Instrumenting Courseware and Leveraging Data with the Open Learning Initiative (OLI)

Author(s): Norman Bier
Carnegie Mellon University
nbier@cmu.edu

Author(s): Stephen Moore
Carnegie Mellon University
stevenmo@andrew.cmu.edu

Author(s): Martin Van Velsen
Carnegie Mellon University
martinv@andrew.cmu.edu

ABSTRACT: Founded in 2002 as part of the Hewlett Foundation's inaugural open education grants, the Open Learning Initiative (OLI) is a recognized leader in adaptive courseware and learning engineering, combining leading research in cognitive and learning science with state-of-the-art technology to create adaptive, open courseware that enacts instruction. By rigorously capturing and evaluating learner data, OLI drives powerful feedback loops that assist learners, support educators, improve courses, and drive learning science research. This workshop will provide an overview of creating instrumented courseware with OLI's tools, aligning measurable, student-centered learning outcomes with active learning activities and assessments. We will provide examples of the data generated by OLI learner interactions and show how this data is used to provide feedback to learners and drive analytics for both instruction and course improvement. Finally, we will show how OLI data is made available for research, teaching participants how to access this information and providing examples of how this data has been used to support primary research, secondary analysis, and ongoing analytics work. Participants will leave with the ability to build their own OLI courses, the ability to access OLI data for their own work, and contacts for ongoing engagement with the OLI team.

Keywords: OER, Instrumented Courseware, Iterative Imrovement, Learning Engineering, Learning Analytics

1 INTRODUCTION – ENGINEERING LEARNING

The Open Learning Initiative (OLI) serves as a combination research and development project at Carnegie Mellon University (CMU), integrating with the larger work of the university's Simon Initiative. OLI focuses on developing, using, improving, and researching science-informed, open courseware as a key element of a community-based research activity focused on understanding and improving human learning.

Central to the Initiative is an approach, born at CMU, that Nobel Laureate Herbert Simon dubbed *Learning Engineering*: the use of learning research and the affordances of technology to design and deliver innovative, instrumented educational practices with demonstrated and measurable outcomes. This close integration of research, data, and instructional practice contrasts with the approaches of many other institutions, where instructional design is frequently based on intuition rather than research, and where technology is often implemented for its own sake rather than as a reasoned,

Creative Commons License, Attribution - NonCommercial-NoDerivs 3.0 Unported (CC BY-NC-ND 3.0)

supportive part of a larger instructional research agenda. From its home in the Simon Initiative, OLI offers an exemplar of the success of the learning engineering approach.

2 THE OPEN LEARNING INITIATIVE

Founded in 2002 as part of the Hewlett Foundation's inaugural, pioneering open education grants (Kernohan, & Thomas, 2018), OLI is a recognized leader in adaptive courseware, learning engineering, and open education, combining leading research in cognitive and learning science with state-of-the-art technology to create adaptive, open courseware that enacts instruction. By rigorously capturing and evaluating learner data, OLI drives powerful feedback loops that assist learners and educators, improve courses, and contribute to our larger understanding of how humans learn. Developed by multi-disciplinary teams, OLI courses can be used to support independent learners, but are primarily designed to support a hybrid instructional model and toolset that maximizes faculty time and expertise. This approach makes OLI unique in the open educational resources (OER) space; while many open projects focus on a loose collection of openly licensed assets, or on developing static OER textbooks, OLI's courseware offers a fully designed learning experience. This experience combines expository content, dynamic activities, and specialized technologies (including labs, simulations, tutors, and other domain-specific learn-by-doing activities). While the expository materials can be downloaded from OLI to create a traditional OER textbook, the complete courseware offers a much greater set of benefits. Data from learners' interactions with these activities, in conjunction with the model of expertise developed as part of the course's design, supports a wide variety of opportunities to adapt to learners' needs. These can include targeted feedback and hints that address demonstrated learner misconceptions, as well as sequencing of problems and activities based on learner achievement, all presented within the context of developing better metacognitive skills and awareness on the part of the student. This same information supports faculty as they design their classroom instruction, with an advanced analytics dashboard that provides detailed learning estimates in relation to the skills and learning objectives specified in the cognitive model. Many analytic systems focus only on engagement or performance metrics; the OLI dashboard estimates learning based upon all aspects of a learner's interactions with assessments (Lovett, 2012). In addition to the benefits for learners and classroom educators, these data also offer benefits in the aggregate, providing insights on course performance that can support faculty in empirically improving the design of a course over time.

