Knewton adaptive learning

Building the world's most powerful recommendation engine for education

Table of contents

Introduction

Adaptive learning, n.:
A teaching method premised on the idea that the curriculum should adapt to each user.

On a basic level, the definition of adaptive learning seems simple. But dig a bit deeper, and the nuances of the term begin to reveal themselves.

There are many degrees and types of adaptive learning (single-point vs. continuous adaptivity, adaptive testing vs. adaptive learning), but often these distinctions aren't made clear. As the quest for personalized learning gains traction among educators, and more and more products claim "adaptive learning" capabilities, a certain fuzziness has emerged around the term.

Here at Knewton we spend our days and nights thinking about adaptivity, as we iterate on and improve the Adaptive Learning Platform™. This white paper is intended to clarify our interpretation of adaptive learning; provide a glimpse into some of the theories behind our recommendation engine; and explain the effect Knewton adaptive learning can have in a classroom setting. We hope you find it informative.

What Knewton means by "adaptive learning"

Adaptive learning is a term that has been tossed around a good deal recently in edtech circles.

When most people use this phrase, what they're really discussing is either a) single-point adaptivity, which evaluates a student's performance at one point in time in order to determine the level of instruction or material she receives from that point on, or b) adaptive testing, which determines a student's exact proficiency level using a fixed number of questions.

When Knewton refers to adaptive learning, we mean a system that is continuously adaptive — that responds in real-time to each individual's performance and activity on the system and that maximizes the likelihood a student will obtain her learning objectives by providing the right instruction, at the right time, about the right thing. In other words, while adaptive testing answers the question, "How do I get the most accurate picture of a student's state of knowledge with a fixed number of questions?", adaptive learning answers the question, "Given what we understand about a student's current knowledge, what should that student be working on right now?"

To provide continuously adaptive learning, Knewton analyzes learning materials based on a multitude of data points — including concepts, structure, and difficulty level — and uses sophisticated algorithms to recommend the perfect activity for each student, constantly. The system refines recommendations through network effects that harness the power of all the data collected for all students to optimize learning for each individual student.

Whether they offer single-point or continuous adaptivity, most adaptive learning companies today are building adaptive applications based off their own content. Knewton, on the other hand, is not an app but a platform that will make it possible for anyone to build her own adaptive applications, using whatever content she likes. Currently we are working with partners in private beta, but soon we will open up a major toolkit that will make it easy for anyone to use the platform to create courses with continuously adaptive learning built in.
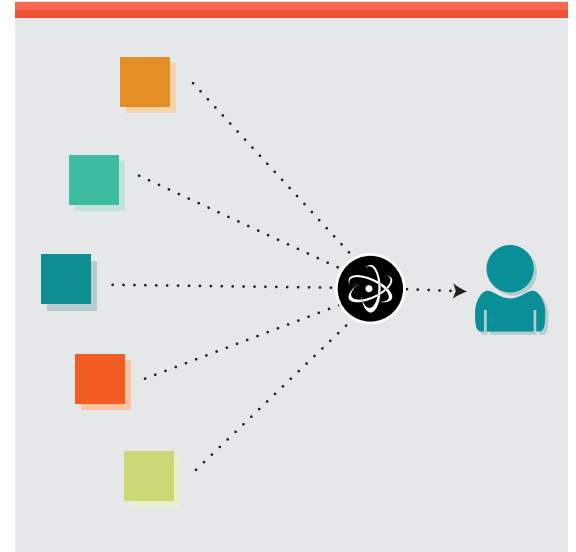

figure a.
A KNEWTON RECOMMENDATION

Theories & approaches behind Knewton recommendations

No two students are identical — they learn and forget at different rates, come from different educational backgrounds, and have different intellectual capabilities, attention spans, and modes of learning. As a result, designing a real-time recommendation engine that is sensitive to the characteristics of each student is an immensely complex task.

Here at Knewton we address this challenge head-on, using educational path planning technologies and advanced models of student ability. These technologies and models ensure that every student progresses through the course material in a way that maximizes his learning.

Here's a quick look at some of the theories and approaches behind the Knewton recommendation engine:

Item Response Theory (IRT)

Imagine that you're teaching a fourth grade math class. You've just administered a test with 10 questions. Of those 10 questions, two questions are very simple, two are incredibly hard, and the rest are of medium difficulty. Now imagine that two of your students take this test. Both answer nine of the 10 questions correctly. The first student answers an easy question incorrectly, while the second answers a hard question incorrectly. Which student has demonstrated greater mastery of the material?

