Looking Back to Move Forward

Co-Authors: Jeremy Roschelle and Adam Schellinger

January 2024

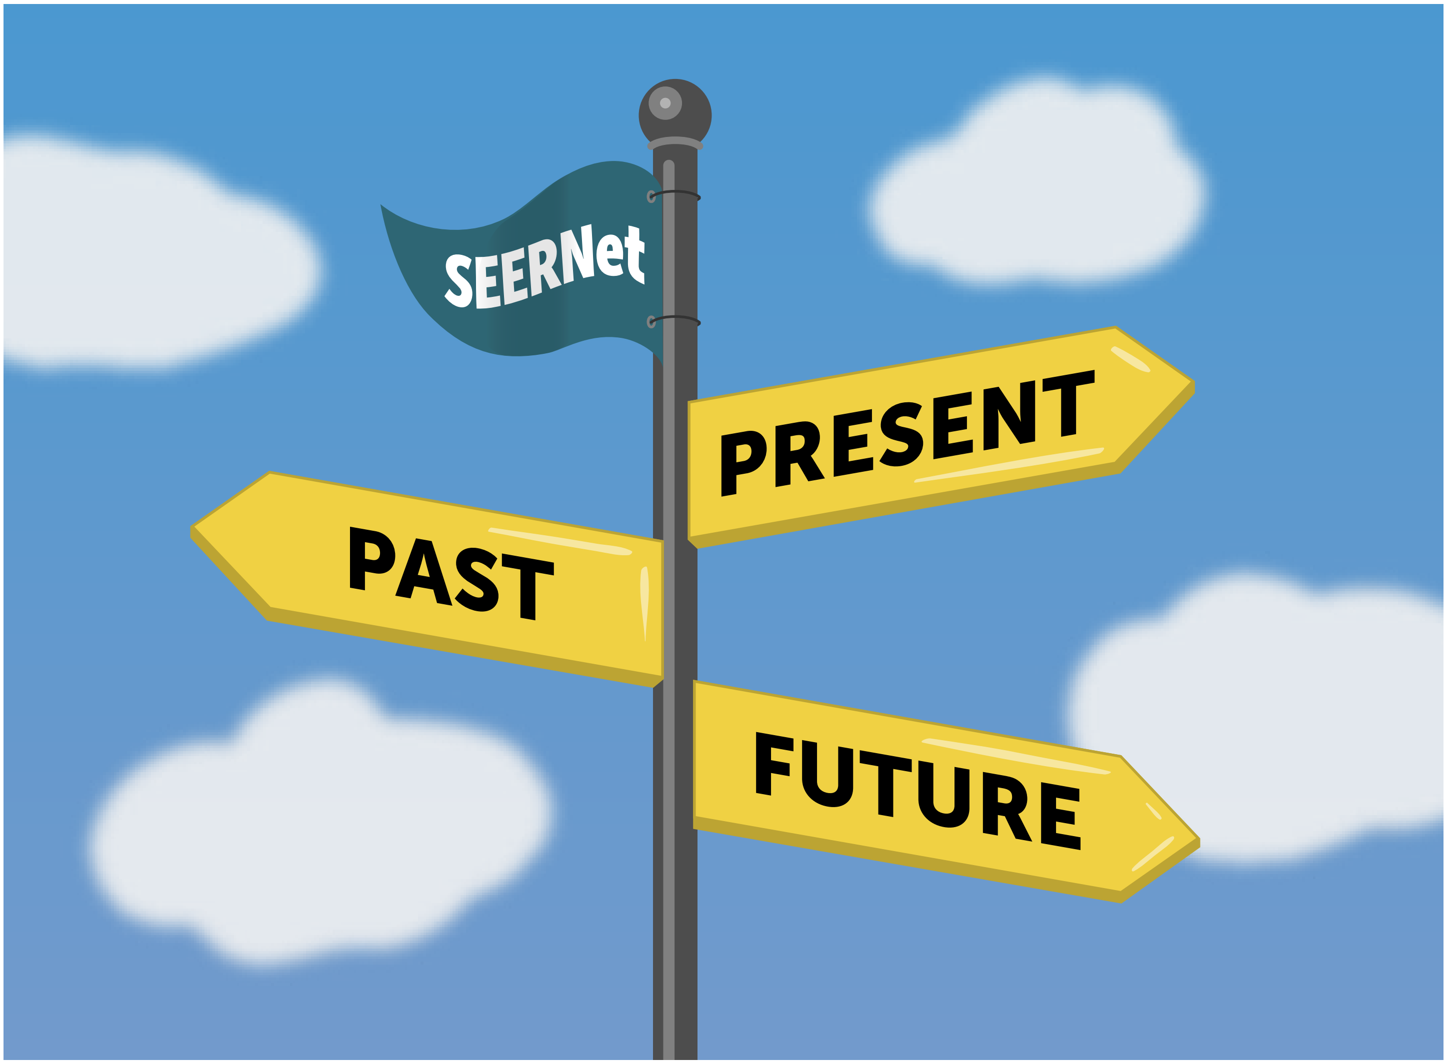

Digital Promise
Accelerating Innovation in Education
Empirical Education
IES Institute of Education Sciences

Acknowledgements

We would like to thank team members from each of the five SEERNet Digital Learning Platforms for their contributions and collaboration, as well as the Digital Promise and Empirical Education staff who provided valuable feedback and support.

