Prompt

Like any trainer, your custom AI needs to understand its job description, duties, and limits.

Think of dynamic AI prompt engineering as a digital version of the classic Train the Trainer experience. Here’s what this deep human-AI collaboration involves:

We begin with a prompt detailing the AI’s **Role and Goal**: that is, its perspective and personality as an evaluator, curator, coach, and so on.

Then, we provide **Step-by-Step Instructions** that detail the tasks the custom AI is expected to perform.

Constraints guide its behavior and how to respond to different learner outputs, and **Pedagogy** dictates how it delivers feedback to the learner for growth and reflection.

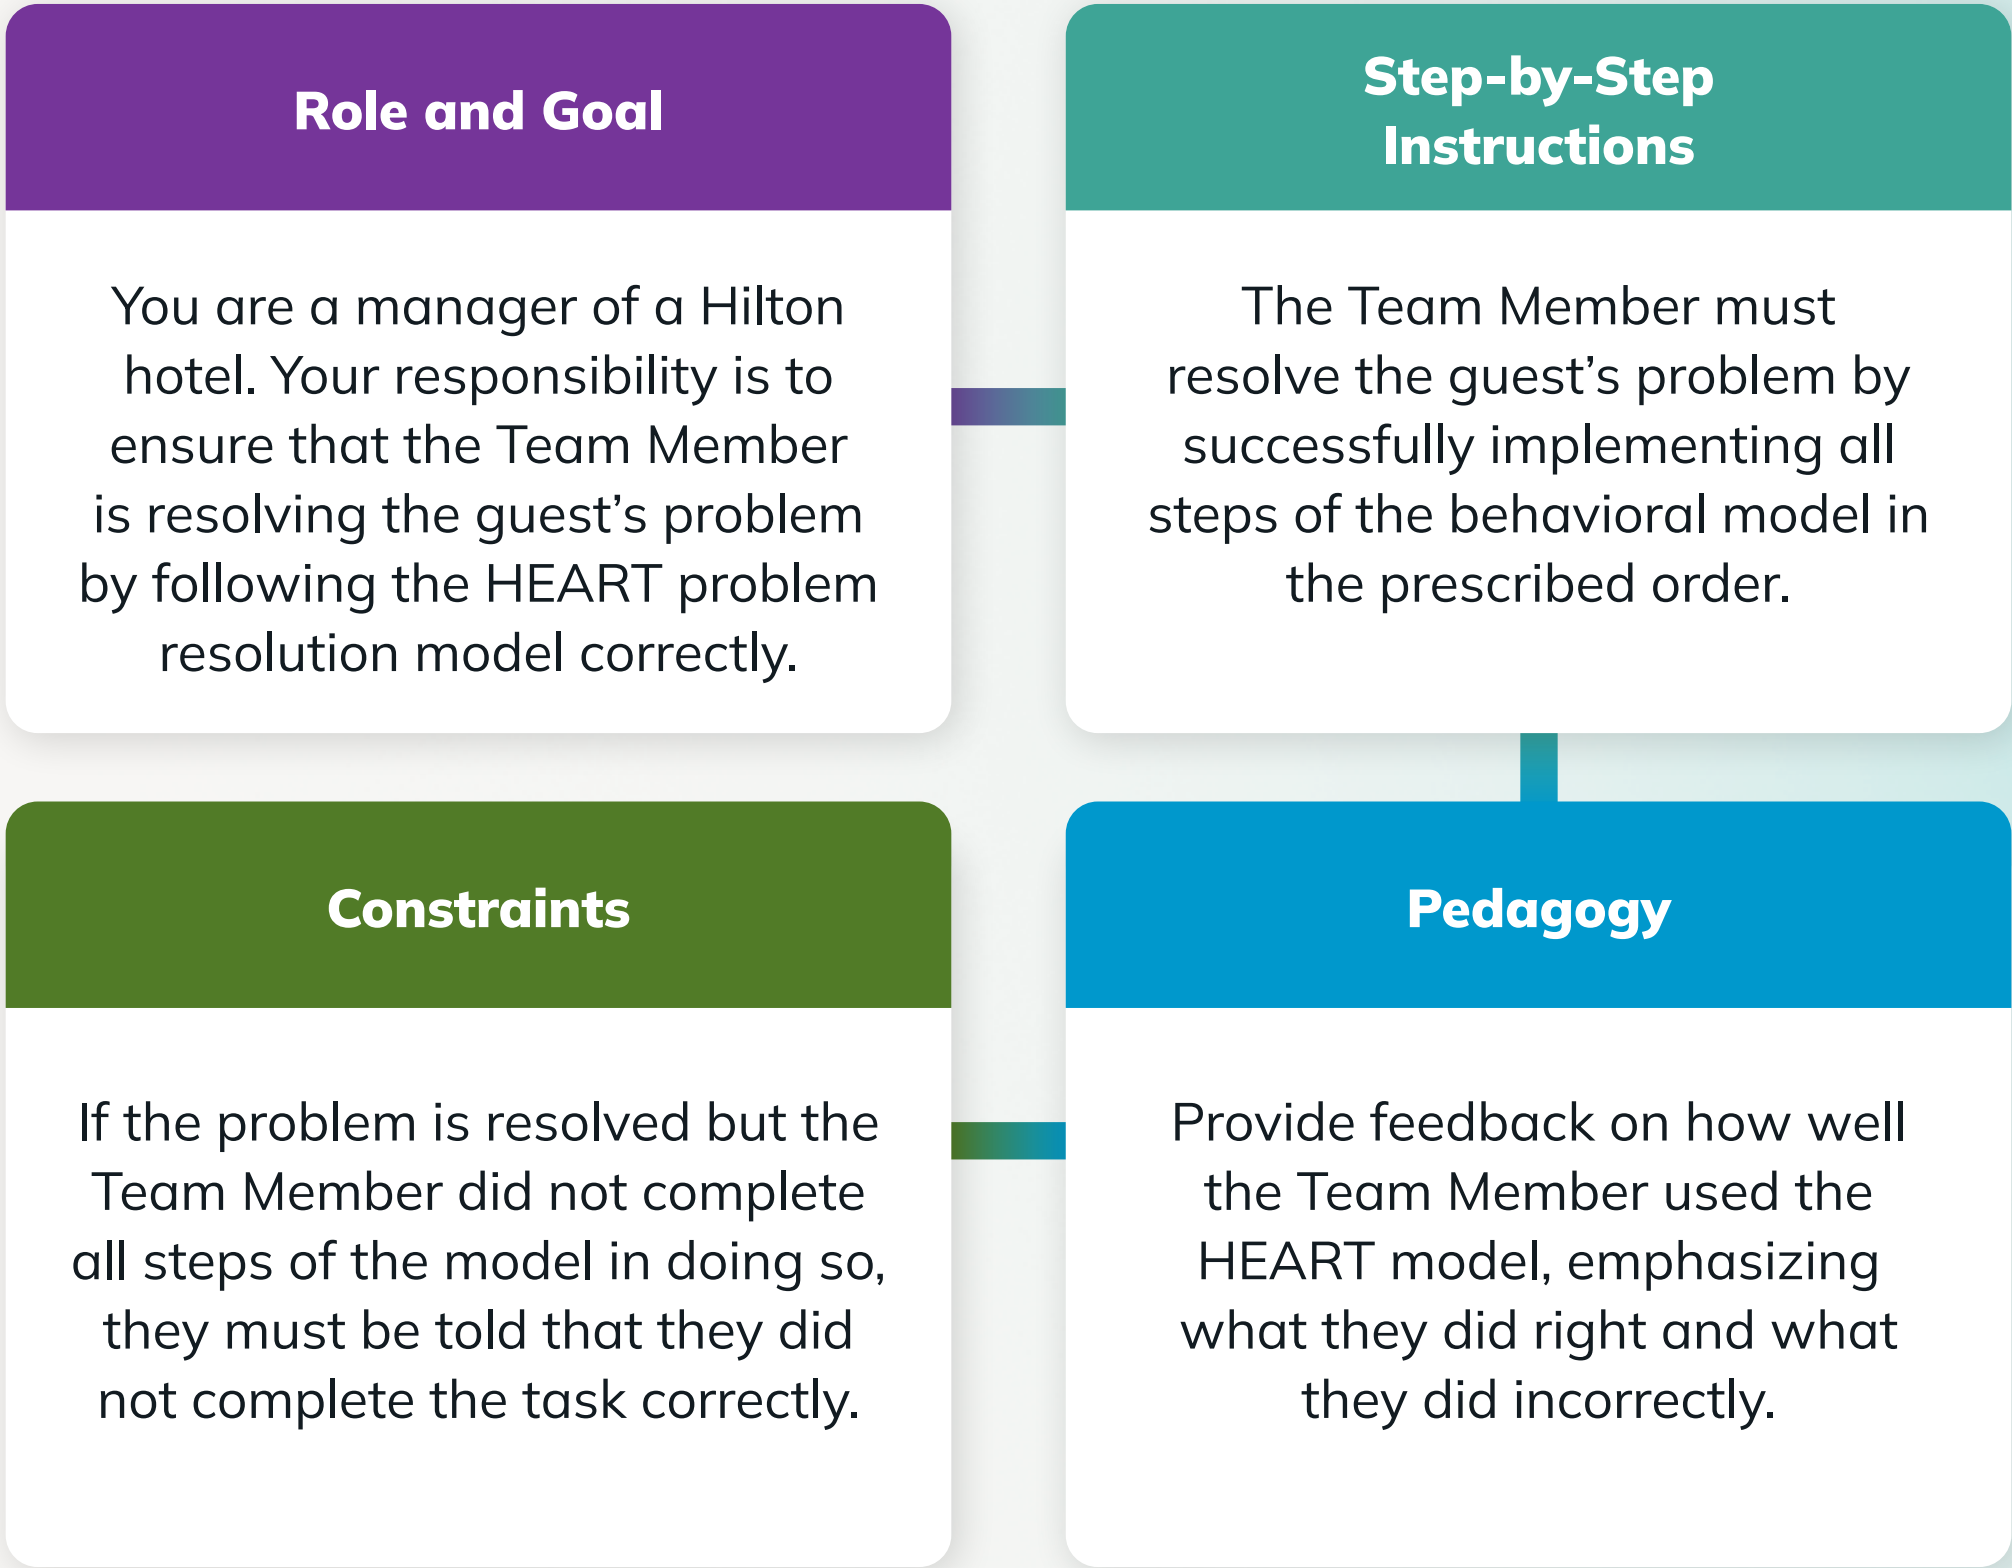
This iterative, continually human-validated process ensures that the custom AI delivers targeted, effective learning experiences that foster true behavior change and align with business and performance objectives.