Under a traditional grading approach, you would assign both students a score of 90 out of 100, grant both of them an A, and move on to the next test. But this approach illustrates a key problem with measuring student ability via testing instruments: test questions do not have uniform characteristics. So how can we measure student ability while accounting for differences in questions?

IRT models student ability using question-level performance instead of aggregate-test-level performance. Instead of assuming all questions contribute equivalently to our understanding of a student's abilities, IRT provides a more nuanced view on the information each question provides about a student. It is founded on the premise that the probability of a correct response to a test question is a mathematical function of parameters such as a person's latent traits or abilities and item characteristics (such as difficulty, "guessability," and specificity to topic).

Figure B shows two item response function curves generated by an IRT model. The curves illustrate how an IRT model relates a student's ability with the probability of answering a question correctly, given that question's difficulty, discrimination levels, and "guessability."
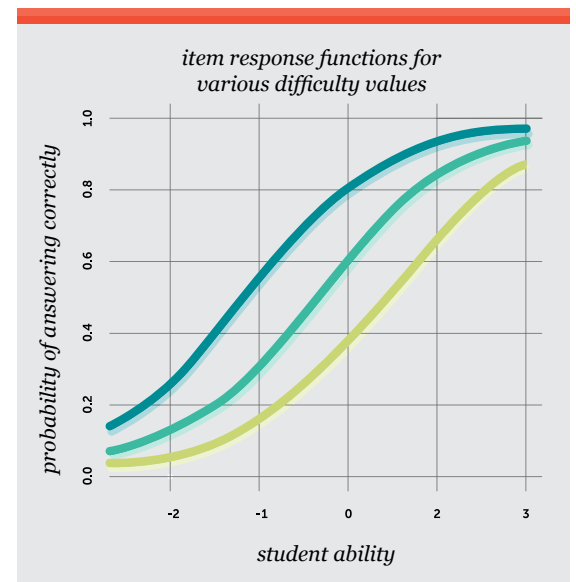

item response functions for various difficulty values

probability of answering correctly

student ability

figure b.
ITEM RESPONSE THEORY

While IRT models are atemporal and reliant upon a single measure of ability (and thus reflect only one facet of the science behind Knewton recommendations), they help us better understand how a student's performance on tests relates to his ability.

Probabilistic Graphical Models (PGMs)

This framework, which encompasses statistical methods such as Bayesian networks and Markov random fields, allows data scientists to code and manipulate probability distributions over multi-dimensional spaces in which there are hundreds or even thousands of variables at play. In other words, PGMs allow Knewton analysts to build complex models one effect at a time, relating the many learning activities they observe to estimations that are useful for recommendation.

One of the ways in which Knewton applies PGMs is by using a student's known proficiencies to determine which other topics he may be ready to master. For instance, such a model might help the platform discover to what degree a mastery of fractions helps students master decimals and to what degree a mastery of decimals helps students master exponentiation. Knewton data scientists can thus determine the relationship between mastery of fractions and mastery of exponentiation. Ultimately, the discovery of these types of relationships allows the Knewton Adaptive Learning Platform™ to continually refine its recommendations.
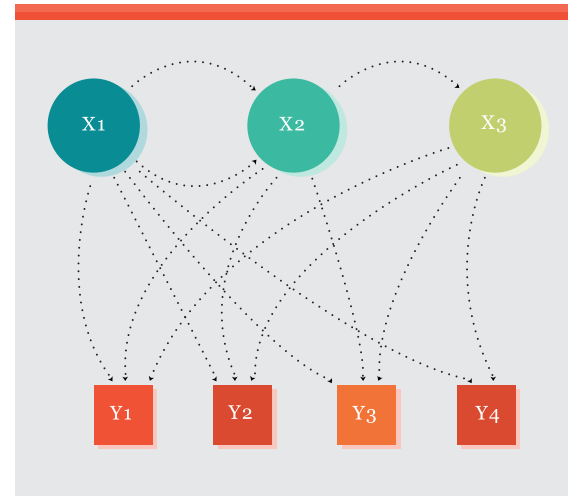
Hierarchical agglomerative clustering