Abstract

SEERNet digital learning platforms (DLPs) are developing new infrastructure to support research in authentic contexts where student learning is happening. In order to contextualize this work within the larger field, we trace historical precedents along four main categories: data repositories, data collection services, research design interfaces, and research communities. By situating this innovative movement alongside its predecessors, we can identify the opportunities for SEERNet and others to progress and sustain the mission of making research more scalable, equitable, and rigorous.

Suggested Citation: Schellinger, A., & Roschelle, J. (2024, January). *Looking Back to Move Forward.* Digital Promise. https://doi.org/10.51388/20.500.12265/203

This project is supported by the Institute of Education Sciences, U.S. Department of Education through Grant R305N210034 to Digital Promise. The opinions expressed are those of the authors and do not represent views of the Institute or the U.S. Department of Education.

Table of Contents

Introduction

SEERNet was funded by IES in 2021 as a community of five digital learning platform (DLP) teams (E-Trials/ASSISTments, UpGrade/Mathia, Terracotta/Canvas, OpenStax Kinetic/OpenStax-Rice University, and Learning at Scale/Arizona State University) connected by a community hub (SEERNet) and later joined by independent researchers. A key goal is to enable research on the large-scale DLPs that today's learners actually use, which would lead to a set of related benefits. For example, researchers could leverage the scale of the platforms to collect larger data sets. Platforms would be intentionally designed to make it easier to conduct research. And lessons learned from the research could yield both generalized knowledge and also practical improvements to tools already in wide use. The five platforms serve both K-12 and higher education and focus on important problems across a broad variety of learning challenges.

Presently, the project has been underway for two years. During this time the focus has been to develop the five platforms as research infrastructure while also building the community that enables the platforms to collaborate and share knowledge with each other and with external researchers. To date, each DLP has made advances to its capabilities and processes for working with external researchers. Meanwhile, the community has hosted and facilitated interactions among the platforms to understand the common challenges, as well as unique differences (Pautz, 2023). Now, by reflecting on the work, we can better understand what kinds of engineering have been required to create DLPs that are ready to serve as research infrastructure.

In this working paper, we take a step further back in order to move forward. We've found that by understanding the historical precedents that came before SEERNet, we can better describe the engineering work of developing a DLP to serve as research infrastructure. Further, we can better express some of the ongoing issues and opportunities that will require continued attention in the future.

We've found that SEERNet builds on four prior advances in educational research:

1. **Data Repositories** provide and regulate access to stored teaching and learning data, and also introduce the key constraint that researchers can only ask and answer the questions that the available data can reasonably address.

2. **Data Collection Services** automatically collect the inputs to and outputs from learning tools, while also introducing the constraint that the data must be collected relatively unobtrusively in the course of a learner using the DLP in their routine educational setting.

3. **Research Design Interfaces** standardize and support the process of deploying contrasting learning experiences to learners, and in conjunction with the data collection services, allow researchers to ask and answer research questions, but also introduce constraints on what types of contrasts can be proposed and in the scope of possible studies.

4. **Research Communities** build the capacity for researchers and practitioners to propose and conduct relevant research using specialized (and limited) capabilities described above, and recognize the constraint that it's not just about making research easier or faster, because there is a depth and breadth of knowledge that researchers must develop to use an existing DLP well.

We won't attempt a complete history of the four advances; our intention is not to conduct a literature review or research synthesis. However, we will provide well-known examples of each of the four advances, especially as each relates to digital learning platforms that large numbers of students use. And we'll highlight examples of important prior research conducted using each advance. Our point in doing so is to see SEERnet as embedded in a history of advances and trying to solve the same underlying problems:

- *How can independent researchers safely access existing data?*
- *How can learning process and short-term outcome data be collected efficiently?*
- *How can research design interfaces simplify the work of fielding an experimental contrast?*
- *And how can the field develop capacity to use these capabilities to ask and answer meaningful research questions?*

The history also helps us to express what may be especially unique about SEERNet. To our knowledge, it is the first and only time many independent DLPs have been intentionally funded to engineer tools that may address all four problems. The work of SEERNet members has been to engineer existing DLPs to create access to data, to provide data collection services, to provide interfaces to field research designs, and to build a community to support independent researchers to use these facilities. We'll provide examples from the DLPs and the community of what each of these engineering efforts looks like.