2.1 OLI Results

Extensive research has demonstrated the success of the OLI approach in postsecondary education. Studies show dramatically improved outcomes, savings in cost and time, and improved learning productivity over time. Perhaps the best known of this work has focused on the use of the OLI Statistics course; this accelerated learning study demonstrated improved outcomes for CMU learners spending less than half the time of their traditional peers (Lovett et al., 2008). Studies of OLI in collaboration with larger public universities have also demonstrated the scale interventions to large numbers of learners, improving outcomes while lowering costs (cf., Bowen et al., 2014; Griffiths et al., 2014). Recent studies have found that the impact of OLI's learn-by-doing activities can be six times that of other instructional approaches (Koedinger, Kim, Jia, McLaughlin, & Bier, 2015), and follow-on studies have indicated that this doer-effect is both causal and is observable in a multiple number of domains and learner contexts (Koedinger, Jia, McLaughlin, & Bier, 2016).

Creative Commons License, Attribution - NonCommercial-NoDerivs 3.0 Unported (CC BY-NC-ND 3.0)

Ongoing studies continue to investigate the role of OLI with different learner populations, and results suggest that the use of OLI activities can help to smooth out expected negative outcomes often associated with vulnerable and under-prepared learner populations (Evans, Leinhardt, & Yaron, 2008; Kaufman, Ryan, Thille, & Bier, 2013; Ryan, Kaufman, Greenhouse, She, & Shi, 2016). Over the past decade, 40 OLI courses have seen enrollments from over four million independent learners. These same courses have been used to support academic classes in hundreds of institutions of higher education and high schools, with more than 500,000 enrollments in these types of credit-bearing contexts. This effort has also contributed to extensive research in understanding how human beings learn, including the generation of hundreds of learner interaction datasets that have been used for primary and secondary analysis. This represents an exceptional community of educators, learners, and researchers with whom workshop participants can engage (and who will benefit from the project).

2.2 Scaling OLI Course Development

As the OLI project has grown, it has become increasingly clear that the need for large teams, extended timelines, and deep technical expertise has been a barrier for scaling the community involved in OLI course development (Herckis & Smith, 2018). Similarly, though learning engineering tools and approaches to leveraging data for iterative course improvement are remarkably sophisticated, these tools have often required more time and expertise to implement than is reasonable for most faculty. This challenge has been compounded by the multiple systems and interfaces required to leverage these improvement tools (Bier & Jerome, 2012). To address this need, OLI has made significant investments in developing an integrated authoring suite to support a broader community in the development, improvement, and refinement of open courseware for the OLI system. Preliminary development efforts focused on easy, WYSIWYG authoring capabilities that allow any faculty member to easily develop OLI course materials, not merely as a set of content, but as an integrated learning experience that provides appropriate semantic context to the materials and supports the easy tagging of skills and learning objectives to all learning activities, providing a foundational cognitive model of expertise for the course.

Subsequent development has focused on upgrading this suite into a more thorough workbench for supporting all faculty in learning engineering. This has focused on two major components: 1) embedding into the system elements of instructional design intelligence and learning engineering support that have traditionally been provided via human consultation, thereby scaffolding the authoring process to encourage best practices for learning from the beginning; and 2) embedding analytics for course improvement directly into the authoring view, making learning data actionable for faculty and lowering barriers for continuous, iterative improvement.

2.3 OLI Course Improvement Analytics

These **course improvement analytics** build on the successful prototypes developed under NSF Grant 1418244 (Data-Driven Methods to Improve Student Learning from Online Courses) and provide a range of insights into the underlying design and effectiveness of the course. These improvement analytics include three core elements:

- **Course Design Analytics,** showing the breadth of learning activities and assessment opportunities in relation to the skills and learning objectives that constitute the cognitive

Creative Commons License, Attribution - NonCommercial-NoDerivs 3.0 Unported (CC BY-NC-ND 3.0)

model of the course. This view supports improvements in the robustness of the course's design and can be used even before student learning data is available. Such use supports more effective preliminary design and ensures that the data gathered from student use will offer a fuller set of actionable improvement opportunities.

