


Using Research in Ed-Tech

How ed-tech companies use
research to make better products

Authored by

Christa M. Simone
Kelsey Gross

January 30, 2017



Digital Promise

Accelerating Innovation in Education

www.digitalpromise.org

Table of Contents

- Introduction 3
- Learning Science 4
- User Research. 5
- Evaluation Research 7
- Conclusion and Recommendations. 9

Introduction

As educational technology (ed-tech) becomes more ubiquitous in K-12 education, schools must be able to differentiate between high- and low-quality products. However, it can be difficult for school administrators to make informed, evidence-based purchasing decisions, as many products do not have strong research support.

In a 2015 [study](#) by the Joan Ganz Cooney Center, for instance, only 24 percent of top literacy apps mentioned research in their app store descriptions. Yet, the ed-tech

field is growing and changing, and more companies are using research in innovative ways to inform and validate their products.

In this report, we share a snapshot of the current state of research use in ed-tech, and highlight some promising practices for using and conducting research in product development. This information can help ed-tech companies improve their research and development (R&D) processes and help schools make more informed decisions about which products to purchase.

Methods

In summer 2016, Digital Promise ran a campaign in which we asked ed-tech companies to share how they use research to inform product development and evaluate product efficacy. Companies could submit responses in three categories. **These categories represent research use in early, mid, and ending stages of product development:**

Learning Science

Using existing scientific research on learning to inform product design

User Research

Gathering information about how users interact with the product or prototypes of the product to make improvements

Evaluation Research

Conducting a formal research study to determine whether the product is effective in promoting learning or solving a problem in education

We received submissions from [53 companies](#). Researchers at Teachers College, Columbia University and Digital Promise reviewed companies' responses to the [submission form questions](#) and scored each one based on a rubric that assessed three main areas: 1) how rigorous and appropriate the company's research methods were; 2) whether the company worked with a researcher or had a researcher on staff; and 3) whether the company shares its research publicly.

Then, within each category, we assessed additional criteria. In the Learning Science category, we reviewed the rigor of the cited studies, and how well the cited research connected to the product. In the User Research category, we reviewed whether the company

chose relevant and rigorous methods to gather feedback, and how well they got to know different stakeholder groups. In the Evaluation Research category, we examined the research design and methods, and how well the research context matched the product's real-world use.

Multiple reviewers scored each application, and scores were averaged. [One exemplar and two honorable mentions](#) were selected in each category based on the final scores.

We analyzed submission data and performed descriptive statistics to gain insight into trends across the set of submissions. In the sections below, we share the results of these analyses as well as recommendations for ed-tech companies as they use research to design, develop, and evaluate their products.

Learning Science

A [large body of academic research](#) exists about how people learn, and this research continues to grow as more is uncovered from the fields of education, neuroscience, psychology, and many others. Ed-tech tools are a promising way to apply this knowledge in real-world education practice, but it can be difficult for developers — many of whom are not research experts — to navigate the academic literature and apply relevant research in product design.

The 33 companies that submitted in the 2016 Learning Science category consulted a variety of research resources, and cited many of the ideas outlined in [our 2015 report](#) on finding and applying Learning Science research.

Specifically, one of the most popular approaches among this year's submissions was to work with an experienced researcher. Eighty-five percent of the companies either had a researcher on staff or worked with an external researcher to find high-quality research relevant to their goals, and to apply research in their product development. Because many different fields of study can be relevant to product development, including cognitive science, social psychology, pedagogy, curriculum, and assessment, several companies also formed advisory boards with experts from across these different fields.

Companies that scored well in this category used research findings from several different disciplines to build their products. For example, BrainQuake, the top selection in this category, incorporated multiple streams of research when designing and developing its mathematics product. They consulted theories of math pedagogy when crafting content for their games, and used research on mindset to create features that tracked students' persistence.

Further, they incorporated findings from media theory and gamified learning research to develop a product that was educational and engaging.