The history also helps us look forward. We envision DLPs-as-research-infrastructure as an important long-term direction for the educational research that is both more relevant to everyday teaching and learning and also more capable of analyzing realistic variability in the inputs, processes, and outputs to learning. Four important directions for the future are also latent in SEERnet's current work but not fully realized. For example, better ways to provide access to data can still be engineered (for example, through secure enclaves). Understanding how to ethically and unobtrusively collect the data that educational researchers most need is still not fully understood. Research design interfaces (RDIs) are oversimplified when described as being roughly the same as “A/B testing” interfaces in usability research, and yet it will take time to understand the nature of RDIs that yield generalizable, replicable insights about learning in heterogeneous environments. Finally, it requires intentional work to attract and support researchers to a new way of doing things, as new ways of doing research could make some things easier but also introduce additional constraints and requirements that need to be understood. Looking forward, we see the need for long-term, continued attention to the data repository, data collection services, research design interface, and research community aspects of the initiative.

This paper will trace the history of similar developments in educational research using DLPs, identify common themes that persist through the present, characterize the opportunities of the field for future work, and discuss the following questions:

- What are the long-term precedents for third party research using data from widely used DLPs?
- How do the five SEERNet DLPs address existing challenges inherent in this work?
- How does understanding the historical context of research on DLPs inform next steps for SEERNet?

Foundational Developments Preceding SEERNet

We identified four key building blocks (Figure 1) that have supported the creation of the SEERNet DLPs—data repositories, data collection services, research design interfaces, and research communities—and will review each in turn.

Figure 1 - Foundational Developments

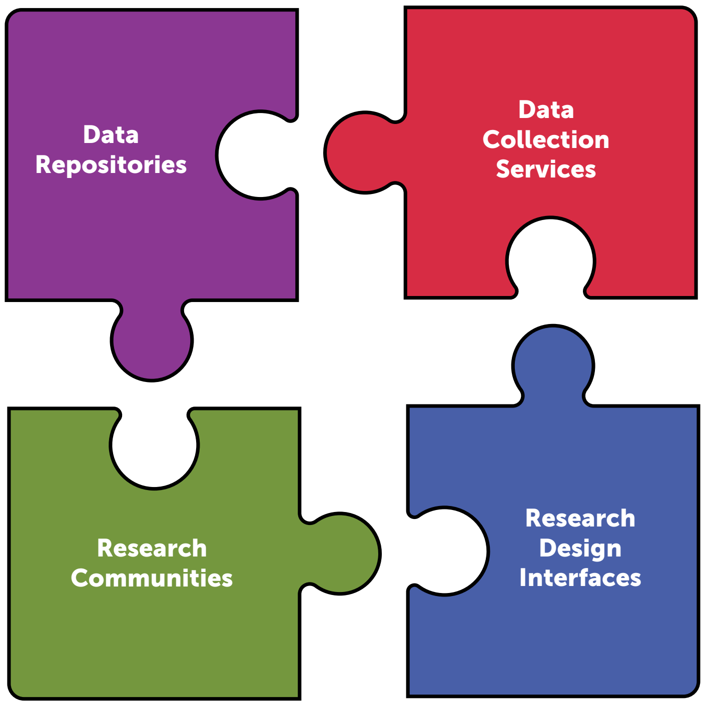

Data Repositories

Before we narrow the scope to DLP data, we must also acknowledge the longer history of secondary analysis in educational research. Collaborations such as the Inter-University Consortium for Political and Social Research (ICPSR) were created with the goal of housing and organizing similar types of educational and social science data and have also created communities of researchers, practitioners, and policymakers with a common resource, while at the federal level, assessment programs like National Assessment of Educational Progress and survey programs like National Longitudinal Surveys (NLS) have provided rich, nationally representative measures of academic achievement and student growth trajectories for researchers to analyze and link to other sources of data. This systematic approach is also exemplified by state longitudinal data systems that compile detailed data from multiple sources across large learner populations. The affordances of these foundational infrastructures have enabled researchers to advance education, as survey-based programs like these preceded by decades the types of repositories for data collected by DLPs that we now highlight.

In one example that is more specific to DLPs than those above, Carnegie Mellon University's DataShop was created out of a 2004 National Science Foundation grant to the Pittsburgh Science of Learning Center as an online repository and analysis platform designed to support research and analysis of educational data. Much of the early data came from early cognitive tutoring systems and used novel formats that allowed logging, formatting, and tagging of interaction data for analysis. DataShop now serves as a central hub for storing, sharing, and analyzing data collected from a variety of educational technology systems and tools; after 20 years, it now holds over 700,000 hours of student learning data across over 1000 unique data sets. Owing to this breadth, the types of research questions and study designs are as varied as the data themselves. From fine-grained interaction data from intelligent math tutoring systems to classifications of tones from foreign language learners, DataShop exemplifies open science principles as a hub for diverse stakeholders to share data and analysis tools across disciplines.