In data mining, hierarchical clustering is a method of analysis that aims to construct a hierarchy or structure of clusters. At Knewton, the technique is used to detect latent structures within large groups and build algorithms that determine how students should be grouped and what features they should be grouped by. An implementation of this technique is incorporated in Knewton Math Readiness, a web-based developmental math course, which provides a dashboard that allows teachers to group students according to their level of concept mastery of the material they are working on.
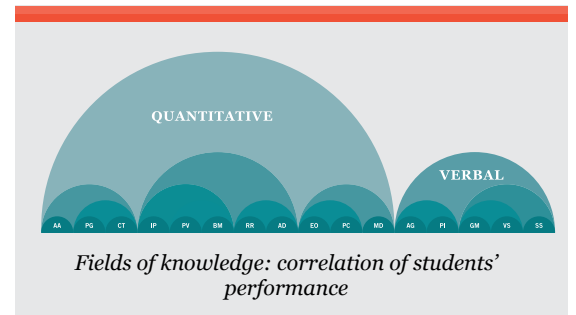

figure c.
PROBABILISTIC GRAPHICAL MODEL

Fields of knowledge: correlation of students' performance

figure d.
HIERARCHICAL AGGLOMERATIVE CLUSTERING

Why the Knewton Adaptive Learning Platform™ is so effective

The science behind recommendation is enhanced by the disciplinary range and tremendous scope of Knewton adaptivity. A cross-disciplinary knowledge graph, continuous adaptivity, lifelong student learning profiles, and vast network effects combine to produce powerfully personalized learning for every student who takes a Knewton-powered course.

Knowledge graph

Knewton-powered courses are linked by the Knewton knowledge graph, a cross-disciplinary graph of academic concepts. The knowledge graph takes into account these concepts, defined by sets of content and the relationships between those concepts. Knewton recommendations steer students on personalized and even cross-disciplinary paths on the knowledge graph towards ultimate learning objectives based on both what they know and how they learn. The more content that teaches or assesses each concept that is added to the system, the more precise the adaptive course experience becomes.

When visualized, the knowledge graph can provide a sense of a student's potential flow through the course material.

Within the knowledge graph, concepts have prerequisite relationships that contribute to defining a student's path through the course. Special relationships that define content as either "instructional" or "assessment" determine what kind of content to deliver to students at any given point.

Continuous, as opposed to single-point adaptivity

A single-point adaptive learning system evaluates a student's performance at one point in time, and from there determines the type of instruction she receives. An example of single-point adaptivity would be a course that includes a diagnostic exam, the results of which determine subsequent course content, with little or no further data mining and personalization.

Knewton's continuously adaptive learning system, on the other hand, constantly mines student performance data, responding in real time to a student's activity on the system. Upon completion of a given activity, the system directs the student to the next activity. For example, when a student struggles with a particular set of questions, Knewton will know where that particular student's weaknesses lie in relation to the concepts assessed by those questions and can deliver content to increase the student's proficiency on those concepts. In this way, Knewton's continuously adaptive system provides each student with a personalized syllabus at every moment.

Lauren

- Equations and expressions
- Statistics and probability
- Algebra
- Functions
- Ratios and proportions
- Geometry
- The number system
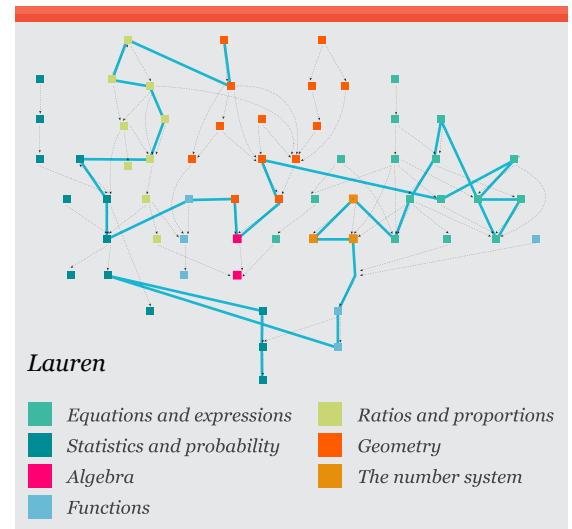
figure e.
LAUREN'S PERSONALIZED LEARNING PATHWAY

