WHITE PAPER

Deliver Personalized Learning At Scale With Adobe Learning Manager's New AI-Based Recommendation Engine

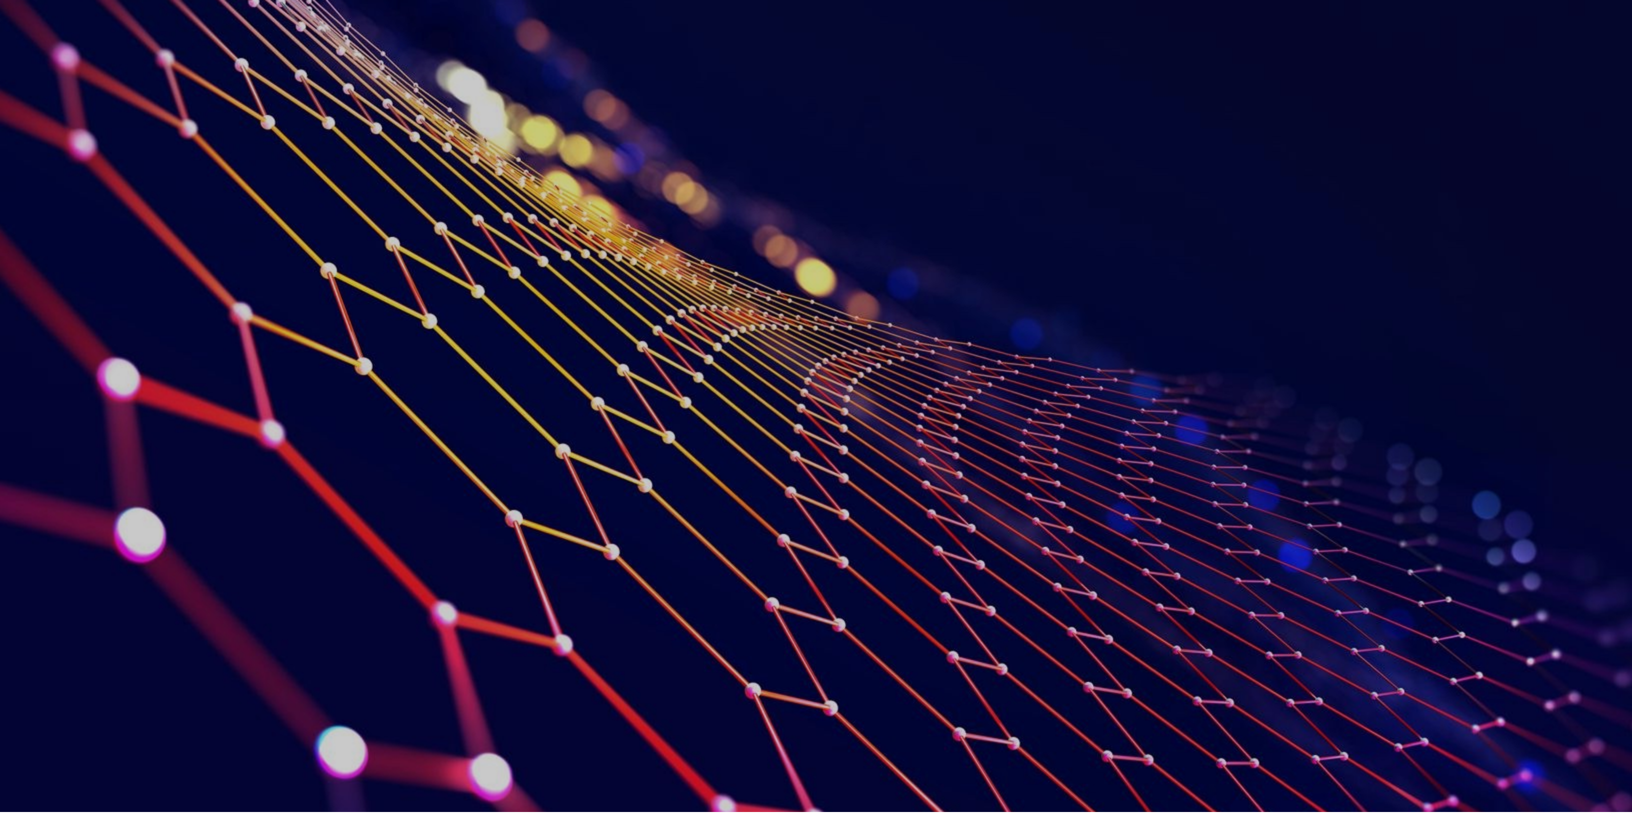

Introduction

If you are an organization looking to invest in customer or partner education, you likely have a set of complex products or services portfolio that caters to multiple user personas/roles. Further, it is likely that your users have different levels of proficiency, such as beginner, intermediate, and advanced, playing those roles in their work. In such a scenario, how do you construct an education experience that boosts self-driven training adoption, thereby leading to increased product adoption?

Adobe Learning Manager's new AI-based recommendation engine is designed to do precisely this—personalize the learning experience for your customers and partners. It does this by using machine learning algorithms behind the scenes to create (and keep afresh) a 'Learning Profile' for every individual and a 'Content Profile' for every piece of content on the platform. Other AI algorithms are then used to match, in real-time, learner to content so that only the courses/learning paths that are most relevant to an individual and have been upvoted by their peer's actions, are recommended.

Table of Contents

How Does Learning Manager's New Recommendation Engine Upgrade Your Learner's Experience?

With the new recommendation engine, when a learner comes to the platform, Learning Manager's AI analyses and associates the learner with a set of products, roles, and levels. There are multiple ways by which the recommendation engine achieves this. Metadata provided by admins, prompting learners to input their learning preferences at key intervals etc. are a few examples of explicit data collection. Then there are implicit methods to deducing this information–for example, by parsing the learner's activity on the platform, importing metadata from 3rd party systems like a CRM or a CDP.

The content profiles of all the pieces of content available on the platform are then used to match and shortlist relevant content for that specific learner from the hundreds of courses/learning paths that may be available in the account. This is the first step. In the second step, the recommendation engine uses a ranking algorithm. The algorithm uses Learning Manager's proprietary 'Score' for a course/learning path to rank and sort relevant content based on the likelihood of enrolment and completion of the content. This ranked list of content is then displayed in the recommendations strips on the learner homepage (and available via APIs for headless LMS environments.)