DataShop has been used across research studies in multiple domains. In an early example, Pavlik and Koedinger (2009) published work on using comparisons of learning curves to build a model of transfer between content domains. The study used data from the Bridge to Algebra Cognitive Tutor from Carnegie Learning. Another early iteration of a SEERNet DLP, ASSISTments, provided data to DataShop. The ASSISTments data was used to evaluate a new model for differentiating next-problem correctness predictions between varying sequences of problems within a particular skill (Wang, 2016). More broadly, researchers use learner data from digital resources at CMU and elsewhere to build models of student learning and iterate across different domains. In one important example, Koedinger and Stamper (2013) published work demonstrating that a cognitive model of problem-decomposition planning skills built from statistical data analysis could be used iteratively. Their DataShop-derived model was successful in improving learner mastery when used to redesign a cognitive tutor unit in geometry. The fine-grained interaction data obtained from cognitive tutoring systems are combined to form a cognitive model of knowledge components (KCs) that learners must possess in order to successfully solve a particular problem or complete an activity. This and many similar studies (Baker, 2010; Goldin, 2018) featured large samples of students and generalizations derived from the analysis of hundreds of thousands of problem solving steps. Thus we see that DataShop enabled efficient storage and retrieval of innovative types of learner data, which allowed research studies that advanced the field by generating new models of student learning and building out research tools across different content domains.

Learnsphere is another educational data infrastructure project led by researchers at Carnegie Mellon University's Human-Computer Interaction Institute that launched in 2014. Similar to DataShop, it aims to create an open and scalable platform for collecting, analyzing, and sharing educational data by focusing on addressing challenges related to educational data management, interoperability, and collaboration. The project is creating infrastructure to enable collaborative data sharing and analysis, supporting researchers, course developers and instructors. One example of research topics studied is curriculum pacing. Patel (2018, 2019) published findings and analytic tools for building visuals of how learners progress through segments within the curriculum of a particular course using historical data. The study used data from over 150,000 learners using a DLP and used plots to demonstrate visual differences between types of student learners, i.e., lockstep pacing vs. cram-to-complete. Building on the work of DataShop, LearnSphere added an open analytic method library, which enabled contributors to document and share new research approaches.

The Open Learning Initiative (OLI) is an outgrowth of both DataShop and LearnSphere which "aims to create a community of use, research, and development to allow for the continuous evaluation, improvement, and growth of courses and learning materials." The OLI project, encompassed by the larger Simon Initiative, is active in the Open Educational Resources community and supports better learning and instruction with high-quality, scientifically-based, classroom-tested online courses and materials. Data from OLI courses and other initiatives at the Pittsburgh Science of Learning Center are accessible through DataShop and LearnSpere. With regard to SEERNet, this approach to evidence-based instruction that in turn advances learning research is a Learning Engineering approach that continues to be developed in SEERNet's E-Trials/ASSISTments and UpGrade/Mathia.

Data Collection Services

One prominent online learning platform that offers a wide range of Massive Open Online Courses (MOOCs) and online degree programs is edX. It was founded in 2012 as a joint venture between Harvard University and the Massachusetts Institute of Technology (MIT) and has since grown to include partnerships with numerous universities, institutions, and organizations worldwide. It has served over 40 million learners since then and now offers over 3,000 courses—whether free, paid, credit bearing, or certificate granting—created by over 160 organizations. Designed from the ground up with research at the forefront, the terms of service enable both secondary data analysis and experimental studies.

The very first course offered by edX in March 2012 was "Circuits and Electronics," a joint offering from MIT and Harvard. Over 155,000 learners registered for the first instance of the course, and the data they generated were eagerly analyzed by researchers. Breslow and Pritchard (2013) looked descriptively at how learners spent time on the different course modules and examined the relationships between learner demographics and outcomes by evaluating predictive models for persistence and success in the course. This is a common theme in the types of studies enabled by edX and other MOOCs—the scale and scope of interaction data generated by DLPs can be used to classify learners along numerous dimensions.

These themes were again explored in one of the largest experimental studies done on edX. Kizilcec and Reich (2020) tested multiple interventions on a sample of over 270,000 learners over a 2-year period. After learners took an initial survey, they were randomly assigned to one of several established behavioral interventions with medium-to-large effects in prior studies of persistence and completion in online courses, or to a control condition. One example was a value-relevance intervention—learners identify important values and reflect on how the course may connect to achieving these goals. While the study was not able to replicate the effect size of impacts seen in earlier studies, subsequent exploratory analyses identified important contextual considerations that inform practice moving forward.