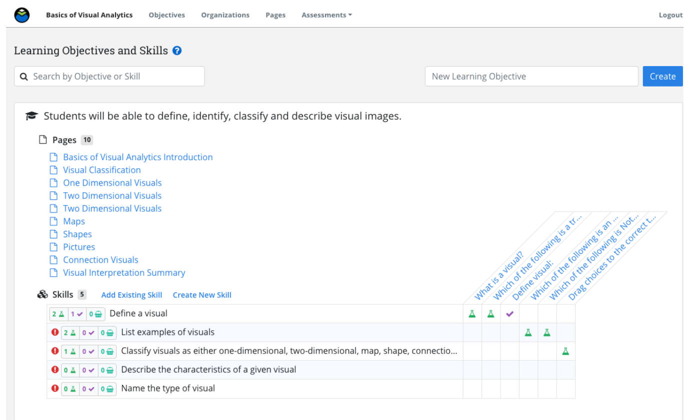

Figure 1: Analytics for Course Design

- **Effectiveness Analytics,** offering insights into both larger learning activities and individual questions. These analytics include traditional item-response theory (IRT) models, difficulty analysis, and views of student use and engagement patterns. These views can also provide insight into the role of individual activities in relation to the larger learning objectives and skills with which they are associated. The interface supports improvement and investigation at a variety of scales and levels of detail, from course-level ("Show me the units with the largest disconnect between practice and exam success") to objectives ("Which objectives are students not succeeding in") to individual questions ("This question is one that most students are not getting correct, even after multiple attempts"), and includes summary-level dashboards and analytics embedded directly in the authoring interface.

- **Cognitive Model Analytics,** building upon ongoing work at the Simon LearnLab (Pittsburgh Science of Learning Center) to understand learning model behavior and identify mismatches between expected and actual learner behavior, offering opportunities to improve the underlying model of expertise that is represented by the course. These elements build on decades of tools and methods for optimizing learning through cognitive model discovery and refinement (Stamper & Koedinger, 2011), particularly Learning Curve analysis, a method to identify latent variables in a logistic regression model called the Additive Factors Model (AFM), which is a generalization of IRT (e.g., Wilson & de Boeck, 2004)

By embedding these analytics tools directly into the authoring interface, we make it more likely that instructors will use them. By carefully leveraging design and user experience expertise from the Simon Initiative community, we build these tools so that any faculty member can interpret and act upon the insights they provide. And by engaging with a larger community, these tools are successfully tested, improved, and used by faculty at a diverse array of institutions (Shestak, 2017; Richie, 2018). Together, this suite enacts core elements of instructional design, guiding and scaffolding authors in the development of instructional materials that are as robust as possible, which will provide sufficient data to engage in an iterative improvement process which leverages that data for an empirical approach to course improvement.

Creative Commons License, Attribution - NonCommercial-NoDerivs 3.0 Unported (CC BY-NC-ND 3.0)

3 OLI LEARNING ACTIVITIES AND DATA

The OLI platform is a collection of tools for creating and delivering online instruction that embeds core learning science principles in the system's design, capabilities, and navigation. Content in the system combines *structure*, *learning objectives,* and traditional *expository materials* (text, examples, images, videos, etc.) with *native activities*—learn-by-doing interactions which offer practice, targeted feedback, and robust hints. Together, these components provide a structured, complete, and supported learning experience. As part of the course development process, the *semantic context* for each of these elements is also captured; OLI defines a learning taxonomy using a series of DTDs[1] that provide additional structure to the learning environment and capture the pedagogical intent of specific components. For example, exposition is captured not merely as a series of textual elements but rather is specified as worked-examples, theorems, learn-by-doing opportunities, self-assessments, and many other semantic elements. This semantic context informs the data that is collected from learns' use of the course, allowing for more meaningful research and analysis than that offered by more free-form design and click-stream collection approaches. The design of the system has been further enhanced by UDL principles to increase flexibility, address learner variability, and allow learners multiple ways to recognize, act on, and engage with knowledge. These pre-defined capabilities may not always provide the full capabilities necessary for new approaches, domain-specific activities, or experiments. Therefore, the system also provides mechanisms for incorporating other non-core technologies, via APIs. Such non-core technologies include standard elements that are used frequently in courses, including certain types of labs, simulations, and cognitive tutors. These technologies can also include less standard, more experimental elements; as technologies and their associated pedagogical approaches become less experimental and better tested, their use becomes more standardized, eventually moving towards integration with the core system.