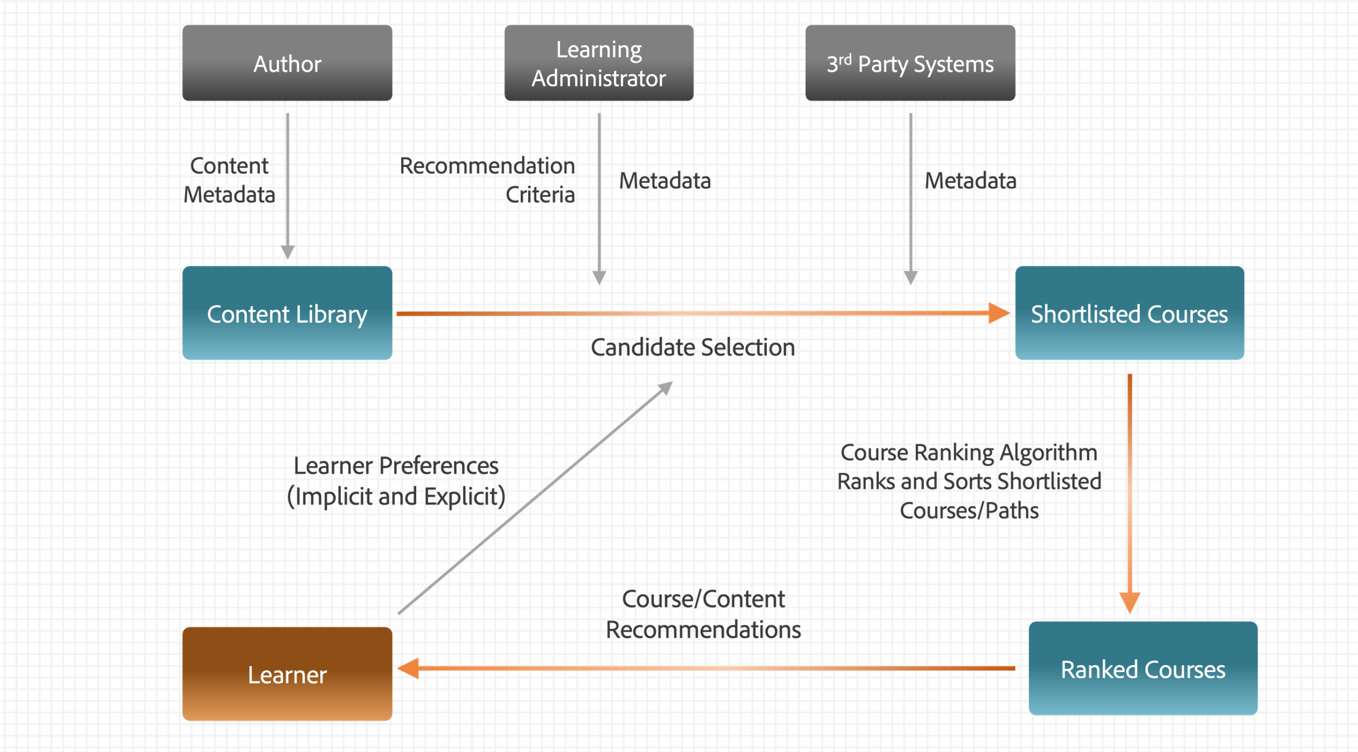

Figure 1: Recommendation engine workflow visualization

The **Course Ranking Algorithm** is at the heart of the new recommendation engine. The algorithm has used **50 million data points** and **five years of aggregated learning data across millions of users** to score and rank courses based on their **likelihood of enrolment and completion.** To ensure confidentiality and security, this analysis was conducted solely with anonymized data. The only objective was to uncover universal "content features" that serve as predictors for enrolment and completion.

Key Features of the New AI-Based Recommendation Engine

1. Configurability for Learning Admins

Let us take an example of a SaaS company that provides custom IT solutions to banks. The company has a diverse portfolio of solutions which include fraud detection systems, secure cloud storage, data analytics tools, loan origination systems, and so on. At a customer bank's end different roles may be associated with the solutions purchased by it, for example the bank's own IT team, UI designers, data scientists, loan sales reps and so on. Further, these roles might have different levels of competency such as beginner, intermediate and advanced.

As a learning leader of this SaaS company focused on improving the customer education experience, you need to provide personalized learning to your customers based on the products they have purchased, and roles + levels applicable to them.

Learning Manager's new AI-based recommendations engine enables configuring learning recommendations based on the parameters—Products, Roles and Levels (PRL). The parameters "Products" and "Roles" can be renamed to adjust to an organization's needs. For example, "Products" can be renamed to "Topics" and "Roles" can be renamed to "Regions."

2. Balance Between Configurability and AI-Driven Dynamism

Learning leaders know best what parameters will drive successful business outcomes in their learning platform. Learning Manager's new recommendation engine strikes a right balance between configurability and the dynamic nature of AI-driven recommendations.

- The Product, Role, Levels parameters are defined and configured by the learning leaders/admins.
- Learners provide their product, role and level preferences.
- Learning Manager's AI then analyzes learner and course profiles to match and shortlist content, rank courses, and feed highly tailored content in the recommendation strips on the learner homepage.