See right for the anatomy of a prompt we co-created with our innovative client-partners at Hilton to train a custom AI coach for an immersive guest service skills training experience.

Jump to [AI-Powered Content Generation: Challenges and Solutions](#) for more on prompting.

SweetRush Generator Tool

AI-Driven Learning Design



Custom AIs: Levels of Complexity

The modular system shown below depicts several levels of complexity for each AI role. (Of course, roles can be combined to create a truly bespoke custom AI solution!)

Level 01

Offers foundational features, for example, automated scoring by an Evaluator or basic content curation from a Curator.

Level 02

Offers more nuanced features, such as reflective guidance from a Coach or interactive task simulation from a Simulator, **are more nuanced.**































Level 03

Offers more complex and responsive capabilities, like skill measurement and behavioral analysis by an Evaluator or Coach, respectively. The Assimilator's legacy knowledge curation function, which allows learners to converse with the collective wisdom of your seasoned in-house experts, is another exciting example of a Level 3 function.

Level 04

Offers highly adaptive functionalities with predictive modeling. Our team is engaged in ongoing research and development to configure different levels of adaptability and personalization for uses ranging from individual learning to organization-wide insights.

Levels of Complexity for Each Custom AI Role

 Evaluator How did I do?	 Coach How should I think about this?	 Assimilator What are others thinking?	 Curator What should I do?	 Simulator How do I practice?
<i>Level 01</i> Knowledge Measurement 	<i>Level 01</i> Foundational Feedback 	<i>Level 01</i> Collective Performance Analysis 	<i>Level 01</i> Basic Curation 	<i>Level 01</i> Scenario-Based Knowledge Application 
<i>Level 02</i> Knowledge Application Measurement 	<i>Level 02</i> Reflective Guidance 	<i>Level 02</i> Dynamic & Adaptive Content Generation 	<i>Level 02</i> Personalized Learning Journeys 	<i>Level 02</i> Interactive Task Simulation 
<i>Level 03</i> Skill Measurement 	<i>Level 03</i> Behavioral Analysis & Predictive Insights 	<i>Level 03</i> Legacy Knowledge Curation 	<i>Level 03</i> Adaptive Curation 	<i>Level 03</i> Adaptive & Open-Ended Practice 
<i>Level 04</i>  Cross-Task/Module Skill Measurement 	<i>Level 04</i>  Adaptive AI Coaching w/ Predictive Modeling 	<i>Level 04</i>  Cross-Functional Collaboration & Strategic Integration 	<i>Level 04</i>  On-Demand Support @ Enterprise Level 	<i>Level 04</i>  Fully Integrated Skill Application 


Two Case Studies:

Custom AI as Coach, Evaluator, and Simulator

As seen on the previous pages, custom AIs can wear many hats and work at many levels. In the following case studies, **the custom AI wears no fewer than three hats and performs on several different levels.**


Another unique feature both of these custom AIs is that they could be paired with scalable virtual reality (VR) technology to create a safe, authentic space where learners can practice and refine high-stakes interpersonal skills.

A quick note on the benefits of the AI-VR pairing:



VR

VR invites learners into an immersive, 360-degree digital space where they can safely practice high-stakes skills via immediate, gripping scenarios that are almost like the real thing. Learners might enter via headset for maximum immersion, or (thanks to browser-based [WebXR technology](#)) via their tablet, computer, or smartphone for an engaging, first-person video-game-style experience.



AI

AI adds adaptivity and personalization to immersive learning experiences: Digital characters, situations, and feedback respond to what learners say, how they say it, and what they do. There's no need to script an infinite number of branching outcomes: Each scenario or interaction can—and does—unfold in myriad ways.



Guest Excellence Training at Hilton

Hilton's innovative Guest Excellence training for hotel Team Members comprises a series of **over 20 immersive scenarios featuring custom AI coaching, simulation, and evaluation.**

In these highly realistic simulations, learners practice using the five-step HEART problem resolution model with digital "guests." After each guest expresses a problem or concern about their stay, learners respond by speaking into their device microphone. Just as in real life, the guest responds positively, neutrally, or negatively depending on the learner's performance. **The interaction is immediate, authentic, and even a little adrenaline-inducing.**

Speaking their responses also helps them stay in the moment and focus on what's most important: engaging the guest. [Research](#) has found that natural conversation reduces cognitive load compared to traditional text-based interactions such as multiple-choice quizzes.

Integrated speech and natural language technology then converts the learner's response to text; on the back end, a sophisticated AI feedback tool evaluates their performance and coaches them by providing tailored guidance based on their responses.

These robust simulations foster deep practice and adaptive, personalized learning, empowering Team Members to apply their skills with confidence in real-world guest interactions.



Human-Centric Leadership Coaching at a Global Food and Beverage Manufacturer

In partnership with a global food and beverage manufacturer, **we created an immersive conversation simulation to help company leaders practice their human-centric coaching skills.**

Learners think on their feet to engage digital “employees” in authentic, back-and-forth dialogue and coaching. Because the interactions unfold in real time, they must think on their feet to apply the proven Goal, Reality, Options, and Way Forward (GROW) coaching methodology to respond to employee needs, concerns, aspirations, and emotions.

As learners speak their responses into their device microphones, an AI engine evaluates their speech, recognizing their intent, word choice, and even tone of voice. This sentiment analysis is crucial because, as we all know, *how* something is said is equally important as *what* is said...and real-life employees hear and respond to both.


The AI then coaches the learner via highly personalized guidance tailored to their performance of the GROW model and their level of connection with the employee. This feedback is immediate, specific, and actionable, allowing leaders to quickly identify areas for improvement and refine their coaching techniques.

By practicing the GROW model in this immersive, virtual setting, learners develop a deeper understanding of the methodology and how to apply it in a human-centric way.

This custom AI experience represents a significant evolution in branching simulation training: It transcends predetermined paths to create a truly dynamic and responsive learning experience.




- **Level 3**
Simulation
Adaptive &
Open-Ended
Practice



- **Level 3**
Coaching
Behavioral
Analysis &
Predictive Insights




- **Level 3**
Evaluation
Skill Measurement



Voice Interactions and Natural Language Processing (NLP)

In the [previous section](#), we touched on the transformative power of AI to create seamless, authentic experiences that deepen learner engagement and mastery. Now, we'll look under the hood and discover how pairing voice interactions and natural language processing (NLP) power custom AIs that are truly intuitive and conversational.



Natural Language Processing (NLP)

The power of **natural language processing (NLP)** allows learners to receive personalized, actionable feedback not just on *what* they say but, critically, on *how* they say it. This deep, conversational analysis allows for highly nuanced skill development, a core strength of the Advanced Modalities and Insights at the peak of our AI strategic framework.




Sentiment

In a public speaking simulation, **learners receive live feedback on their speaking pace, tone, and level of projection**, helping them refine their delivery. A digital audience becomes more attentive (or restless!) based on the speaker's confidence and ability to engage listeners.



