

Last update: August 12, 2026

At Schoox, we are pioneering a new standard of responsible, ethical, and transparent use of AI functionality in the Learning & Development industry. Our commitment to our customers goes beyond buzzwords and compliance. Our mission is to provide you with AI tools that help your organization train faster, perform better, and drive measurable business impact—while keeping you fully in control of your data.

Core AI Principles

Customer Content is Never Used to Train AI Models: At Schoox, we do not and have never trained AI functionality on customer content. Our AI functionality is trained on licensed datasets and information available in the public domain. Such functionality is rigorously tested and refined using proprietary methods developed by Schoox's internal team.

Built for Trust: Enterprise-grade security and privacy by design are the foundation of Schoox's platform. Personal data is not used in any of Schoox's AI models.

Human-Centered: AI functionality supports your people—it doesn't replace them. Our platform enables L&D leaders, instructional designers, and Admins to achieve greater business impacts more efficiently without eliminating human expertise.

Fully Configurable: You decide how, where, and when AI is used. AI functionality in Schoox's platform is always optional and fully configurable by Admins at multiple levels of system settings. You control which AI tools are enabled and what data is processed.

Transparent & Explainable Outputs: Schoox's AI functionality provides clear, contextual reasoning so that Admins can evaluate, edit, or override outputs as needed. For example, the functionality explains its reasoning for specific skill suggestions and course recommendations, which turns the "black box" of AI into an understandable tool that Admins can customize as needed. There's no guesswork – just clarity and results that are relevant to your business.



Responsible & Ethical Use: We are committed to preventing bias and harmful outputs in our AI models. We regularly audit our systems using anonymous data to verify that outputs align with ethical standards.

Continuous Innovation & Improvement: Relentless innovation is in our DNA. Our platform enables users to easily rate AI output and provide voluntary feedback, which we use to refine and improve the performance, output, and quality of our AI systems.

Obsessed with Your Success: Your success is more than a priority – it's our obsession. To ensure you always have a voice, we maintain a Customer Advisory Board, Customer Community, and a dedicated Customer Success Team to collect real-time feedback and align our products with evolving customer needs.

Security & Privacy You Can Trust

Your content remains yours—always private, always protected.

Zero Data Sharing: Customer content is never used to train external models or shared across customer Academies.

Secure Architecture: Schoox's AI functionality combines industry-leading enterprise models with private, secure data retrieval. We leverage pre-trained enterprise-class large language models (LLMs), including OpenAI's GPT-5 and Amazon Bedrock (RAG), that provide logical segregation of data across client instances, ensuring your content is never exposed outside of your Academy.

Enterprise Protection: Schoox's platform is SOC 2 Type 2 compliant and utilizes data encryption in transit (TLS 1.2+) and at rest (AES-256).

Built-In Safeguards: Our AI functionality employs a multi-layered proactive approach to prevent prompt injection, unauthorized access, and other misuse. Such protective measures include input sanitization, guardrails, role-based constraints, multi-factor authorization, and template-locked prompts that individual users cannot directly manipulate.

Regulatory Alignment: Schoox's platform is aligned with GDPR, CCPA, and other global data privacy, security, and AI regulatory frameworks. We also provide clear, accessible documentation about the use, storage, and retention of customer and user data. For more



information about Schoox's security and compliance programs, please visit the Trust Center at www.schoox.com/trust-center/.

Dedicated Governance: Schoox maintains an internal AI Governance Committee that oversees all AI deployments, conducts rigorous risk assessments, and evaluates every AI feature to ensure alignment with industry best practices, legal and regulatory obligations, and ethical commitments.

AI Functionality in Schoox's Platform

Schoox and the Learning Impact Suite uses AI to transform your L&D data into a continuously evolving learning ecosystem where every course, skill, and training activity connects to performance improvement and measurably impacts your bottom line. Some of the capabilities in the Schoox core platform include:

Content Creation: Generate full-length and/or micro-courses, assessments, reinforcement exercises, and L&D visual assets, including video and audio training materials, achievement badges, and platform icons that promote and maintain brand consistency throughout your organization. You may configure the AI functionality to generate content based on existing materials in your Academy only, or using publicly available resources outside of our Academy, or a combination of both.*

Content Delivery: Optimize learning engagement and knowledge retention through AI-recommended course delivery schedules.

Skills Mapping: Intelligently map job roles to skills, and skills to business impact areas to identify key development opportunities that drive performance.

Micro-Assessments & Adaptive Reinforcement: Conduct real-time skill assessments to personalize each learner's training plan and create targeted micro-courses for continuous, low-friction reinforcement that sharpens skills and improves retention.*

Learning-in-Action (LINA) Chat Assistant: Answer staff questions in real-time using information that is available to users from Academy materials. L&D leaders can leverage data from the AI assistant to develop insight into training programs and opportunities.



Some of the capabilities in the Learning Impact Suite include:

Enterprise-Wide Training Plans: Automatically analyze training data, content and goals, identify gaps, and build curated learning paths and course recommendations for your entire organization.

Skills Clustering: Identify the skill combinations that optimize proficiency and directly impact your defined business metrics.

Content Gap Analysis & Remediation: Evaluate your content library, assess whether existing courses address critical skill gaps, and generate new full-length and micro-courses with a few more clicks.*

*Course creation functionality based on third-party licensed content may be limited to the extent prohibited by such third party. For more information, please see Schoox's AI Functionality Policy or contact your Schoox Customer Success Manager.