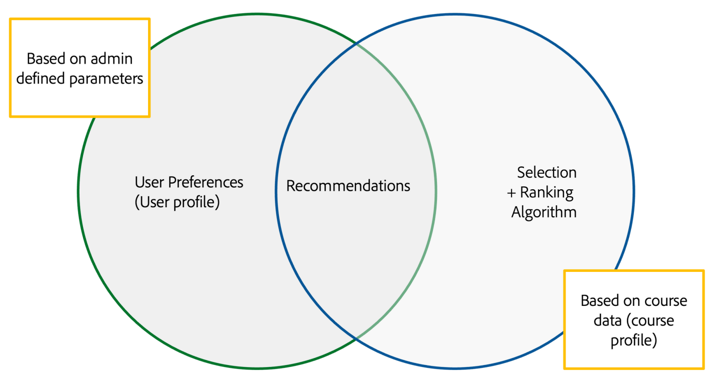

Figure 2: Balance between configurability and AI-driven dynamism

Key Features of the New AI-Based Recommendation Engine (continued)

3. Structured Data and Management

The ability to capture learner preferences data and course metadata in a structured format can have a wide array of applications for the organization.

With the new recommendation engine, Learning Manager is introducing the concept of 'Learning Profiles' and 'Content Profiles' that allow learning leaders to capture learners' preferences across products, roles and levels, and course metadata in a structured format within the platform. This data is easy to manage, search, and monitor for assessing and improving the quality of recommendations via a dashboard.

Recommendations on the Learner Homepage

With the new recommendation engine set up, when a learner logs into the platform, the following recommendations 'strips' are displayed on the learner homepage.

Recommendations Strip	Logic
Super Relevant Strip	Displays personalized content based on all three learner preferences–Products, Roles, Levels and ranked by Learning Manager's AI-based ranking algorithm. The algorithm is built on a model that uses 50 million data points and five years of aggregated learning data across millions of users.
Product/Topic Strips	Displays personalized content based on learner's Products/Topics interests, ranked by Learning Manager's AI-based ranking algorithm.
Discovery Strip	Displays trending content from the account that may be outside of the learner's PRL preferences. All courses in the account are ranked by Learning Manager's AI-based ranking algorithm to drive recommendations to this strip.

Figure 3: Types of Recommendations Strips on the Learner Homepage and their logic.

The following is an illustration of what the three types of recommendations strips translate to on the learner homepage.