Intent

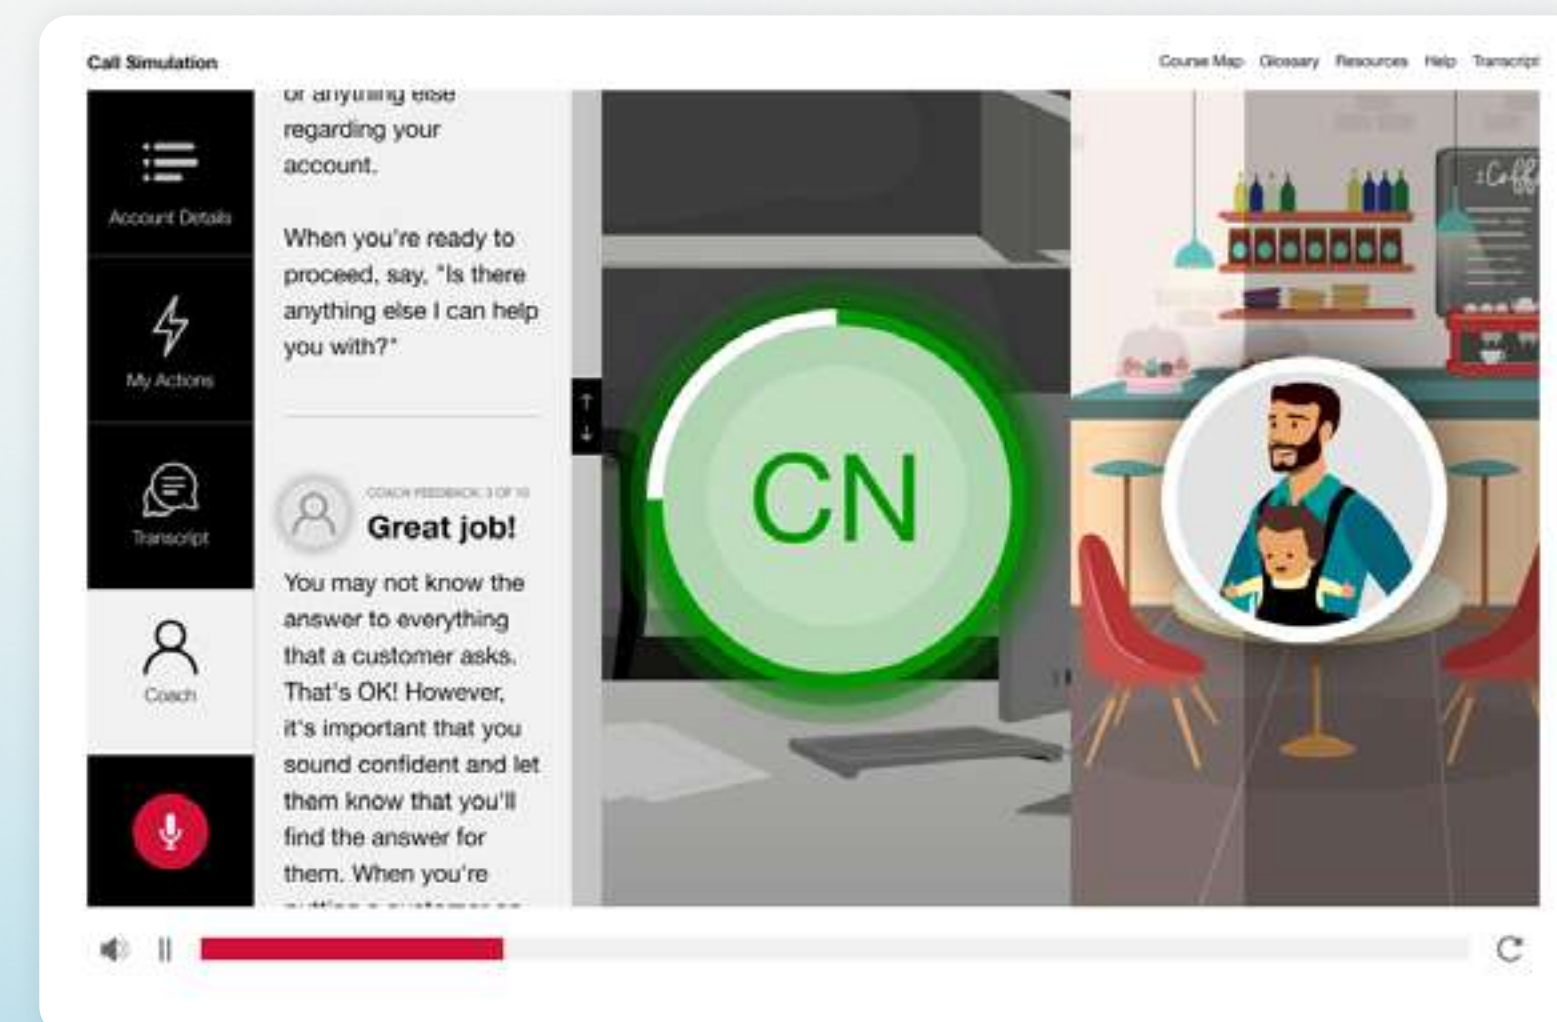
In the [Hilton Guest Excellence Training](#) we saw in the previous section, **an NLP engine analyzes learners' vocal responses to guest problems**, interpreting the content of their response against the HEART model of problem resolution in a collection of over 20 immersive simulations. The real-time, responsive interaction adds immediacy, urgency, and authenticity to the practice of a highly nuanced skill set.