William

- Equations and expressions
- Statistics and probability
- Algebra
- Functions
- Ratios and proportions
- Geometry
- The number system
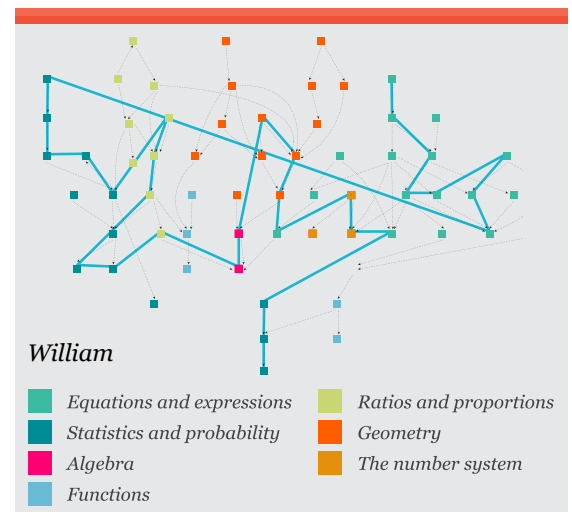
figure f.
WILLIAM'S PERSONALIZED LEARNING PATHWAY

The following are specific examples of approaches that allow Knewton to offer truly continuously adaptive learning:

Spaced reinforcement

In contrast with massed reinforcement, the standard method of drilling which requires students to apply new concepts or skills in a short period of time until they demonstrate mastery, spaced reinforcement (also referred to as distributed reinforcement) is a learning method in which new concepts or skills are absorbed while previously-taught concepts and skills are reinforced. Because new material is introduced incrementally and woven into familiar material, spaced reinforcement typically occurs over an extended period of time. Spaced reinforcement allows Knewton recommendations to help students build their skills in a cumulative way and retain understanding once it is gained.

Retention & learning curves

The Knewton recommendation engine needs to be able to take the degradation or diminishment of skill (or forgetting) into account. That is, it needs to be able to detect such occurrences and provide actionable recommendations as a result.

Inspired by Hermann Ebbinghaus' work on memory retention and learning curves, Knewton data scientists have used exponential growth and decay curves to model changes in student ability while learning and forgetting. These curves are governed by the following premise: each time students are exposed to content associated with a given topic, they receive a "bump" in their virtual ability level for a topic; likewise, if they are not exposed to some other topic, they likely "forget" that topic over time. The forgetting curve itself that governs rate of retention is roughly described by the following formula:

$$R = e^{\frac{-t}{s}}$$

where R is memory retention, S is the relative strength of memory, and t is time.

By integrating this curve into engine validation efforts, Knewton data scientists can capture the way a student's knowledge waxes and wanes, depending on how and when she is exposed to content. Ultimately, the process allows Knewton data scientists to test the algorithms that govern a student's flow through the course.

Student learning profile

With Knewton, students can maintain a continuously updated learning profile that contains information on what the student knows and how she learns best. The profile is progressive, which means it keeps getting smarter the longer the student remains on the platform.

For instance, if a student who has already taken a Knewton-powered course enrolls in another, the course starts "warm" with that student's data
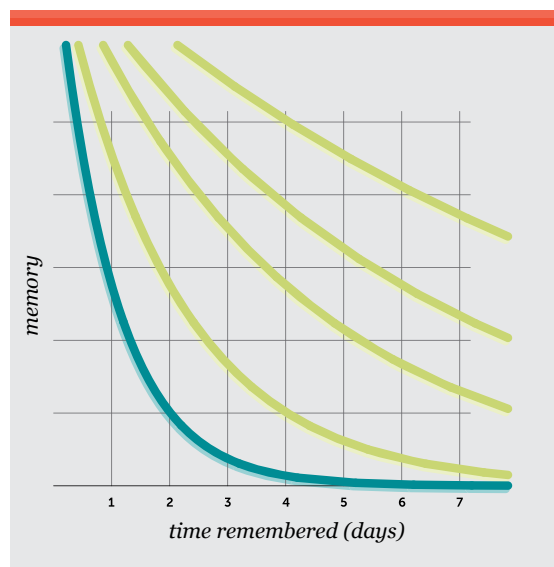

figure g.
THE FORGETTING CURVE

(as opposed to starting "cold" with no data). The course takes into account the student's recently mastered concepts and skills and unique trajectory through the material, and uses this knowledge to maximize student learning continuously from that point forward. Once enough data is collected, the platform will uncover patterns in the student's learning, such as blind spots, modality and medium preferences, and granular strengths and weaknesses. The more often a student uses Knewton-powered courses, the more effective the platform becomes at serving up targeted learning material.