In addition to identifying relevant research from multiple fields, top companies in this category successfully translated what they learned from the research into product features that function in real-world settings. BrainQuake used research demonstrating that symbolic representation of numbers and other mathematical concepts is essential for more advanced mathematics, but can be very difficult for early math learners (Kirkpatrick et al 2001, Nunes et al 1993). As a result, they focused on providing learners with symbolic representations of math concepts, and opportunities to work through multi-step problems to improve symbolic comprehension. Another example is TenMarks Math, which referenced a finding that the most effective learners are able to explain content back to themselves periodically, "making inferences and justifications not presented in the material" (VanLehn, Jones, & Chi 1992, Chi 2000). Because of this, TenMarks incorporates self-explanation exercises as a key feature of its math curriculum, and gives learners the opportunity to internalize and reinforce math concepts.

User Research

Once a working prototype is ready, companies use research to learn more about users' experience with the product. This process, called user research, is generally very iterative. Users interact with an early prototype of the technology and give feedback. Then, the product developers use this feedback to make improvements, and the process begins again with an updated version of the product.

Methods of User Research

The 32 companies that submitted in this category utilized several different methods to conduct their user research (Table 1).

Table 1:
Methods used to conduct user research.

Method	Description	Numbers of Companies (% of total submissions)
Interview	One-on-one sessions with a student, teacher, parent, administrator, etc. Can use a set of specific questions or more open-ended themes.	23 (72%)
Survey	Online or paper-based questionnaire about a user's experience with the product.	20 (63%)
Observation	A researcher observes a user as he/she uses the product and notes where misunderstandings occur or improvements could be made.	18 (56%)
Focus Group	Similar to an interview, but with multiple users all together.	16 (50%)
Data Analytics	Analyzing data collected automatically during product usage, which provides insight into users' learning, performance, and progression.	16 (50%)
Pilot Study	A small-scale study to evaluate the product in a real-world setting, or demonstrate "proof-of-concept."	12 (38%)
A/B Testing	Comparing two different versions of the product to see which one users prefer	7 (22%)

Compared to other types of products, technology products offer the benefit of automatically tracking and analyzing real-time user data, called data analytics. The term data analytics here refers to data collected through the product about a user's interaction, such as clicks, activities, or scores. Many companies in this campaign used this data to understand the user experience. For example,

Oxford Learning used data analytics to analyze each student's progress in their Easy Read System. In particular, they studied the data for users who did not progress through the entire system successfully. After analyzing the data and talking with students, teachers, and parents, the company developed a map of eight different causes of reading difficulty. They used this model to make their product

more personalized for each student’s reading ability, leading to an increase in the program’s completion rate.

In addition to the more common methods in Table 1, companies described some other innovative methods of gathering user feedback. For example, two companies formed an “Educator Advisory Board” or “Customer Advisory Board” of teachers that met regularly with the product developers to give feedback on future product development plans. This innovative approach allowed the companies to proactively receive educator input before developing new product features, as opposed to asking for their opinions after creating a finished product.

Another innovative user research method was eye tracking. Eye tracking is a technology that uses a specialized camera to detect the direction of a user’s gaze (e.g., where one is looking). One company used this technology to understand how and why some children were having difficulties with their reading app. A second company used eye tracking in user testing sessions to collect data about where students were focusing their attention, allowing them to better understand the features that evoked the highest student engagement.

User Audiences

Nearly all (94 percent) of the companies engaged with more than one stakeholder group when conducting user research. The most common stakeholder groups were students, teachers, administrators, and parents (Table 2). Other audiences included academics/researchers, curriculum developers, librarians, corporate partners/stakeholders, counselors, and admissions staff.

Table 2:
Most common user groups consulted during user testing.

User groups	Number of companies receiving feedback from each group (% of total submissions)
Students	28 (88%)
Teachers	26 (81%)
Administrators	11 (34%)
Parents	8 (25%)

High-scoring companies in this category used multiple methods of gathering data, and spoke with multiple stakeholder audiences. This approach likely increased the quality of user research and the feedback companies received, providing a well-rounded picture of the product.

Evaluation Research

Evaluation research helps answer the question, “Does the product work?” These studies are conducted in a real-world setting, and must be carefully designed to measure whether the product delivers on its promise. For example, if the product is meant to enhance math learning, then an evaluation study should measure whether students’ mathematical ability improved after using the product.

Thirty companies submitted in the Evaluation category. The companies ranged in age from less than one year to 35 years old, though two-thirds were less than five years old. The mean score was 8.3 out of 12 possible points.