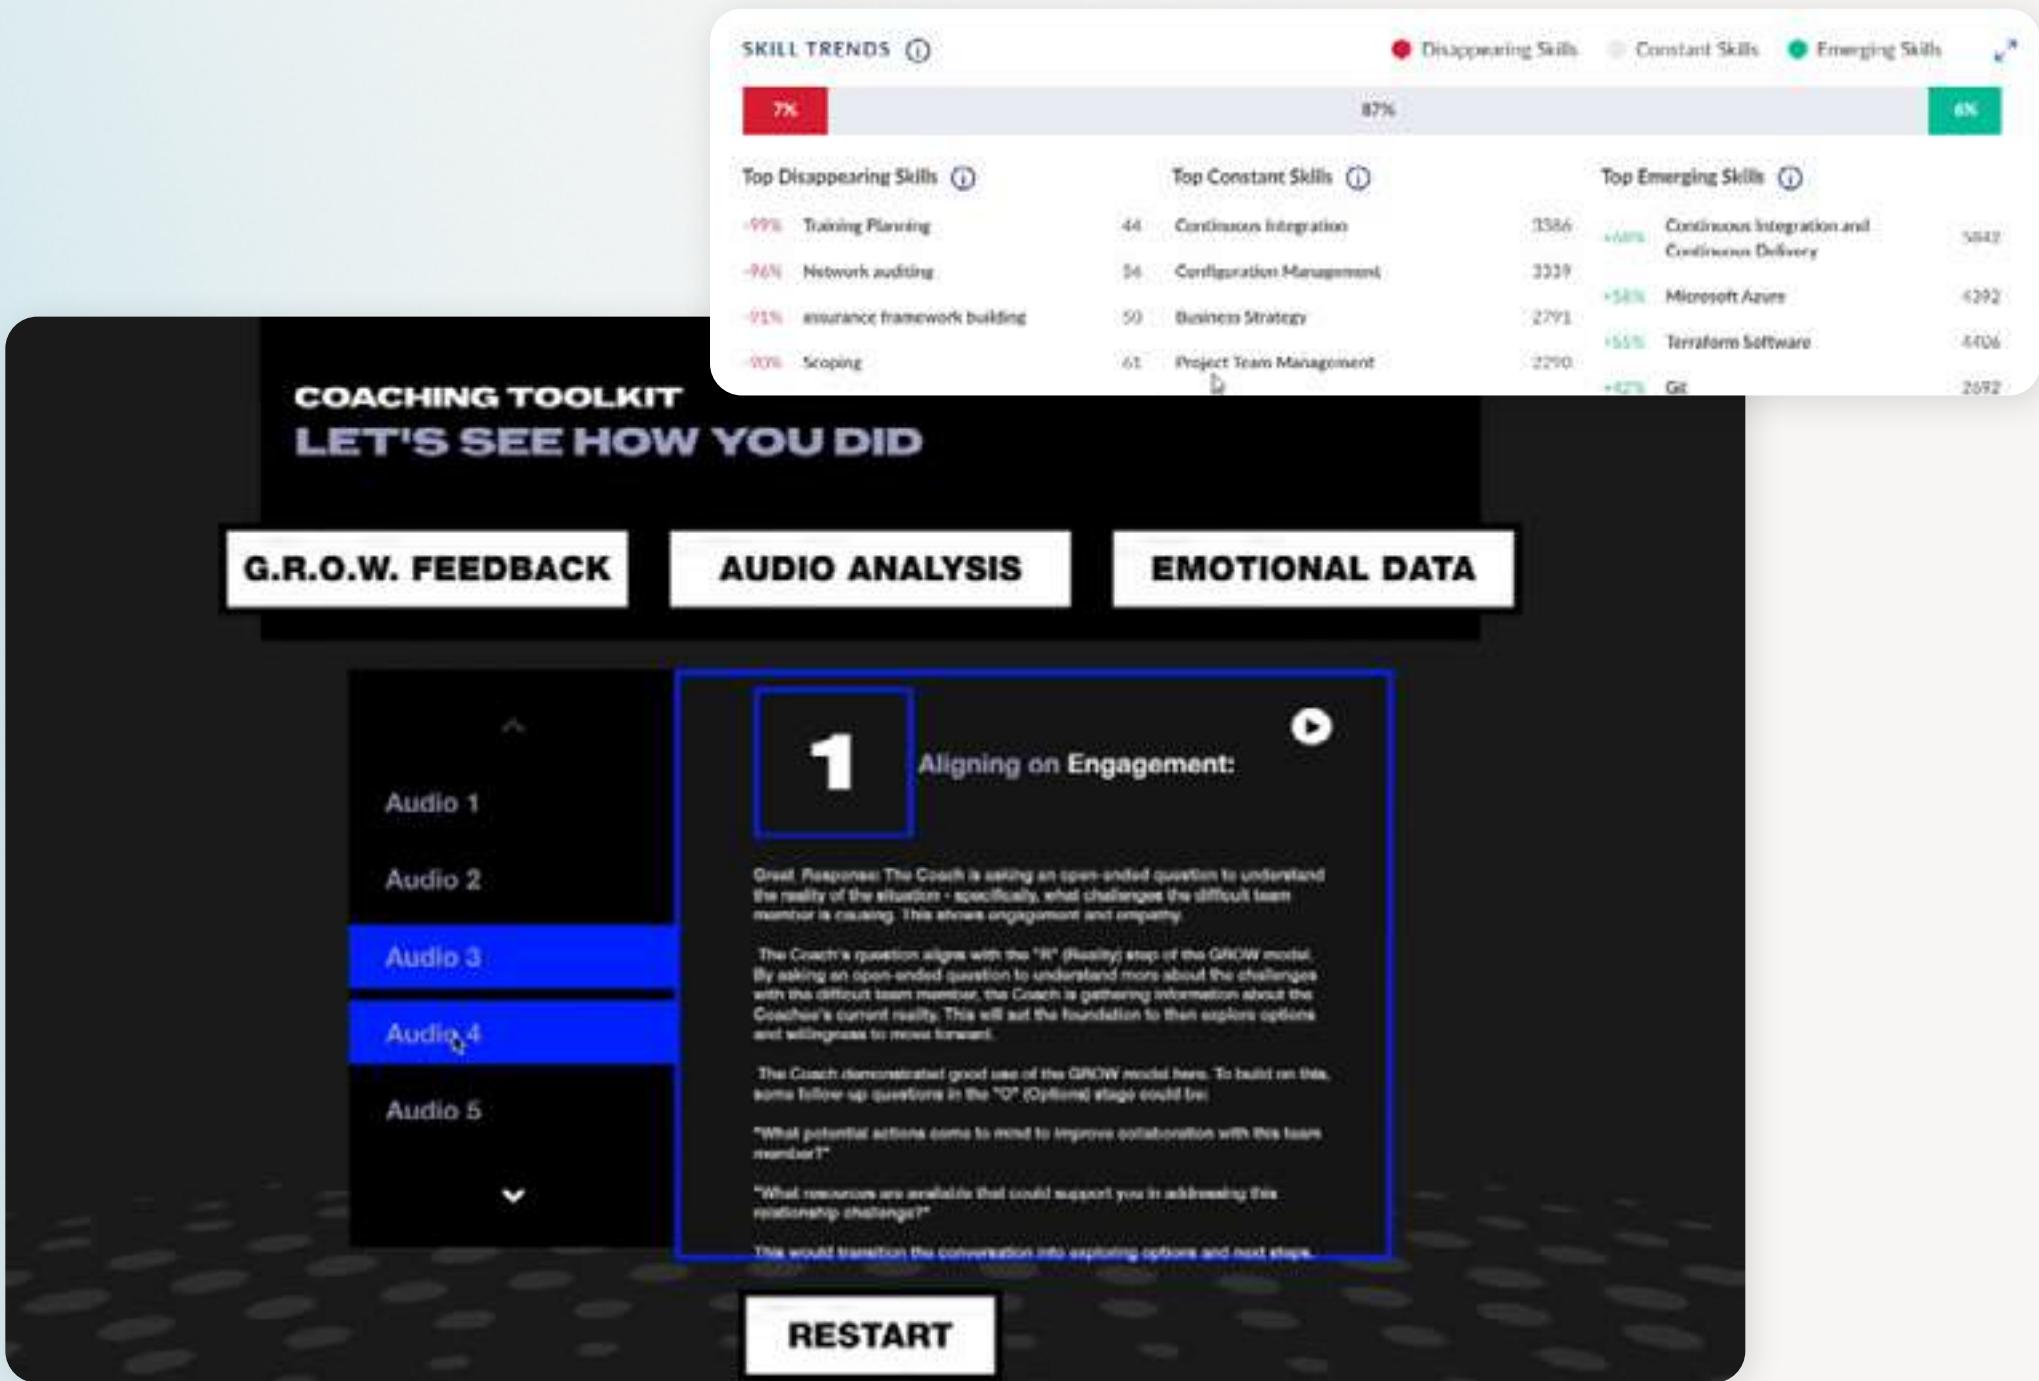
Word Choice

In a customer service training scenario, **learners respond vocally to callers and receive real-time coaching on whether the words and phrases they choose are operationally and culturally appropriate.**

We see all three—sentiment, intent, *and* word choice—measured by the [immersive coaching simulation](#) we explored in the previous section. In this experience, the AI engine analyzes the learner's spoken responses, recognizing their sentiment, intent, and word choice—all critical components of a human-centric coaching conversation.



Transforming Data into Actionable Insights



If a skill matters enough to coach, simulate, evaluate, and so on, it's important enough to measure. And thanks to the rich performance data generated by solutions that leverage machine learning and the AI-NLP pairing, in-depth learning analytics are now at your fingertips.

Here are just a few of the capabilities these data offer your team:

- **Personalize learning:** Track progress and tailor content to individual learners' needs and gaps.
- **Optimize content:** Analyze data to enhance training materials; for example, just-in-time materials to guide learners through a simulation where many get stuck.
- **Map and audit skills:** Pinpoint talent gaps and longer-term development needs to address in future learning strategies and roadmapping.
- **Demonstrate ROI:** Measure the impact of L&D programs and solutions on business objectives.
- **Ensure AI quality:** Continually analyze AI-generated content for quality and bias.

True data superpowers emerge when you combine AI and immersive learning technologies. **Read on to discover how AI-powered immersive experiences can live, thrive, and launch from your organization's LMS**—and even share the complex learner metrics of your choice.

AI-Powered Content Generation: Challenges and Solutions

A Guideline-Based Approach with the SweetRush Content Generator

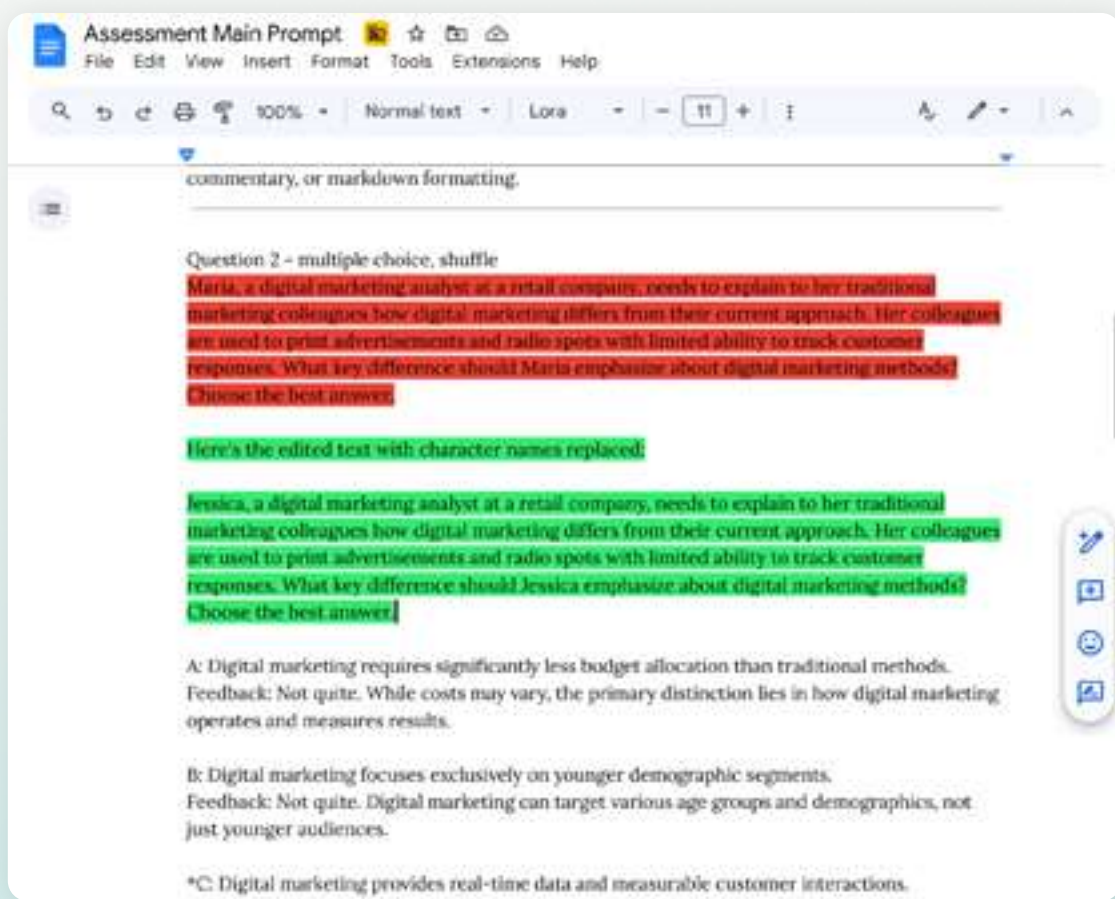
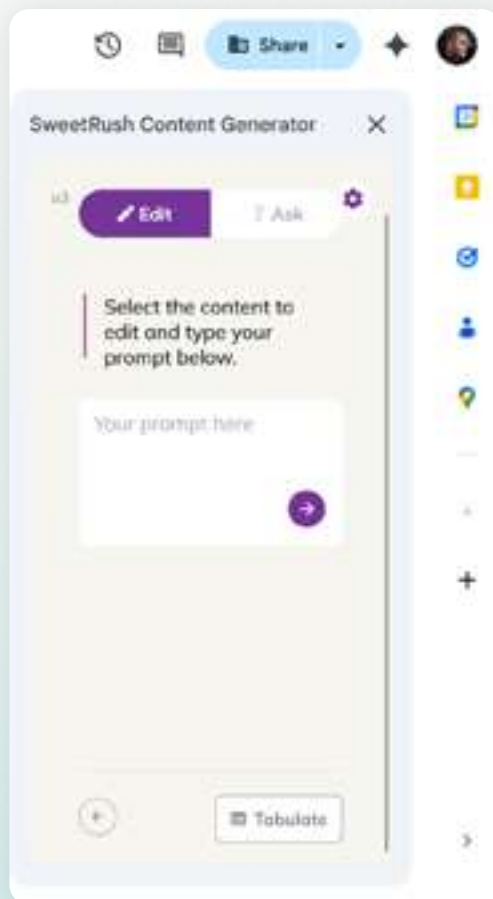
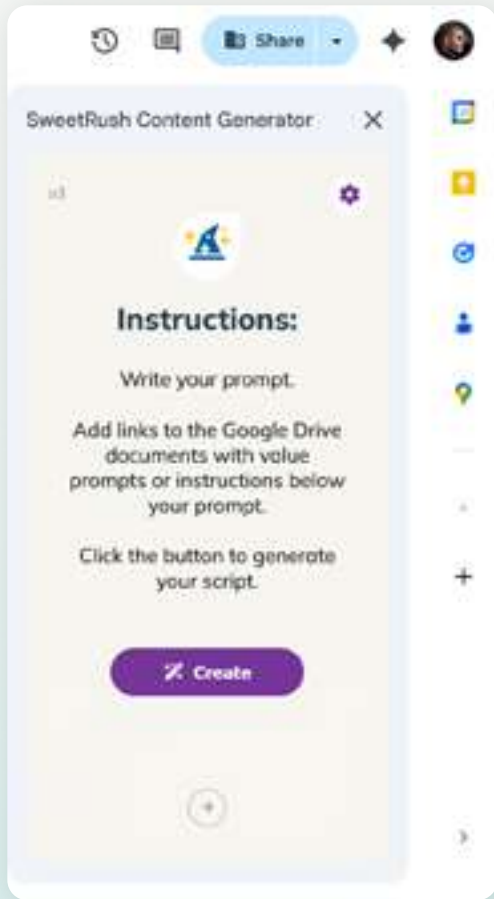