AI Architecture & Models

Schoox's AI functionality leverages the following models to deliver powerful generative AI capability that can be customized to each customer's organization while maintaining critical enterprise-grade security:

GPT-5: Content generation and general AI tasks leveraging publicly available resources via API

Amazon Bedrock + Retrieval-Augmented Generation (RAG): AI functionality securely using customer-specific data and content

Benefits of the AI models:

- Customer content is not used to train any AI model.
- AI models provide enterprise-grade security and protective guardrails.
- State-of-the-Art generative AI capability is delivered at scale and speed.
- AI models are customizable based on rapidly evolving skill requirements, content development, and business priorities.

How does Schoox prevent prompt injections?



Schoox mitigates prompt injection by design, not just detection, by implementing multi-layered preventative measures, including input sanitization, guardrails, role-based constraints, RAG content controls, and structured prompt engineering. Users cannot directly manipulate prompts because generative AI functionality uses system-generated, template-locked prompts with embedded instructions and output constraints that are only configurable by an Admin. Additionally, users may only access permissioned content for RAG features, and responses are grounded in verifiable, read-only sources. Role-Based Access Control (RBAC) further restricts content generation actions to authorized roles. All prompt interactions are logged for Admin review, and suspicious patterns trigger rejection or human review by an Admin.

How did Schoox build the AI functionality?

Schoox's AI architecture utilizes the enterprise-grade AI models described above. Such models enable structured prompt engineering and validation pipelines. Amazon Bedrock further provides Retrieval-Augmented Generation (RAG) for grounded output based on internal materials without exposing such content to external LLMs or sacrificing security. Schoox's training of such functionality is limited to licensed datasets and information available in the public domain only. AI functionality is never trained on customer content.

Can customers modify the underlying algorithms in the AI functionality?

No, customers cannot directly modify the underlying algorithms. However, customers maintain extensive control over AI outputs and configuration, which allows Admins to manage the application of AI functionality in their Academy, including by disabling access to certain features. Admins can also override, edit, or regenerate any AI-generated output such as:

- Content Creation (full-length and micro-courses)
- Skill Mapping & Clustering
- Micro-assessments and Skills Remediation
- Learning Path & Training Plan Recommendations

The ability to override AI suggestions enables outputs to be customized and aligned with your organization's specific standards, brand, and learning objectives.

How do customers provide feedback?

The AI functionality enables multi-level feedback using:

- Feedback Capture: Users can rate responses, flag issues, and provide qualitative input.



- Feedback Routing & Analysis: Schoox’s internal teams will log, categorize, and review feedback.
- Prompt & System Optimization: Insights will inform improvements to prompt logic, instructional templates, and retrieval context weighting.

Such process allows for continuous improvement of the quality, precision, and contextual alignment of AI outputs through prompt optimization and structured review cycles. In addition, Schoox leverages human quality assurance, structured audits, and customer collaboration to refine the quality of AI output and ensure the functionality meets enterprise-grade standards.

How does Schoox protect against hallucinations?

The AI functionality minimizes hallucinations through upfront controls, including prompt scaffolding and strict generation boundaries, and uses RAG for grounded content. In addition, built-in validation framework verifies output for accuracy, depth, and structure through scenario validation, sequencing rules, and quality checks. If issues are identified, a post-generation remediation pipeline regenerates flagged content and adjusts prompts. Governance and role-based controls ensure only authorized users can generate content with traceable outputs.

How does Schoox measure and monitor accuracy, bias, and performance?

We use a combination of automated validation, human oversight, telemetry logging, and customer-driven voluntary feedback to measure and monitor the accuracy, fairness, and general performance the AI functionality. We continuously monitor performance across generation success rate, output rejection rate, and time to output. We also track content generation, skill-to-KPI link strength, user engagement, and security measures.

Can Schoox’s generative AI functionality explain its reasoning?

Yes—while the foundational models do not inherently “explain” their reasoning in the way a human might, Schoox surfaces structured, contextual reasoning around many generative AI-powered outputs. Our architecture combines prompt engineering, rule-based decision layers, and explanation frameworks to deliver transparent, explainable outputs, especially in high-impact business scenarios. Reasoning is transparent across multiple use cases within the platform, including skill recommendations for job roles, course suggestions aligned to business goals, KPI-to-skill mapping, assessment design, and the source attribution behind Learning Chat Assistant responses.

What are the intellectual property (IP) implications of using Schoox’s AI functionality?



The use of AI functionality is governed by the Service Agreement between Schoox and the customer, which incorporates Schoox's General Terms of Service and AI Functionality Policy. In accordance with such Terms and Policy, as between Schoox and a customer, the customer retains ownership and control over all AI-generated output, except for Licensed Material provided by Schoox or a third-party content provider (see the AI Functionality Policy for more information). Schoox does not claim ownership of AI-generated output constituting customer content. Additionally, separate contracts with third-party providers of the underlying AI models are not necessary in order for a customer to access Schoox's AI functionality.

Does Schoox maintain formal policies regarding the ethical use of AI?

Yes, Schoox adheres to strict ethical guidelines regarding the use of AI functionality within our platform. Our AI policies focus on transparency, fairness, user privacy, and customer control over whether, when, and how AI functionality is used within their Academy. We proactively prevent biases in our AI recommendations and regularly audit our systems to maintain ethical standards. We also emphasize the responsible deployment of AI functionality to support human decision-making rather than replace it, ensuring that our AI functionality is an assistive tool that enhances human judgment.

Why Schoox Leads the Pack

Our goal at Schoox is to set the industry standard with a comprehensive responsible AI framework covering security, privacy, ethics, and governance—going beyond typical policies to deliver true transparency and control to our customers.