The researchers worked to operate under open science principles emblematic of our guiding SEER standards, pre-registering their analyses and publicly sharing analysis and code. (This paradigm is carried through in SEERNet's Kinetic). By enabling researchers' extreme flexibility in implementing different types of interventions across open educational resources where learning is actively happening, DLPs massively lower the cost of recruitment, allowing experiments that run authentically at scale.

Another example is a study from 2015 that used existing course and interaction data from ASU to evaluate guided learning. While the courses developed were touted as a revolutionary way for students to gain course credit for participation in newly revamped MOOCs, a very small percentage of students completed the course requirements (Reich, 2017). Overall, we see that edX both generated massive amounts of learner data and opened its platform to active experimentation, allowing researchers to validate new statistical models from experimental data all while preserving learner data privacy.

Research Design Interfaces

The technical advances that made it easy to store and share large amounts of fine-grained, longitudinal, and extensive data from student interactions with technology enabled corresponding advances in research methods In particular, DLPs were able to create specific interfaces for researchers to conduct experimental comparisons between two or more conditions. These interfaces make it easier to instrument comparison conditions in the DLP, to randomly assign the variations to a population of students, and to collect data. In one sign of the importance of research design interfaces to research communities, SEERNet and other researchers have hosted four annual workshops at the Learning@Scale conference about A/B testing and platform-enabled research.

ASSISTments provides one important example of a RDI within SEERNet. The mission statement for ASSISTments is “to improve education through scientific research while not compromising student learning time.” Designed from the ground up as a tool for teaching and learning that would produce data for re-searchers, the platform originated in 2003. While built as a tool for teachers and students to improve math instruction, the platform now counts students, practitioners, and researchers among its users. “We now think of ourselves more as the creators, and maintainers, of an instrument that others can use to do their science. The scientists who use the tool do not have to be the ones who made the tool” (Heffernan, 2014). One study using ASSISTments demonstrated the impact of spacing in mathematical expressions (Harrison, 2020), showing that students performed better when spacing matched the order of precedence.

The MOOC Replication Framework (MORF) is a close contemporary of SEERNet that aims to facilitate more effective use of the data generated by online platforms for replication across studies, courses, and data sets (Gardner, 2018). The platform provides a large, diverse data set collected from multiple MOOC offerings and a technical solution for running predictive models in a controlled environment. There are also parallels to SEERNet in the development of infrastructure built to offer evaluation as a service. By linking student-level data from edtech platforms to demographic and achievement data obtained directly from educational institutions, all while protecting student privacy, these tools aim to improve the efficiency of efficacy studies evaluating the impact of edtech on outcomes of interest to local stakeholders and the broader field. LearnPlatform and Empirical Education developed tools that can do this type of work, such as evaluations of an early reading program or a teacher productivity tool.

Research Communities

The term "educational data mining" was introduced in the early 2000s and the International Educational Data Mining Society was founded in July 2011. Stamper and Ritter discussed the rise of the "super experiment" in 2012, indicating the shift from lab scale (under 100 students) to internet scale (more than 10,000 students). Also in 2011, the Society for Learning Analytics Research (SoLAR) was founded, and the first International Conference on Learning Analytics and Knowledge (LAK) took place, while the Learning@Scale community and its accompanying conference began in 2014. Learning@Scale arose out of the growth of MOOCs but attempted to broaden the scope of how teaching and research on teaching could take advantage of the scale and types of data being generated by DLPs. The field of learning analytics has grown to encompass a wider array of data analysis techniques and approaches beyond just data mining, and the workshops on A/B testing that began in 2020 have shown considerable growth in the number of participants and papers in the succeeding years. A paper from 2021 by Jensen and Umada used data from nearly 50,000 students in Florida that used AlgebraNation, another math DLP, to analyze the benefit of retrieving learned content using nudges for short, formative quizzes compared to more passive modes of learning. This year's conference included a study involving nearly 20,000 students with over 15 million learning events that aimed to replicate the "doer effect"—the finding that those who actively do practice while reading show better outcomes than those who only read (Campenhout, 2023).

One important way in which research fields, societies and communities add to the DLPs-as-research-infrastructure movement is by focusing attention on important scientific and practical challenges. For example, Baker (2019) used a keynote address at the Learning Analytics and Knowledge Conference to challenge the community to work on important, unsolved problems. One unsolved problem is creating generalizable "detectors" of students' affective states, such as boredom, while they use a DLP. Knowing when a student is bored with educational content in a DLP could allow for targeted interventions to improve their engagement and effort, and although affective detectors have been built for specific platforms, no generalizable approach exists to date.

How Do the SEERNet DLPs Address Existing Challenges?