Among the companies ten years old and younger, there was a wide range of scores, which indicates there is no direct relationship between company age and score (Figure 1). The large proportion of young companies with high scores may undermine a belief we often hear that younger ed-tech startups do not have the time and resources to conduct evaluation research.



Figure 1: Relationship between companies’ age and score in the evaluation category.

Over half (57 percent) of companies in this category had a researcher on staff to conduct evaluation research. Even more (77 percent) partnered with a university, external researcher, or research organization to conduct evaluations (Table 3). Companies reported that partners provided advice about research design, and some leveraged their networks to gain access to schools that could test the product.

Table 3: Research Partnerships

Type of research relationship	Number of companies reporting (% of total submissions)
Researcher on staff	17 (57%)
University partnership	19 (63%)
Contract Research Organization (CRO) or Non-profit partnership	8 (27%)
Consultant	4 (13%)

Another benefit of partnering with an external researcher is that they are independent of the company, which may help to reduce bias in collecting and analyzing the data, and ultimately provides more credibility to the research. For example, BrainQuake partnered with graduate students of a professor at the Stanford Graduate School of Education to test their math tool in a local elementary school. The students designed and executed the evaluation study, conducted the analysis, and ensured an unbiased portrayal of the study’s results.

In addition, 30 percent of the companies in this category recruited research experts to join their scientific advisory boards. As in the

Learning Science category, research experts can help guide product development, as well as conduct product evaluation research. Scientific advisory boards provide companies with a way to get occasional, high-level feedback from a variety of experts.

Among the 30 companies in this category, just over half (57 percent) publically share the results of their evaluation studies on their websites. Of these, only four companies published their research in a peer-reviewed journal, which is reviewed by experts in the field and provides the most credibility. This low number may reflect the frequently demanding and long journal publication process, which often extends beyond the ed-tech product development cycle. Many companies instead offer their research results in the form of a white paper, which provides a comprehensive description of the study and findings.

Two important indicators of a rigorous evaluation study are whether it includes a *comparison group* and uses *randomization*. In a comparison group study, one group of students uses the ed-tech product while the other continues with learning as usual, or uses a different, comparable product or program. Both groups should be as similar as

possible, so any changes can be attributed to the ed-tech tool and not some other factor, like the introduction of iPads into the classroom. Randomization is the process of randomly assigning students to either be in the comparison group or the treatment group (using the tool); this additional step again reduces the chance that some group-level factor other than the ed-tech tool is causing the difference in the outcome being measured.

Among the companies who submitted in this category, 67 percent used a comparison group, and 23 percent used randomization when assigning groups. Randomization can be difficult in school settings because it can be burdensome for teachers to offer different activities for groups of students within a class. WootMath, the top entry in this category, took an innovative approach to this problem as they evaluated the adaptability feature of their product (increasing difficulty as the student progresses) and its effect on student motivation and engagement. They had the comparison group use a non-adaptive version of their math tool, so all the students believed they were receiving the same program. Because of this, students in both groups were similarly motivated and engaged with the product.

Conclusion and Recommendations

Research can guide ed-tech developers to build more engaging and effective products, and provide school leaders with information about which products are grounded in evidence. Though the field is still relatively new, the ed-tech companies that participated in this campaign are using research in innovative ways at every stage of product development, and offer a wealth of promising practices. These include:

- When designing your product, consult multiple learning science research resources across different disciplines that are relevant to all aspects of your product.
- Form a scientific advisory board to get expert advice on the development of your product, and to make sure it is grounded in sound scientific theory.
- Conduct user research using multiple methods and with multiple different stakeholder groups, and iterate quickly.
- Find a partner research organization to help you design and implement your evaluation study. This could be a university-based partner, a non-profit or contract research organization, or consultant.

- If possible, use a comparison group and randomization when conducting an evaluation study.
- Consider partnering with a researcher to conduct a rigorous study that could be submitted to a peer-reviewed journal.
- Share your research results publicly.

Ultimately, every stage of research contributes to a feedback loop that drives continuous innovation and learning. Existing research informs early product development, user research hones the product so it is engaging and easy to use, and evaluation research generates new data about learning, which in turn drives further product improvements.

As more ed-tech companies follow these recommendations to conduct and share their research, the lines between academia and industry, and research and practice, can start to blur. This will allow for faster innovation, broader sharing of knowledge, and ultimately, improved student learning outcomes.