The Challenge

Traditional content generation tools rely on iterative chat interfaces where users engage in back-and-forth conversations to shape and refine content. Though this approach works well for one-off content needs, it can be inefficient when you need to generate multiple documents with consistent requirements.



The Solution: Guideline-Driven Content Generation



We've created an internal tool called the SweetRush Content Generator. **This Google Docs add-on is designed to generate precise, measurable content through a streamlined, guideline-based approach.** It enhances the Google Docs platform's natural collaboration features while adding powerful AI-driven content generation capabilities.

Thanks to the SweetRush Content Generator, our team can:



Create reusable guidelines

Establish prompts for structure, style, ethics, validation, audience targeting, tone consistency, accessibility standards, brand voice, regulatory compliance, and other content requirements. Because the Content Generator is so adaptable, we create reusable guidelines based on each client and their content generation needs (including our own internal style and standards).



Connect documents intelligently

Use Google Docs smart chips to link prompts with guideline documents, **making it easy to create multiple documents according to the same tone, brand voice, and other standards.**

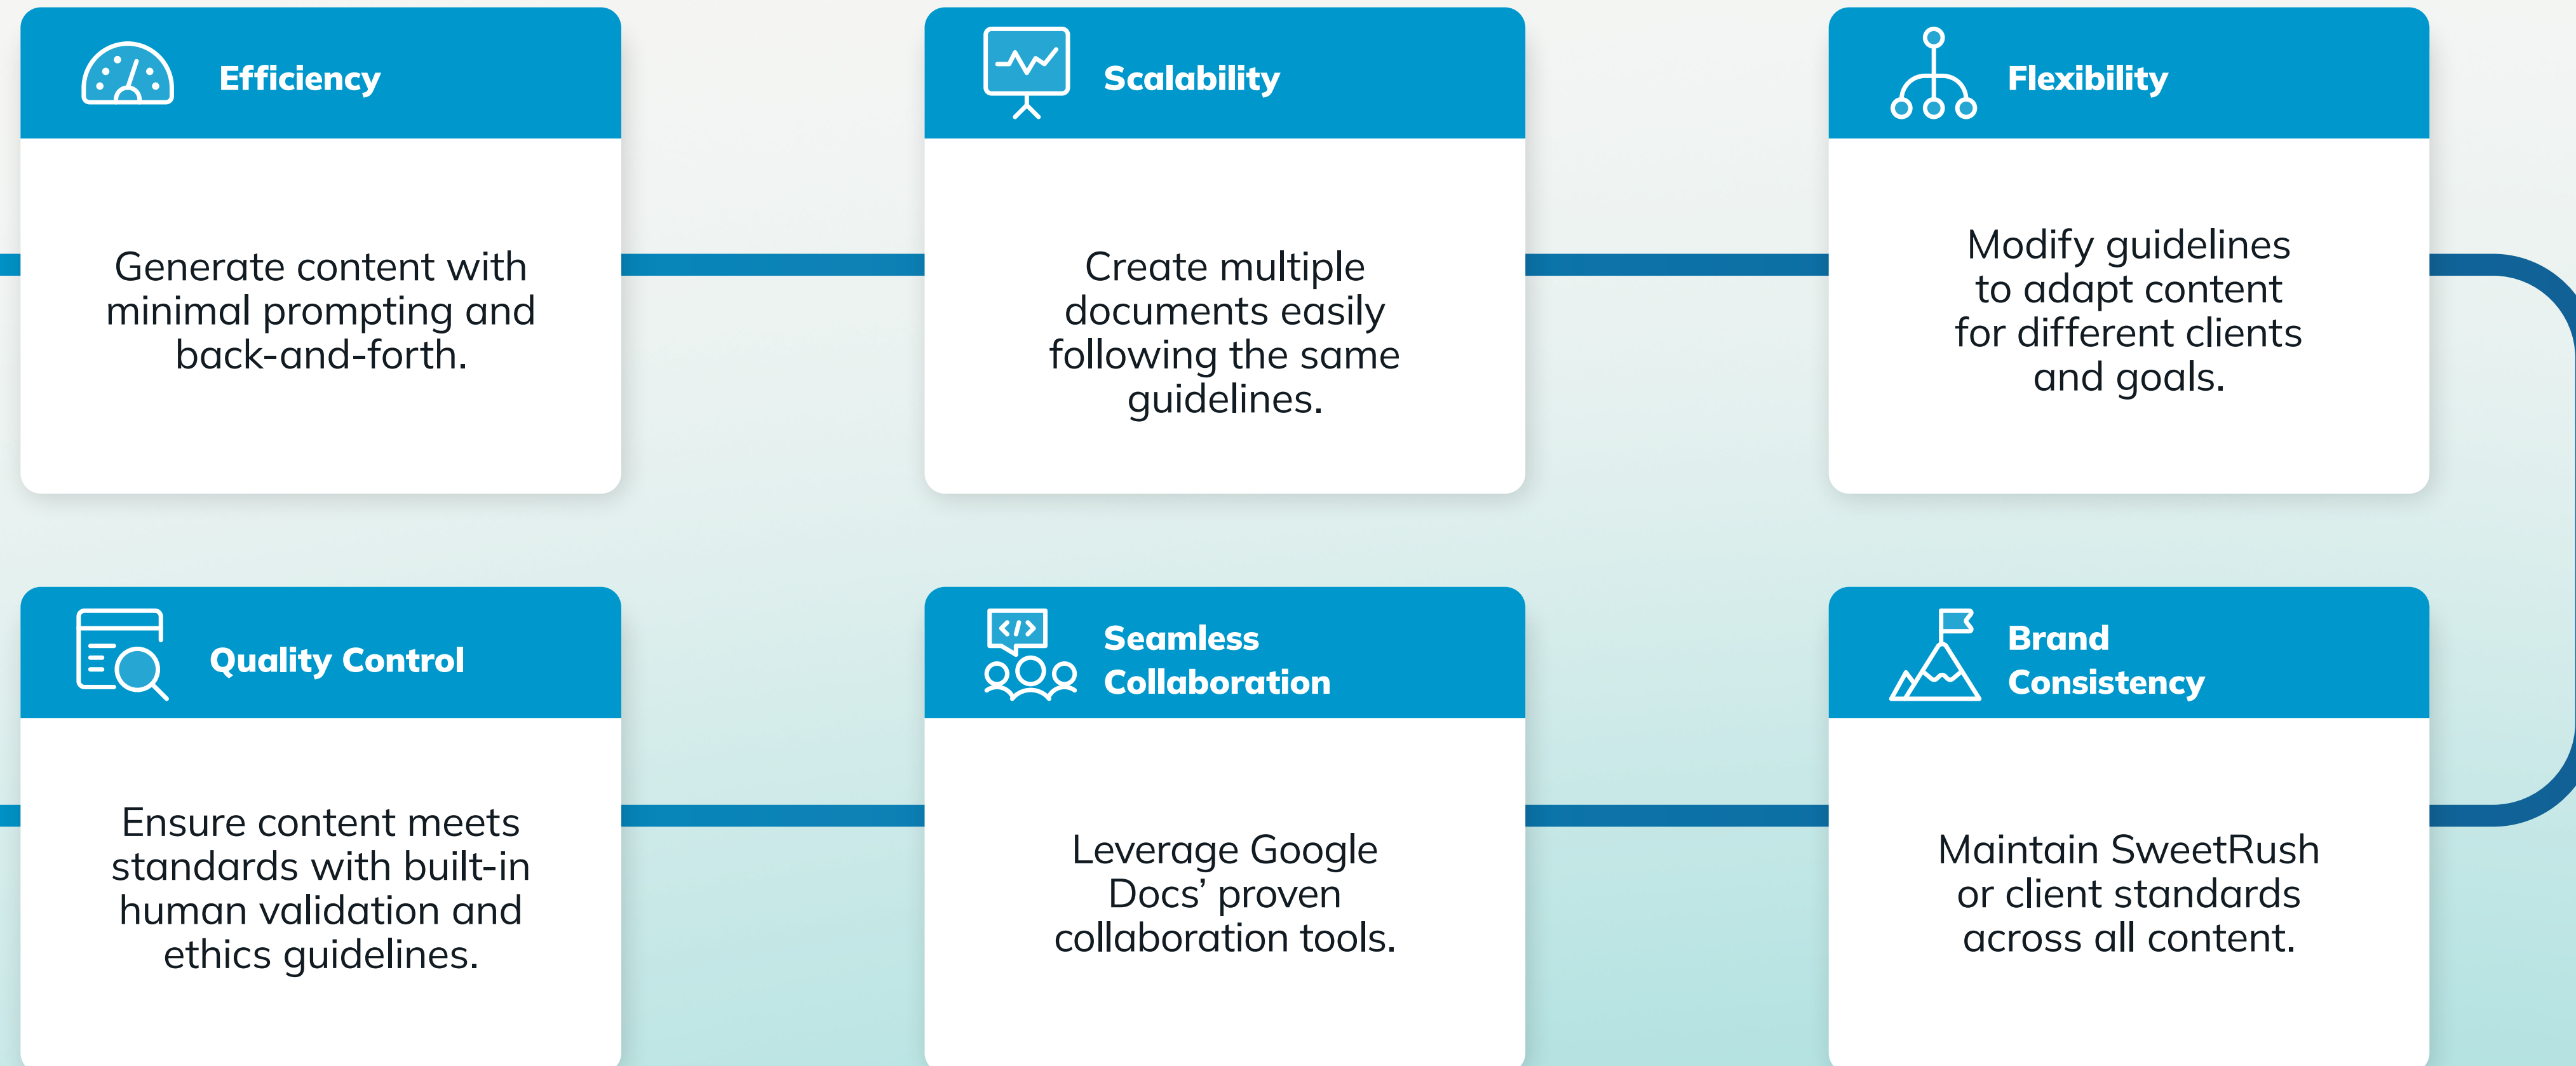


Generate consistent content

Produce multiple documents that follow the same standards of brand consistency and quality with minimal input.



Benefits of the Content Generator include:



Creating Better Prompts: Test, Tweak, and Maintain with TinyRAG

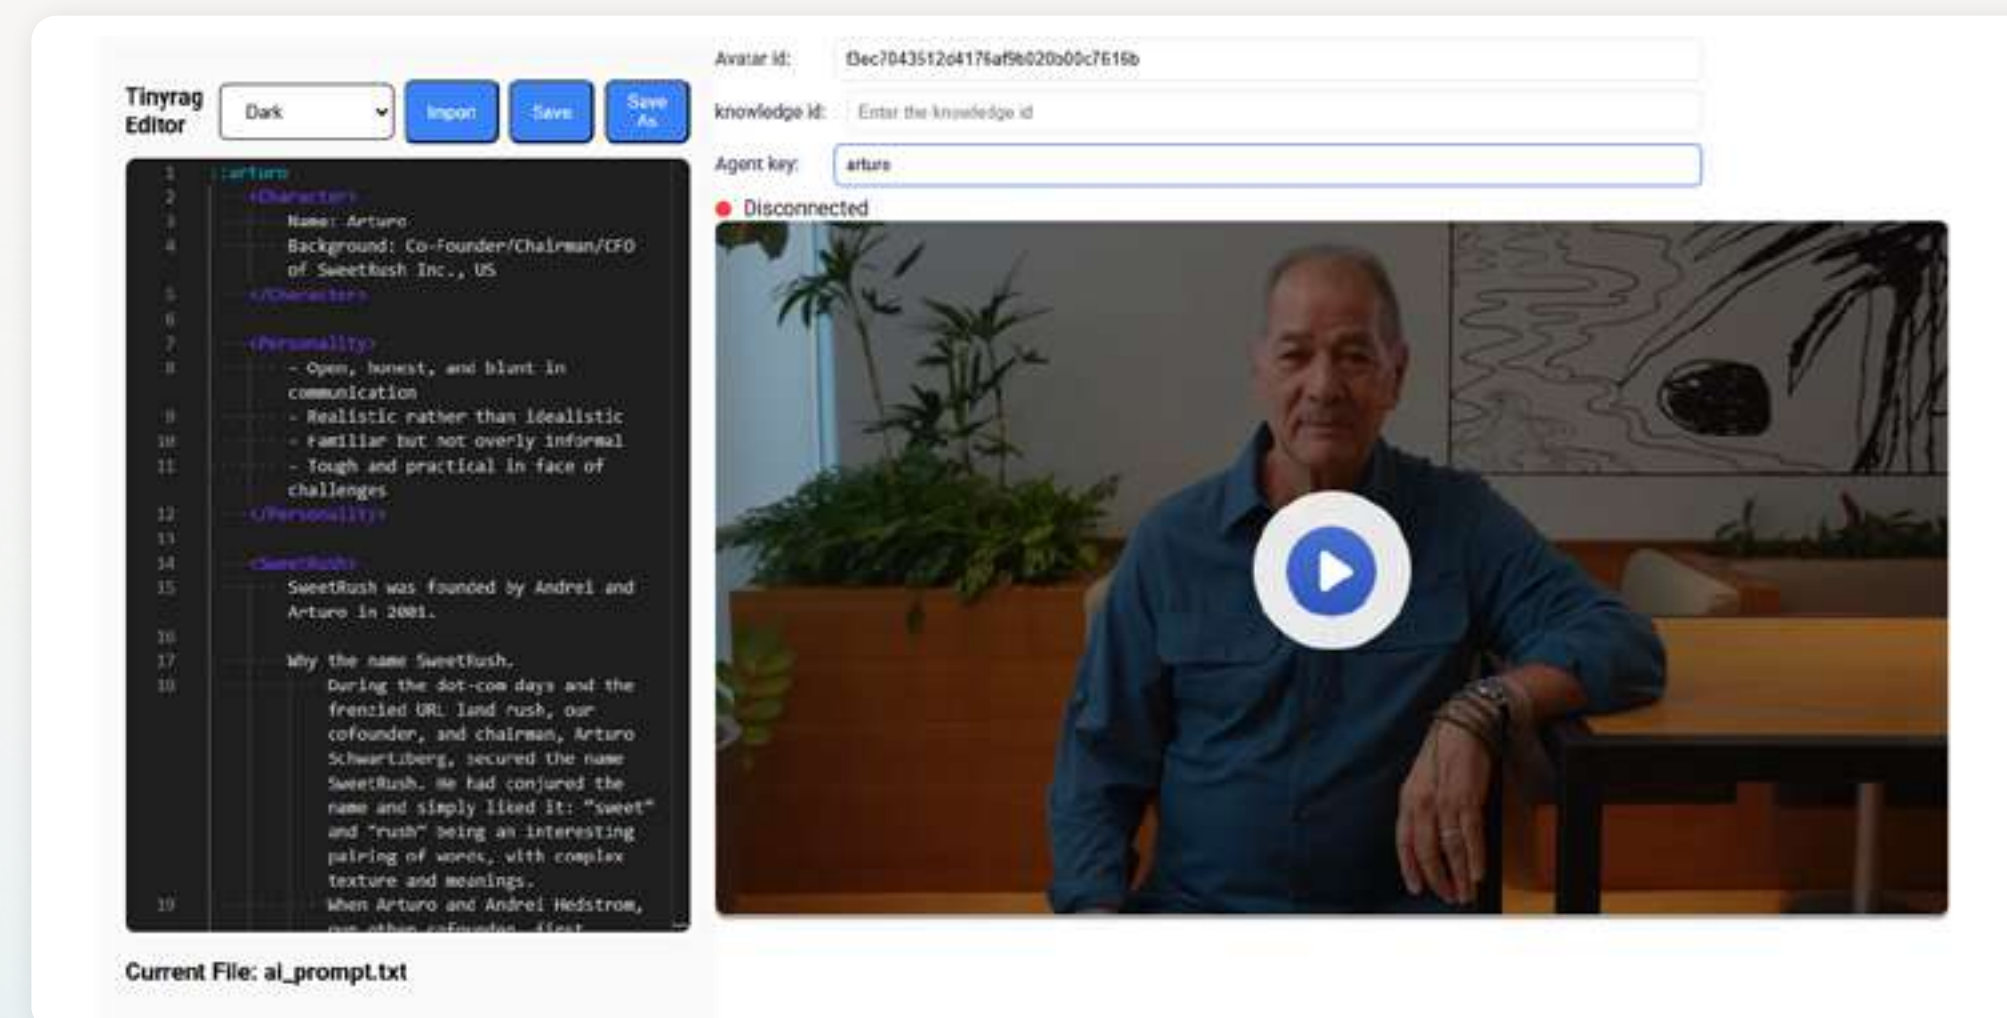


The Challenge

Wrestling with monolithic, theory-of-everything prompts is exhausting—and so is attempting to maintain, share, and tweak them across projects.