The goals of SEERNet are to create technical solutions that tackle the challenges of doing education research in authentic learning contexts at scale, and to further the field by growing a community of developers and researchers. The work is guided by the SEER principles: committing to open science through pre-registration; increasing access to data, methods, and findings; examining inequities; and generalization and scaling. The challenges in this work are not new and will not be easily solved: how to protect learner privacy when storing and sharing detailed data from multiple sources, how to transform granular learning outputs into meaningful outcomes, how to experimentally manipulate learning activities while preserving authentic classroom environments, how to grow a scientific community while working towards a common goal. This movement is built on the work of countless researchers, learning scientists, data scientists, and practitioners that have come before, and SEERNet is addressing present challenges across each of the four domains we highlighted, as discussed below.

Data Repositories to Provide Safe and Secure Access

Two of SEERNet's DLPs in higher education focus on privacy-protecting data access. First, ASU's Learning at Scale is increasing access to institutional data from ASU for external researchers. Navigating across departments with unique review processes and policy requirements such as FERPA and GDPR, the project must balance preserving learner privacy with sharing the rich cross-sectional sources of data of interest to researchers. This raises challenges that are both administrative and technical, and involves a diverse set of stakeholders.Second, the Kinetic team is developing a secure enclave, which would allow researchers privacy-protecting access to student-level information to run their analyses virtually without exposing the underlying data. This technical approach is of great interest to those with equity-relevant research questions, echoing the prior work published under SEERNet's central hub remarking on the tension between equity-relevant research and open science and the importance of centering practitioners (Pautz Stephenson, Banks & Pakhira, 2022; Zacamy, Roschelle, 2022). The three other DLPs are also exploring how to answer questions about what works for whom, given the different types of contextual information already available in their platforms, all in ways that actively protect students' identities.

Data Collection Services Link Hypotheses to Research Insights

SEERNet DLPs are working on sharing data that enables measuring the relationships among varied inputs, observable processes, and immediate outcomes. E-Trials and UpGrade, particularly, have the ability to expose granular data from their DLPs to trace learning trajectories over time. Both platforms have a rich history of research to draw upon—including experimental research, efficacy studies, and secondary data analysis—and provide rich, hierarchical data to access when investigating prior learner performance. With E-Trials, researchers can use content in ASSISTments' content library and create studies without having to recruit participants. If researchers use their own content, they will need to recruit their own subject pool. By centering the teacher in its research design process, Terracotta enables the exploration of authentic, relevant research questions. The two aforementioned higher ed platforms also provide data collection services.

Research Design Interfaces Enabling Testing Conjectures

All five SEERNet DLPs are engineering RDIs. A comparison of what they offer is available on the SEERNet website. By way of example, we'll focus on Terracotta and Kinetic, each of which open research on widely used platforms (Canvas and OpenStax, respectively) to instructors and researchers. Terracotta puts the teacher in the role of researcher. Within their own Canvas course, instructors have the ability to implement countless types of activities to manipulate different inputs that are relevant to their unique learning context. Centering practitioners affords the power to rapidly test hypotheses around a particular assignment and the flexibility to replicate a similar intervention in another activity or course, or with a different sample. Kinetic has also designed their system with a high degree of flexibility in the types of interventions that can be implemented and outcomes that can be assessed. By utilizing Qualtrics as the back end, researchers can incorporate interventions that aren't limited to the structure and offerings of the underlying OpenStax resources and can collect outcomes across different domains.

Research Communities Building Capacity to Ask and Answer Relevant Questions

SEERNet's capacity-building work—including webinars for interested external researchers, conference presentations and papers; its practitioner advisory and general advisory boards; and web/social media presences—is all aimed at growing and diversifying the research community and generating relevant, actionable research. Increasing the feedback loops between each of these separate systems and documenting the processes along the way should generate a sustainable system of interested parties that endures beyond the initial IES grant phase. We see this work as fitting in with and contributing to existing fields and societies—EDM, SoLAR, Learning @ Scale, and others—that gather large numbers of relevant researchers. For example, SEERNet members have offered workshops, papers and other opportunities to the broader fields at these conferences, and SEERNet welcomes members of these fields to propose studies. Research studies can be proposed to IES, NSF or any other funder.

Informing Next Steps and Growing the Community

In looking backwards, we have found value in seeing the work of SEERNet in terms of the interconnections among four trends that have emerged in research over time:

1. Data repositories
2. Collecting DLP data
3. Enabling experimental research via RDIs
4. Growing a community of researchers

This analysis helps us more clearly see where and how SEERNet can contribute to a longer-term movement. For example, new and stronger mechanisms to enable equity-relevant research while protecting student privacy could contribute to the concept of data repositories. Likewise, to enable research studies, SEERNet enables collecting and interrelating three kinds of data: data about students, learning process data, and short-term learning outcome data. Advances in collecting all three types of data automatically could form a contribution of SEERnet to broader research movements. SEERNet researchers are also pursuing innovation in RDIs and by comparing these interfaces side-by-side, we may be able to inform others who wish to provide research interfaces to their DLPs.