In this way, the Knewton Adaptive Learning Platform™ works to minimize unproductive feelings of frustration and confusion and build student skills in a natural way over time. More fundamentally, it provides penetrating insight into students' own understanding of the material — what they truly grasp and don't grasp, their misunderstandings and misconceptions. Students develop a deeper, more nuanced understanding of their learning style and strengths and weaknesses, helping them to maximize their academic potential.

The implications of all this are straightforward: student engagement can be strengthened if academic work is imbued with a sense of continuity. Nothing is more dissatisfying to students than feeling like the challenges they face are essentially arbitrary and culminate in nothing. The Knewton learning profile answers the student need for continuity and meaning by affording students a sense of long-term investment in the learning process.

Network effects

The more students who use the Knewton platform, the more powerful it becomes — the more refined the relationships between content and concepts and the more precise the recommendations. For each student learning each individual concept, the system finds a population of extremely similar students who have already learned that concept. It then asks, "Who among those similar students learned that concept the best, and what did they do that worked so well?" The system then allows this proven effective learning path to inform the student's learning path going forward. In this sense, a student experiencing a particular challenge (for example, spatial skills as they relate to math word problems) need not be limited by the fact that no one else in her class is experiencing the same difficulty. The Knewton Adaptive Learning Platform™ is able to take the combined data of millions of other students to help each student learn every single concept she ever encounters.

Network effects are a natural consequence of the knowledge graph and the models Knewton uses to determine its recommendations. In isolation, each student's response to each question is only a tiny scrap of information, but when propagated through the entire system and understood in context, the value of that information is amplified tremendously. Every student action and response around each content item increases the system's understanding not only of the student and the content item, but also, by extension, of all the content in the system and all the students in the network.

> *The Knewton Adaptive Learning Platform™ is able to take the combined data of millions of other students to help each student learn every single concept she ever encounters.*

Big data & adaptive infrastructure

The term "big data" is used to describe the tremendous volume, velocity, and variety of data generated by various technology platforms, many of which involve the continuous or ubiquitous collection of data. Big data refers specifically to data sets that are so large and complex that they are challenging to work with using traditional database management tools; specific challenges include the storage, search, analysis, and visualization of data.

Big data & education

The advent of "big data" in areas such as internet search and social media has disrupted existing industries, created new industries, and led to the extraordinary success of companies such as Google and Facebook. Big data unleashes a range of productive possibilities in the education domain in particular, since data that reflects cognition is structurally unique from the data generated by user activity around web pages, social profiles, and online purchasing habits.

One feature that distinguishes the data produced by students (from that of consumers shopping online or engaging in social media, for example) is the fact that academic study requires a prolonged period of engagement; students thus remain on the platform for an extended length of time. Furthermore, there is a focus, intention, and intensity to students' activity: they are engaging in high stakes situations — taking a course for credit, trying to improve their future, expanding their range of skills. The sustained intensity of these efforts generates vast quantities of meaningful data that can be harnessed continuously to power personalized learning for each individual.

Another feature that distinguishes the data produced by students is the very high degree of correlation between educational data and the aggregated effect of all those correlations. If, for example, a student has demonstrated mastery of fractions, algorithms can reveal how likely it is that he will demonstrate mastery of exponentiation as well — and how best to introduce that concept to him. If a student has demonstrated mastery of various grammatical concepts (say, subjects, verbs, and clauses), educational data can optimize his learning path, so that different sentence patterns will "click" for the student as quickly as possible.

The hierarchical nature of educational concepts means that they can be organized in a graph-like structure, which means that the student flow from concept-to-concept can be optimized over time, as we learn more and more about the relationships between them through data. Every student action and response around each content item ripples out and

> *Big data unleashes a range of productive possibilities in the education domain.*

affects the system's understanding of all the content in the system and all the students in the network.

Adaptive infrastructure

Knewton has established an infrastructure that allows the platform to process tremendous amounts of student data. For instance, inference on probabilistic graphical models is one example of a class of algorithms called "graph algorithms." These algorithms are special in that they can be broken down into units of computation that depend only on other specific units and can thus be parallelized very efficiently if the work is split between computers so that limited coordination is required.