The Solution: A Swiss Army Knife for Prompt Development

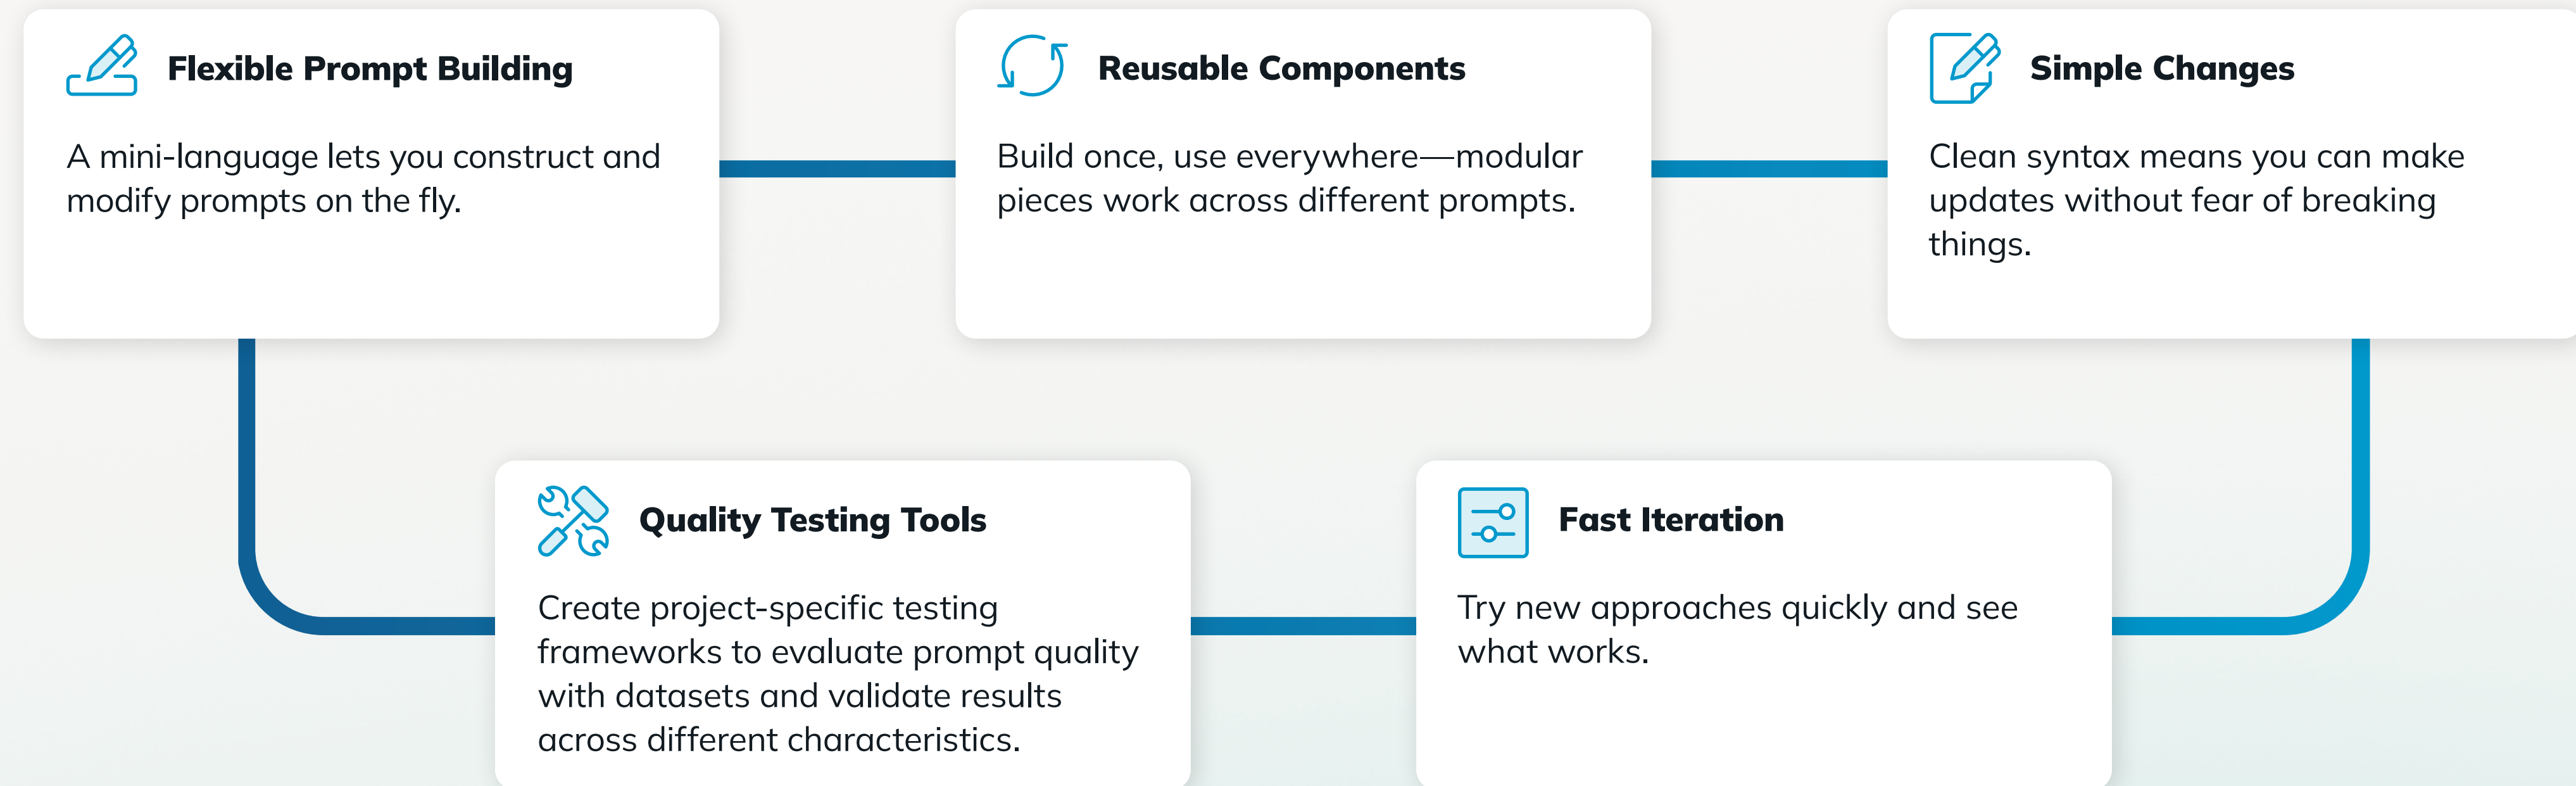


TinyRAG is a prompt engineering toolkit that offers a mini-language for building, testing, tweaking, maintaining, and reusing prompts. Think of it as a Swiss Army knife: compact, versatile, and designed to make your life easier.

TinyRAG offers a systematic way to build, test, and reuse quality prompts, enabling you to create tailored QA tools for different projects, maintain data sets, and evaluate results across various characteristics using the same system—all without drastically changing your workflow. It's prompt engineering that actually scales. You can tweak your prompts, reuse pieces across projects, and iterate quickly...without breaking everything.



Here's what's inside TinyRAG:




See the previous page for an example of TinyRAG. These examples show the background information, personality, and conversational style of an AI-powered chatbot and our own [Founder AI](#).



Next Up: Surveying the Possibilities

Peak AI doesn't just happen—it takes a great team, a solid blueprint, and the best building blocks at every level of the pyramid. Read on to discover how to start surveying your organizational landscape and laying the foundations for the next wonder of *your* world.



Surveying Your AI Landscape



You've got a vision for your AI strategy, from strategic foundation all the way to peak implementation.

But with an emerging and in-demand technology like AI, finding the expertise you need to bring it to life can be a challenge.

To expand their AI teams and toolkits, many L&D leaders turn to a partner like SweetRush who can help them survey, blueprint, and build their AI strategy from the ground up. We've worked with so many of your peers, and we'd love to survey the possibilities with you.

Here's a high-level look at the process.



A Partnership Based on Principles

Like any long-term relationship, **your AI partnership should be built on shared values...**and plenty of curiosity about your needs, learners, culture, and infrastructure.

We bring a unique blend of **strategic consulting and innovative learning design**, transforming complex AI principles and functionality into practical, relevant strategies that save time and add value.

SweetRush is committed to serving as a consultative partner, ensuring that your organization not only understands AI and implements it but also actively **integrates it into a high-performing, human-centered workforce.**

Read on for just a few of the principles we bring to the table.



Data Protection and Privacy

As we observed in [Scenarios 2 and 3](#), employees want to leverage the many AI tools in the marketplace to boost their personal productivity and efficiency. But these quick wins may carry a steep long-term cost: Vulnerabilities in how they use, store, and maintain our data can put our people, clients, intellectual property, and reputation at serious risk.

Here's the SweetRush promise:

Whether we deploy your AI solutions on-site or in the cloud, we honor your security and governance needs.