The analysis also helps us see where SEERNet has limits or challenges. Research is limited by the kinds of data to which DLPs provide access and this can make equity-relevant research hard. Research can be limited by the kinds of data DLPs are willing to collect on behalf of the researcher. And research can be limited by constraints built into the RDI. These are all technical challenges which can be overcome in due time. One important longer term challenge is discussed next: for all this engineering to pay off, we need to grow the research community that wants to use these DLPs as infrastructure for their research.

Building a community is hard, and growing and diversifying that community is an ongoing challenge. An implicit benefit of the community approach of SEERNet is the opportunity for more perspectives to be heard in the development of tools and infrastructure across the platforms. With the central SEERNet advisory board, complemented by the advisory boards of each platform, we can ensure that best practices and learnings from the product-development and research processes are improved and shared among the broader group, documenting not just deliverables and outcomes but challenges and improvements along the way. Rather than several monolithic research studies conducted in a vacuum, the community structure of SEERNet will allow platform developers, researchers, and other stakeholders to build knowledge and amplify findings across all of the rich work being conducted. Yet, each platform is unique in the types of content and learners they support. The condensed time period of the IES research grants includes a 6-month feasibility phase where platform and research teams collaborate to determine the availability of background and outcome data, implementation of interventions, and tools for data processing and analysis. Given the goal of SEERNet to utilize open science principles, some of the DLPs are building workflows around OSF.io to ensure durable storage of not only data and results but also context and process. This resource is not specific to the education sciences like DataShop or LearnSphere, as there are now a variety of ways to preserve and share similar work, so the community can collaborate on finding the optimal audiences for each. How can we identify successful findings and scale them up across larger samples or address similar research questions across platforms?

The community is also focused on opening these opportunities to early-career researchers. Going beyond just a yearbook of studies, summarizing and sharing the infrastructure behind the platforms may enable efforts similar to Ryan Baker's MORF, a framework for addressing challenges inherent in replicating and scaling research across experiments across platforms and contexts. Additionally, we can imagine what it would look like for other platforms and researchers who wish to become involved in SEERNet. Similar efforts, like the Behavioral Intervention Research Infrastructure led by Kizilcec and Baker, will be contemporaneously experimenting with allowing external researchers access to the Realizeit digital learning platform to study interventions across a wide variety of large colleges and universities. Researchers may aim to conduct similar experiments across platforms, utilizing the affordances of each, or investigate a common theme using different research designs. Additionally, programs like the IES Training Grant will ensure that the learnings from SEERNet and similar projects are taken up by the next generation of learning scientists, AI/ML data scientists, and related researchers. By outlining the history of similar efforts and defining the type of themes and challenges addressed by SEERNet, we seek to attract the participation and insight of others who have done, are doing, and hope to do the type of work discussed here and across the community, united under the goal of building a community dedicated to improving teaching and learning for all students. Beyond the 5-year horizon of its initial funding, how will the lessons learned during these studies be applicable across different contexts, both within and outside of the traditional IES research cycle? Interested researchers and practitioners can stay informed about our work and join the effort by visiting SEERNet.org or joining our interest list.

Conclusion

The roots of the SEERNet DLPs and the underlying movement to make education research more relevant, efficient, and scalable traces back over multiple decades of learning science research, educational data mining, and open science frameworks. While some of the tools being built are brand new, many of the underlying people and processes have existed in one form or another during that entire time frame. As we continue to grow the community, we take this opportunity to acknowledge the achievements of the researchers and developers who came before and situate current research in its larger context. As we grow this movement and onboard the initial cohort of external researchers, the community, its funders, and the larger field will expand opportunities to participate in SEERNet, to leverage the DLPs as research infrastructure, to share expertise and and to work to common solutions to the challenges that stand in the way of full realization of the concept of Digital Learning Platforms as Research Infrastructure.

References

Baker, R. (2019). Challenges for the Future of Educational Data Mining: The Baker Learning Analytics Prizes.

Baker, R. S., D'Mello, S. K., Rodrigo, M. M. T., & Graesser, A. C. (2010). Better to be frustrated than bored: The incidence, persistence, and impact of learners' cognitive–affective states during interactions with three different computer-based learning environments. *International Journal of Human-Computer Studies, 68*(4), 223–241.

Baker, R. S., & Yacef, K. (2009). The state of educational data mining in 2009: A review and future visions. *Journal of Educational Data Mining, 1*(1), 3–17.

Gardner, J., Brooks, C., Andres, J. M., & Baker, R. S. (2018). MORF: A Framework for Predictive Modeling and Replication At Scale With Privacy-Restricted MOOC Data. *2018 IEEE International Conference on Big Data (Big Data)*, Seattle, WA, USA, 2018, 3235-3244.

Goldin, I., & Galyardt, A. (2018). Most of the time, it works every time: Limitations in refining domain models with learning curves. *Journal of Educational Data Mining, 10*(2), 55-92.