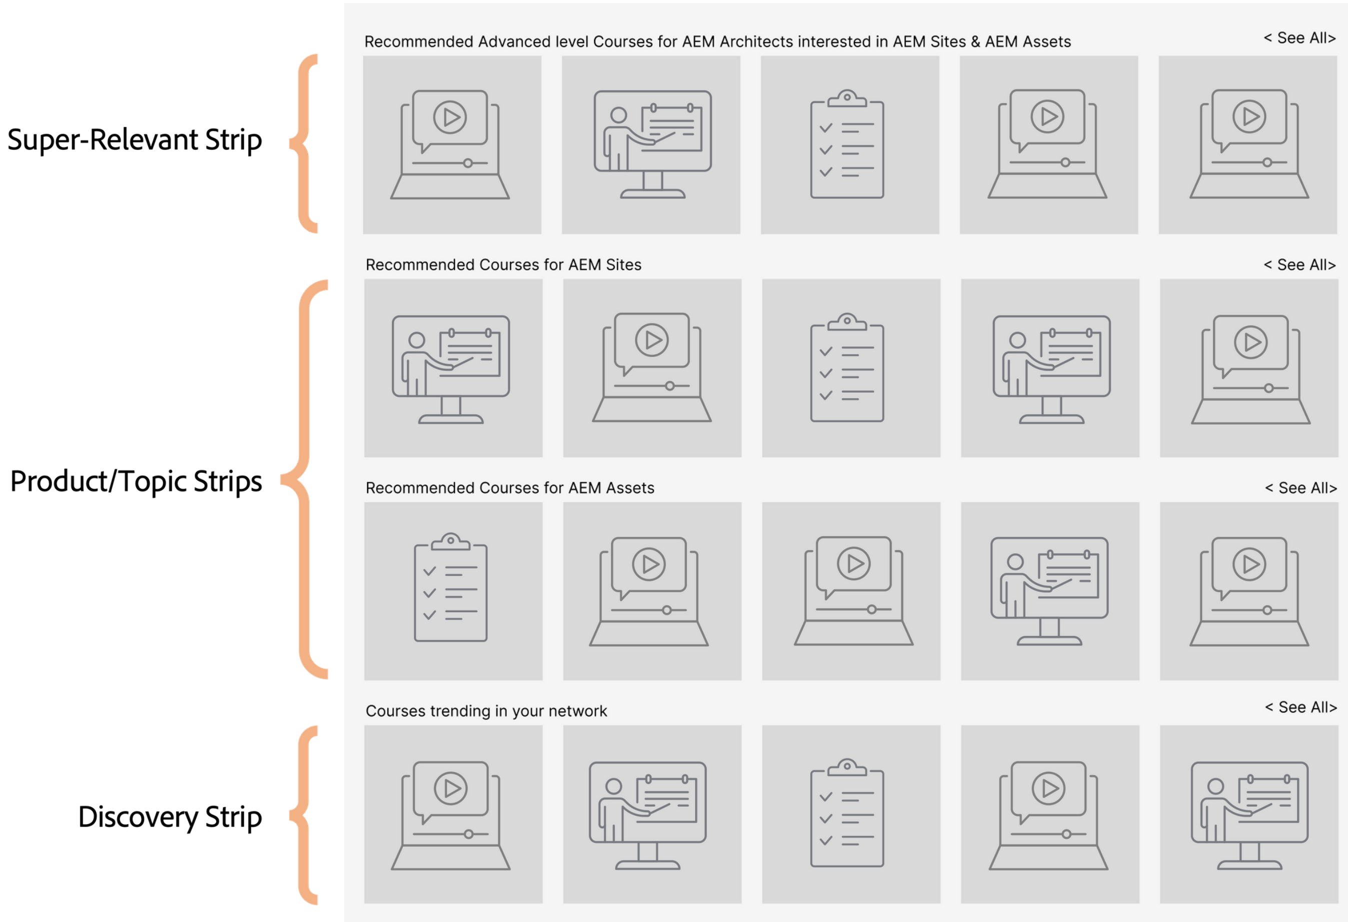

Figure 4: Structure of the Recommendations Strips on the Learner Homepage.

A Glimpse Into How the AI Recommendation Engine Works

Personalization Modes

When learners come to a learning platform looking to acquire new skills, they may expect varying levels of personalization in the recommendations they are provided with, and these can be broadly categorized into three types:

Uber personalization

Personalized Discovery

Trending Content Discovery

Uber Personalization

Very often learners come to the platform looking for specific learning content on their preferred topics or products and within a specific context. Context here would mean that learner, for example, wants to learn about sales but within the context of – 1) sales for an enterprise tech product, and 2) for someone who has significant experience in this function and therefore is looking to learn advanced sales techniques. The learner here does not want to learn about sales in the context of retail sales and is not looking for the basic stuff. The learner in this case is anticipating an "Uber Personalized" experience.

Personalized Discovery

A second type of personalization is where learners are specific about their interest area but prefer an element of exploration in terms of the topics they could learn. They want to be suggested topics within a defined learning area. An example could be of a sales professional who is currently into selling IT services but is looking to broaden her expertise in Retail/FMCG sales. Another example could be of a graphic designer who is great at using the image editing tool/product Adobe Photoshop and is an expert in photo editing and restoration but looking to broaden his Photoshop skills into additional areas such as graphics for social media. In this case, the learners are anticipating a 'Personalized Discovery' of learning content.

Trending Content Discovery

In this third type, learners are keen to know what everyone else is learning about and what courses/topics are trending. For example, currently topics/products such as Generative AI, Design Thinking, Microsoft Outlook Productivity Hacks, Chat GPT could be of interest across various learner profiles. In this case, the learners are open to exploring different topics/products and are looking to learn what's interesting based on their popularity.