Given the absence of robust, public frameworks for accomplishing these computations at a large scale, Knewton has designed its own framework called AltNode which works by dividing work between machines and then sending continuous updates between the minimal necessary number of machines. All significant updates are stored in a distributed database and broadcast between services. If one machine fails, another one automatically takes its place, recovering recent values from the database and resuming work. One unique feature of AltNode is that it allows models to recover from any state and respond to new data as it arrives.

figure h.
ALTNODE FRAMEWORK

Knewton adaptive learning in the classroom

Adaptive learning supports mastery-based learning, a school of teaching founded on the idea that student progression through a course should be dependent on proficiency as opposed to the amount of time spent on academic work. Though it may not always be referred to by name, mastery-based learning describes any situation in which one is given a set of problems, labs, or activities, and in which progression through that material is dependent on successful completion of various tasks rather than seat time.

Online adaptive learning makes it possible to implement mastery-based learning in a scalable way. Knewton Math Readiness, for example, creates a guided, self-paced environment in which live instruction is optimized around targeted group sessions. The course is designed to present students with personal learning paths as it continually assesses their mathematical proficiency and adapts accordingly. Lessons consist of videos, online textbook selections, and lesson quizzes, and students progress through the course by earning badges. Early efficacy reports reflect the success of the program: after two semesters of use with over 2,000 developmental math students at Arizona State University, withdrawal rates dropped by 56%, pass rates went from 64% to 75%, and 45% of the class finished four weeks early.

This implementation model is often referred to as *blended learning*, a term which describes any arrangement in which a student learns in part at a brick-and-mortar facility and in part through online delivery with student control over time, place, path, or pace.

Irene Bloom, a Senior Lecturer at ASU, was originally a skeptic of online learning. But she says that the classroom dynamic has changed for the better since introducing Knewton into her remedial math classes: "I love looking around the classroom and seeing them working in groups, talking to each other and explaining things to each other... Most of the time, different groups are working on different things, depending on where they are in the course. This is very new for me. Before this, I worked on the assumption that all students were at the same place. Now, because they progress at different rates, I meet them where they are."

ASU Remedial Math	Fall '09-Spring '11	Fall '11-Spring '12 (with Knewton)
Pass Rates	**64%**	75%
Withdrawal Rates	**16%**	7%
Students Finishing Early	**n/a**	45%

figure i.
KNEWTON MATH READINESS
RESULTS AT ASU

"Before this, I worked on the assumption that all students were at the same place. Now, because they progress at different rates, I meet them where they are."

How Knewton adaptive learning engages students

Knewton adaptive learning can improve student engagement by increasing self-confidence, decreasing discomfort and frustration, and encouraging productive learning habits.

Instant feedback

Students are less likely to lose focus if feedback is immediate and they can quickly self-correct. A continuously adaptive learning system is able to deliver personalized feedback to both multiple choice and free response questions quickly — that is, instantaneously or near-instantaneously. The result is pacing that is conducive to risk-taking, experimentation, iterative development, and rapid learning.

Community & collaboration

Isolation can exacerbate the challenges students experience in school. An adaptive system can improve student engagement by weaving a social component into coursework. Knewton Math Readiness, for instance, provides a dashboard that allows teachers to group students who are working on the same material. Using the reporting features, teachers can also arrange peer review opportunities and form groups of students whose abilities complement each other.

figure j.
STUDY GROUPS AND PEER REVIEW

Gamification

With countless opportunities for students to demonstrate skill and reflect on action and feedback, adaptive courses naturally have much in common with games. What's more, adaptive courses keep students in a game-like state of "flow" by escalating the difficulty of the work incrementally and unveiling levels one at a time to increase suspense. These and other game elements can be heightened to transform adaptive courses into truly gamified learning experiences.
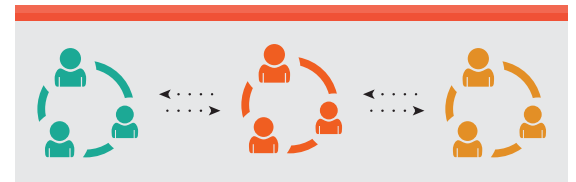

figure k.
GAMIFICATION VISUAL ELEMENTS

How Knewton adaptive learning empowers teachers

Knewton adaptive learning gives teachers insight into the learning process, specifically in terms of efficacy, engagement, and knowledge retention. The platform also provides an unprecedented flexibility of scope; teachers can grasp patterns in student activity and performance across the whole class or drill down into individual student profiles to determine exactly why a student is struggling.

Addressing the diverse needs of students

One of the biggest challenges facing teachers and school administrators today is the growing diversity of the students within their population. A greater diversity of students means a greater diversity of needs to consider. Some struggle because English is not their first language; others have difficulty with focus or organization. Others may be particularly weak in some area but possess unusual strengths in another.