Harrison, A., Smith, H., Hulse, T., & Ottmar, E. (2020). Spacing out!: Manipulating Spatial Features in Mathematical Expressions Affects Performance. *Journal of Numerical Cognition, 6*(2), 186-203.

Jensen, E., Umada, T., Hunkins, N., Hutt, S., Huggins-Manley, A., & D'Mello, S. (2021). What you do predicts how you do: Prospectively modeling student quiz performance using activity features in an online learning environment, 121-131.

Kizilcec, R., Reich, J., Yeomans, M., Dann, C., Brunskill, E., Lopez, G., Turkay, S., Williams, J., & Tingley, D. (2020). Scaling up behavioral science interventions in online education. *Proceedings of the National Academy of Sciences, 117*(26), 14900-14905.

Koedinger, K., Cunningham, K., Skogsholm, A., & Leber, B. (2008). An open repository and analysis tools for fine-grained, longitudinal learner data. In R. S. J. d. Baker, T. Barnes, & J. E. Beck (Eds.), *Educational Data Mining 2008: 1st International Conference on Educational Data Mining, Proceedings*. Montreal, Quebec, Canada. June 20–21, 2008, 157–166.

Koedinger, K. R., Stamper, J. C., McLaughlin, E. A., & Nixon, T. (2013). Using data-driven discovery of better student models to improve student learning. In H. C. Lane, K. Yacef, J. Mostow, & P. Pavlik (Eds.), *Artificial Intelligence in Education. AIED 2013*. Lecture Notes in Computer Science, vol 7926. Springer, Berlin, Heidelberg.

Hansen, J. D., & Reich, J. (2015). Democratizing education? Examining access and usage patterns in massive open online courses. *Science, 350*(6265), 1245–1248.

Heffernan, N. T., & Heffernan, C. L. (2014). The ASSISTments ecosystem: Building a platform that brings scientists and teachers together for minimally invasive research on human learning and teaching. *International Journal of Artificial Intelligence in Education, 24*, 470–497.

Motz, B. A., Üner, Ö., Jankowski, H. E., et al. (2023). Terracotta: A tool for conducting experimental research on student learning. *Behavior Research*.

Ostrow, K. S., Heffernan, N. T., & Williams, J. J. (2017). Tomorrow's edtech today: Establishing a learning platform as a collaborative research tool for sound science. *Teachers College Record, Volume 119*(3), 1–36.

Patel, N., Sharma, A., Sellman, C., & Lomas, D. (2018). Curriculum pacing: A new approach to discover instructional practices in classrooms. In R. Nkambou, R. Azevedo, & J. Vassileva (Eds.), *Intelligent Tutoring Systems. ITS 2018*. Lecture Notes in Computer Science, vol 10858. Springer, Cham.

Pautz Stephenson, S., & Roschelle, J. (2023). Making Waves: Reflections on SEERNet's Progress and a Vision for a New Generation of Education Research. *Digital Promise.*

Pavlik Jr, P., Cen, H., & Koedinger, K. (2009). Learning factors transfer analysis: Using learning curve analysis to automatically generate domain models. *International Conference on Educational Data Mining*, 121–130.

Prihar, E., Syed, M., Ostrow, K., Shaw, S., Sales, A., & Heffernan, N. (2022, July). Exploring common trends in online educational experiments. In *Proceedings of the 15th International Conference on Educational Data Mining*. International Educational Data Mining Society.

Ritter, S., Murphy, A., Fancsali, S. Lomas, D., Fitkariwala, V., & Patel, N. (2020). UpGrade: An open-source tool to support A/B testing in educational software. *L@S Workshop on A/B Testing at Scale*.

Romero, C., & Ventura, S. (2010). Educational Data Mining: A Review of the State of the Art. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews), 40*(6), 601–618.

Siemens, G. (2013). Learning analytics: The emergence of a discipline. *American Behavioral Scientist, 57*(10), 1380–1400.

Stamper, J. C., Lomas, D., Ching, D., Ritter, S., Koedinger, K. R., & Steinhart, J. (2012). The rise of the super experiment. *International Educational Data Mining Society*.

U.S. Department of Education. Institute of Education Sciences. (2022, June 30). *Standards for Excellence in Education Research*. https://ies.ed.gov/seer/

U.S. Department of Education, Office of Educational Technology. (2012). *Enhancing Teaching and Learning Through Educational Data Mining and Learning Analytics: An Issue Brief*. Washington, D.C.

Wang, Y., Ostrow, K., Adjei, S., & Heffernan, N. T. (2016). The Opportunity Count Model: A Flexible Approach to Modeling Student Performance. In *Proceedings of the Third (2016) ACM Conference on Learning @ Scale (L@S '16)*. Association for Computing Machinery, New York, NY, USA, 113–116.