A learner coming to the learning platform could be in any of the three modes or more than one mode simultaneously and the recommendation engine should be able to provide learning recommendations accordingly. This is the key philosophy behind the three types of recommendations strips that can be set up on the learner homepage.

Personalization Mode	Recommendation Strip
Uber Personalization	Super Relevant Strip
Personalized Discovery	Product/Topic Strips
Trending Content Discovery	Discovery Strip

Figure 5: Personalization modes and the corresponding Recommendations Strips. Please refer to Page 6, Figure 3 to revisit the details on the types of Recommendations Strips and their logic.

However, this is just the first part of the equation. The second part is Learning Manager's course ranking algorithm. In short, the course ranking algorithm ensures that within each recommendation strip, the most useful content is displayed upfront in terms of the order of display.

Course Ranking Algorithm

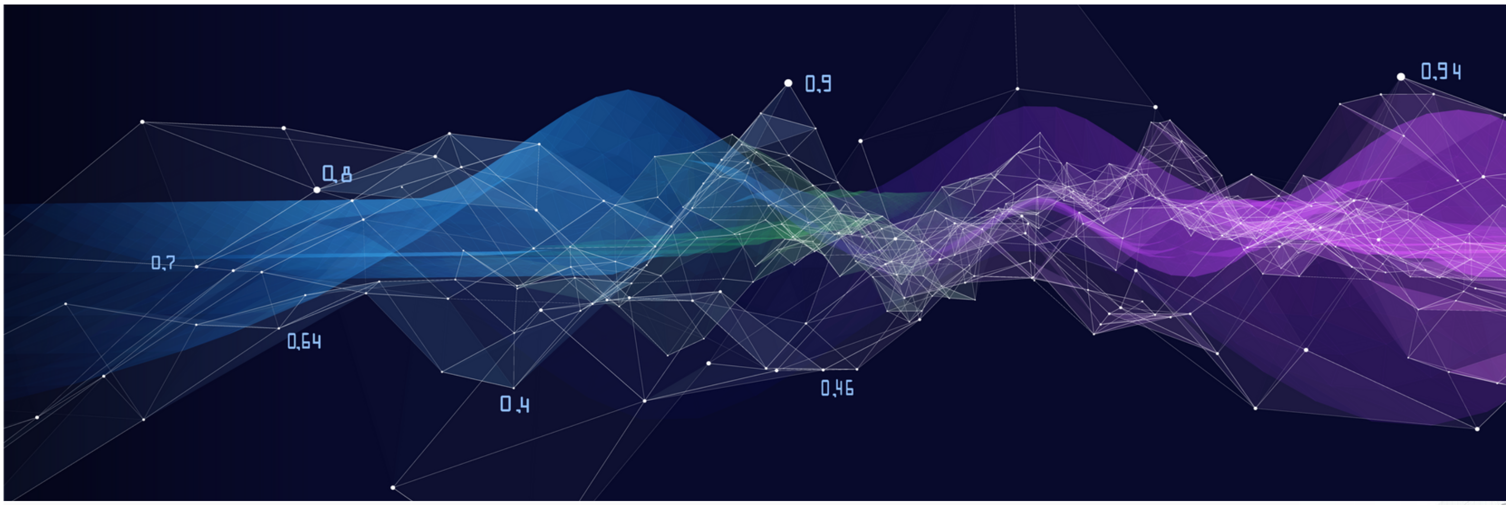

The goal of recommendations in a learning platform is to get learners to learn more. The intention to learn is primarily signalled by enrolling in a course and completing it. By self-enrolling in a course and completing it, a learner effectively signals his/her interest in the course. What then makes sense is that we leverage this signal and transmit it to other learners in the platform as well.

Learning Manager's course ranking algorithm takes enrolment and completion as a proxy for the relevance of that content to a learner, and therefore uses those events as the maximising goal in its AI model. The algorithm generates a ranking score for each course/learning path available on the platform. The higher the likelihood of a course being enrolled and completed–the higher the score generated.