Knewton adaptive learning allows teachers to address the needs of diverse students while gaining insight into the learning process. The platform may discover, for instance, that a student who is weak with math word problems is struggling because he has difficulty with reading comprehension; the system can then direct the teacher to specific material on syntax and vocabulary and suggest ways that the teacher might coach the student individually in these respects. Later, the instructor may be informed that another student who understands mathematical concepts but has trouble with carelessness in arithmetic should receive feedback about how to develop stronger estimation abilities or check work once completed. The instructor can then coach that student with a precise understanding of his particular weaknesses.

For example, the Knewton Math Readiness course dashboard includes an "on-track/off-track" concept to measure student progress through the course. This concept functions as a binary indicator that helps teachers efficiently grasp information about the class as a whole. Using this tool and others, teachers can see reporting data from two perspectives.

The whole class

The Knewton Math Readiness instructor dashboard includes a histogram which provides a big picture assessment of the whole class' "track" status.

Using the dashboard, teachers can also see how students are performing in individual subject areas; which segments of material are the most and least challenging for students; and what kinds of patterns in both performance and activity emerge across the class. After multiple years of teaching the same course, teachers will be able to compare data from year to year. Knewton analytics will help them home in on useful information while leaving them free to interpret the results.
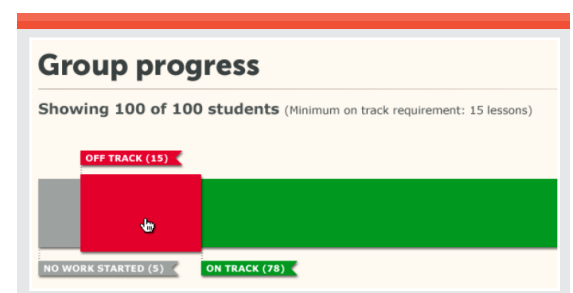

Group progress

Showing 100 of 100 students (Minimum on track requirement: 15 lessons)

OFF TRACK (15)

NO WORK STARTED (5)   ON TRACK (78)

figure l.
GROUP PROGRESS DASHBOARD

Individual students

While the reporting dashboard in Knewton Math Readiness is streamlined so that teachers can focus on the big picture, the dashboard also allows teachers to click into the interface and drill down to each student's work in the system. Teachers can see how students have performed on specific quizzes and exams. If a student isn't grasping the material, teachers can determine (using analytics that guide them to specific data points) where precisely any misunderstanding is occurring.

This capacity allows teachers to both address the diverse needs of students and better understand their content, so that they can refine it from year to year.

Improving & understanding content for long-term curriculum development

An instructor dashboard that measures the efficacy of content can help teachers determine the strongest and weakest aspects of their teaching materials. As described previously, Knewton adaptive learning can help teachers understand precisely what the content they are working with teaches and assesses. This ensures that content can be analyzed for fine-tuned improvements from year to year, and that students are never stuck with outdated or ineffective materials.
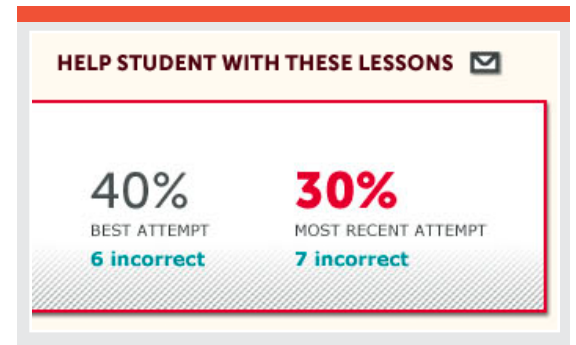

HELP STUDENT WITH THESE LESSONS

40%
BEST ATTEMPT
6 incorrect

30%
MOST RECENT ATTEMPT
7 incorrect

figure m.
INDIVIDUAL PROGRESS DASHBOARD

Knewton is building products and services that will make it easy for all publishers, teachers, and developers to use Knewton adaptive learning.

Whether you have an existing digital learning product or are looking to build a new online course, Knewton can provide continuously personalized recommendations for your students via the Adaptive Learning Platform™.

We are currently working with partners in private beta, but will be expanding our partner program in the coming months.

Please let us know about your adaptive learning needs by filling out our partner survey:

www.knewton.com/signup-platform