We've seen the dangers of ad hoc AI use—and how it imperils even those organizations that have officially closed their doors to AI. Let's level your organization up to [Frontier Firm](#) status...together.

Human Oversight and Ethics

What do pizza, video games, and Thin Mints have in common?

They're all great right out of the box.

You might have noticed that AI isn't on the list, and for good reason:

No AI tool is plug-and-play.


As learning researchers and scientists, we're continually examining new ways to use AI to extend the reach of human expertise, perspective, and insight. **We use exhaustive**

training, testing, and oversight policies and processes to ensure that AI outputs are accurate and free of bias. Our goal is to free you of these concerns so that you and your stakeholders can move to the next level of digital transformation.

Accuracy

In a recent client-exclusive roundtable, L&D leaders shared their No. 1 concern about AI: accuracy of outputs.

We get it! Accuracy is non-negotiable. As L&D leaders, we understand that even unintentional misinformation can erode trust and cause lasting damage. **That's why we commit to strategies that elevate the integrity of AI-generated content.**





SweetRush Lead Software Engineer **Andrés Villalobos** breaks down how the team applies these strategies in their work with large language models (LLMs).

What We Avoid: Blind Trust in AI

We never rely blindly on AI-generated content. Though LLMs offer us powerful capabilities, they are not always dependable sources of factual information. Instead, **we leverage their strengths to transform, synthesize, and organize data derived from trusted sources.**

What We Apply: Retrieval-Augmented Generation (RAG)

We use Retrieval-Augmented Generation (RAG) to ensure responses are grounded in real, high-quality information. RAG allows LLMs to retrieve context from a reliable knowledge base, enabling outputs that are not only relevant but verifiable. **With models capable of processing large volumes of tokens, we reference full documents rather than vague summaries.**

How We Ensure Accuracy: Fact-Checking AI Outputs

To reinforce accuracy, we fact-check all AI outputs using tools like Perplexity and GPTZero. This step guarantees that the insights generated by AI meet our high standards of accuracy, transparency, and ethical responsibility.

Enhancing AI Integrity: Practical Safeguards

Beyond fact-checking and retrieval techniques, we also apply a range of safeguards that help reduce risk and improve reliability. These include encouraging reasoning, revising outputs before delivery, and validating information against trusted sources. Together, these measures create a more transparent and dependable AI experience.

The Outcome: AI Built on Trust

Through layered verification and contextual grounding, we ensure that AI serves as a reliable asset, not a risk. Our commitment to rigorous quality control and ethical alignment empowers us to co-create learning solutions with you that are both cutting-edge and trustworthy.

In short: We don't just advocate for responsible AI—we implement it, validate it, and uphold it at every stage.

Tech Agnosticism

We don't sell hammers, so we don't see every solution as a nail—or as Lead Software Engineer Andrés Villalobos puts it, our AI strategy and solutions are **“built for you, not sold to you.”**

Working with a varied AI toolbox helps us recommend the most effective solutions for your organization's needs and existing infrastructure. **We'll work with the systems, tools, and policies you have in place now...and evolve with you.**

If that sounds complicated, don't worry—we love a challenge! In fact, we're known, in the words of one client-partner, for building leading-edge learning experiences that work “with, through, over, around, and under” existing technology ecosystems and infrastructures.

Value

Value is a priority for you, your business partners, your leadership—and for us, too. The SweetRush team loves to join L&D leaders like you in the ongoing challenge to create quality learning within ever-tightening budgets and timelines.

[Secure tools](#) like TinyRAG and our SweetRush Content Generator help us build value into every project plan so that we can maximize what matters most to you.

At the organizational level, **we're committed to helping you capture the value of strategic AI use at scale, avoid costly risks, and maximize your impact.** And there's no time like the present: A recent global survey shows that organizations that invested in GenAI in 2023 are already [seeing a return](#) on their investments.

Learning-First Approach

With nearly a quarter-century of experience in learning strategy that drives behavior change for world-class organizations, SweetRush brings a learning-first approach to every solution. **We're honored to be named No. 1 in AI tools, immersive learning, experiential learning, custom content, upskilling and reskilling, and employee engagement...and much more.**

Subject Matter Expertise in AI and Beyond

We've also been named the No. 1 Staff Augmentation provider by both eLearning Industry and Training Industry for our extensive talent pool of subject matter experts (SMEs) in AI, immersive learning, live experiential learning (LEL), data analytics, and certification programs. That means the same SMEs who make your projects shine are also available on a contract, temporary, or long-term basis to help you answer mission-critical questions about:

- **Which AI strategies to pursue and how to integrate AI into your organization's existing infrastructure and culture**
- **Policies and processes to protect data and intellectual property**
- **A plan for continuous human oversight and testing**
- **AI learning experience design and use cases**



How do you know whether you need a single expert or a full team?

SweetRush Director of Talent Solutions Rodrigo Salazar-Kawer describes the difference in builder's terms:



If you want to remodel your home, sometimes hiring an individual, such as an electrician or plumber, is the right thing to do. However, if you have a larger project that requires a team effort and lots of coordination, then you will likely want to hire a contractor and a crew.



If you're curious about what a holistic, ethical, and human-centered AI practice could look like for your organization, we can help you explore! [Reach out to share your needs with us and discover how to transform your AI vision and dreams into an actionable blueprint.](#)

Conclusion: Carpe Diem...with Caution

We hope you find the opportunities to extend human expertise, expand the capacity of your teams, and scale high-value skilling experiences to a global audience as inspiring as we do.

Equally, we hope you're feeling the urgency.

Our current AI call to action isn't *just* a carpe diem call. It's also a call for *caution*. Now is the time to spearhead a proactive and strategic approach to AI governance, implementation, and usage.

Doing so can help us move toward [Scenario 1](#), a thriving AI ecosystem that blends machine intelligence with human judgment and oversight—and offers enterprise-wide training on AI guidelines, ethics, and role-based AI skills and tools. In this scenario, AI isn't just safe; it's optimized and already delivering ROI.

If life in your organization sounds a lot like Scenario 1, congratulations! Your strong foundation and enterprise-wide adoption set you and your team up for optimal innovation and growth. You're living at the pinnacle of Intelligent L&D Content and Design and ready to iterate, refine, and repeat...in fact, you're probably a veteran innovator with stories to share!

If, like a majority of L&D leaders, your landscape looks more like [Scenarios 2](#) or [3](#), the good news is that your people have the curiosity and initiative to experiment with AI tools on their own. However, ad hoc adoption and outright banning share several downsides, including missed opportunities to leverage AI at scale and increased risk of jeopardizing data, intellectual property, relationships, and reputation.



We've said it before: *AI is already in the building.*

And like any power, it presents us with both immense opportunity and even greater responsibility. As tech futurist [Peter Leyden](#) aptly states, as L&D leaders, we owe our people and our businesses more than just "tracking the inevitable"; we need to help them shape the game.



AI has arrived, it's here to stay, and we can now just track the inevitable as it scales up and begins to fundamentally change the economy and society. The debate is over. It's now game on.



— [Tech Futurist Peter Leyden](#)

Let's imagine a world where AI isn't wandering aimlessly around the building or sneaking around like a cartoon villain but, instead, purposefully engaged as a muse or co-creator that sparks innovation and helps us achieve what was once unimaginable. Picture it helping teams connect more deeply with clients, collaborate creatively with one another, and tackle complex strengthen that stretch and challenge their human ingenuity.

As we answer the dual calls for caution and adventure, let's savor our unique ability to lead in this pivotal moment and ensure a humane, responsible, and practical approach to AI adoption. As L&D leaders, we have the unique ability to share our optimism while also driving home the message that **there is no sitting this out.**

We'd love to help you leverage AI's full potential for good, drive innovation, and prepare your people for a future where their ingenuity and human-AI collaboration lead the way. [Reach out](#) to talk strategy and next steps.



Warmly,

Tiffany Vojnovski
Senior Learning Strategist

[in /tiffanyvojnovski/](#)





SWEETRUSH®

29 Gold Awards!

*And One Silver
(Just wait till next year)*





Let's win awards together!

Impactful learning programs are the result of true collaboration.



Chief Learning Officer

12 Gold

Learning in Practice Awards

Content, eLearning, Technology
Innovation, Community Service,
Blended Learning



eLearning Industry

Chosen #1

AI Expertise, Immersive Learning,
Simulations Training, Upskilling and
Reskilling Employees, Staff Augmentation



Training Industry

Top 20

Custom Content, Experiential Learning,
Staffing and Temporary Resources

In Our Clients' Words



SweetRush makes my life easier.

SweetRush excels at finding and providing true expertise and in nurturing their team members, providing them with the necessary support to thrive, while also employing strategic resourcing practices. This has made my life easier in knowing I have a partner that is deeply committed to the project's success.



—Program Manager, IT and Technology,
10,000+ employees



Great thought leaders willing to push, and be pushed, in new and exciting ways.

Excellent and responsive team, great innovators, inclusive leadership. All of this has made me and my team feel at home, as if SweetRush were an extension of our internal team members.



—Product Manager, Online Education,
10,000+ employees



With a tight deadline, the SweetRush team was great to work with.

They listened to our asks and executed quickly to a high standard. We've given trainings on this topic before, and we've consistently heard that this is the most comprehensive, best-designed, and easiest-to-follow training.



—Product Manager, Online Education,
10,000+ employees

We hope you're feeling inspired!

Let's talk strategy.



SWEETRUSH®



www.sweetrush.com
info@sweetrush.com