The ranking algorithm has been developed using 50 million data points on learner behaviour and five years of aggregated learning data across millions of users. The two key stages were:

Stage 1	Selection of factors that influence enrolment and completion rates.
Stage 2	Learning the weights for selected features through logistics regression analysis.

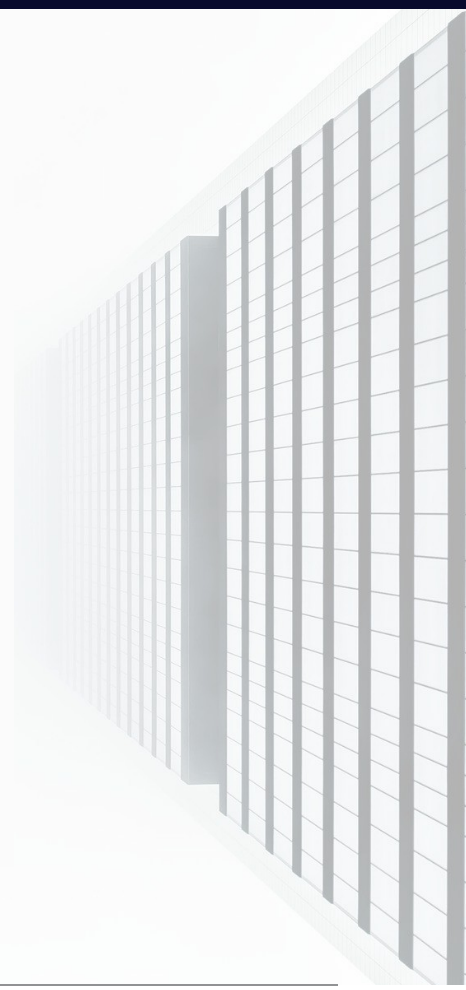

Stage 1

We ran a number of statistical analyses such as Random Forest, XG-Boost among others to identify eight features of any course/path, that are most correlated with the likelihood of a learner's enrolment and completion of that content. These eight features are provided below:

Category	Feature
Enrolments	How many learners have enrolled in the course in the past? Three features based on the enrolment performance of a course are used by the algorithm as described below: 1 Enrolment performance of a course/path in the last 7 days compared to its weekly enrolment performance historically 2 Enrolment performance of a course/path in the last 7 days as compared to all other course/paths in the account 3 Enrolment performance of a course/path in last month as compared to all other course/paths in the account
Recency	How recent is the course? This basically drives more fresh content to the learners. Two features based on course recency are used by the algorithm as described below: 4 Has the course been published in the last 7 days? 5 Has the course been published in the last 30 days?
Ratings	6 How well was the course rated? Better rated courses get recommended more.
Duration	7 What is the duration? Shorter courses require lesser time commitment from learners and therefore such courses are preferred over courses with longer duration.
Completion	8 How many learners were able to complete it in the past? This metric brings in the much-needed element of content quality into recommendations. Courses that have better completion numbers are preferred by the algorithm.

Stage 2

We used logistics regression analysis to generate the weight matrix for the feature vectors. The algorithm calculates the embeddings for each course within our eight-dimensional vector space. This embedding is then transformed into an overall course score by using the weight matrix. This final course score, which is universally calculated for all courses in the account (including any third-party courses that have been imported), is therefore a predictor of the likelihood of the enrolment and completion of the course.

Conclusion

Learning Manager's new AI-based recommendation engine is a powerful tool for organizations who want to construct a highly personalised customer/partner education experience. The recommendation engine dynamically places the best content in front of learners that boosts self-driven training adoption, thereby leading to increased product adoption. The new recommendation engine is lightweight to implement and significantly reduces the Admin workload involved in setting up personalized recommendations.

Ready to learn how to set up the new AI-based recommendation engine for your account?
The following help article provides step-by-step instructions.

https://helpx.adobe.com/learning-manager/recommendations-adobe-learning-manager.html