3.1 Native OLI Activities

Expository content forms an important part of OLI learning environments, but more important are OLI's native activities. These active learning activities provide students with opportunities to answer questions and solve problems, with targeted feedback and help. By aligning these activities with the course's student-centered, measurable learning outcomes, the OLI system is able to continuously assess student learning. These activities support a range of machine-evaluated question types (multiple choice, fill-in-the-blank, short answer, multi-select, ordering, hot-spot, and others), along with feedback and hints. Formative assessment within the OLI system is provided in-line, with these activities presented within the flow of other expository elements. Such activities are considered "low-stakes" within the OLI context—learners do not receive a score for these activities, activities support an infinite number of attempts, and instructors are unable to see the specific results of an individual student's success or failure for a given activity. (Instructors are able to see that a student has completed an individual activity, and they are presented with an aggregate view of their class's success; in this way, low-stakes activities present students with a "safe space" to practice, without being penalized for mistakes.) Low-stakes activities have two different semantic contexts, called purpose types, based on the pedagogical intent of the activity: **Learn By Doing (LBD)** activities are

[1] http://oli.cmu.edu/dtd/

Creative Commons License, Attribution - NonCommercial-NoDerivs 3.0 Unported (CC BY-NC-ND 3.0)

inserted to provide students with opportunities to practice or master new knowledge and skills, and assume that the learner will make mistakes along the way. **Did I Get This (DIGT)** activities are presented as self-assessment opportunities, provided at points where it is anticipated that the learner should have mastered a specific skill (and offering additional context to support the learner in metacognitive skills development and guidance for self-remediation, if needed). By design, the student is assumed to not yet have mastery, so the exercises should be tailored to include some instruction and reinforcement through the question content itself and from the immediate feedback offered for correct and incorrect answers. Questions can include hints to provide support to students as they learn. OLI adds one to three hints where appropriate, following this pedagogical rule-of-thumb: (1) Restatement: What is the question asking? (2) Cognitive hint: Here are steps you should take. (3) Bottom-out hint: Used in numeric or text-input only, where the student may not be able to get to the correct response or feedback on their own. This same approach also supports summative assessment—high-stakes **Quizzes** and **Checkpoints**—which provide more detailed scores and information for instructors and can be used to calculate grades in formal learning environments.

Compared to many online systems, such as learning management systems which focus on collecting navigation and clickstream data, OLI's native activities form the heart of a richer dataset. Each activity is broken down into one or more problem steps (Antonenko et al., 2012; Psaromiligkos et al., 2011). For instance, if a question asks a student to set the value of three dropdown boxes, then that question has three steps. In addition to the traditional timestamps and UI elements the student interacts with, each step is assigned a set of one or more hypothesized competencies or knowledge components (KCs) required by the student to answer the question (Stamper & Koedinger, 2011; Koedinger, Corbett, & Perfetti, 2012). This KC tagging of the questions, in conjunction with their accuracy, time on task, and number of attempts, provides detailed insights into which concepts students struggle with the most. In particular, the KC mapping provides a comprehensive modeling of the student learning process and enables both students and instructors to better assess their learning. Moreover, when used in conjunction with the additional semantic context provided by the OLI course structure, this data can be used to more meaningfully understand demonstrated learner misconceptions, evaluate course design elements, and provide information for primary and secondary learning science research and analysis.

3.2 Integrating Custom and Third-Party Activities

As a result of the increasing specialization of learning technologies, most current learning platforms depend on external learning tools, the consequence of individual companies and organizations tackling a unique type of student interaction or learning domain and its resulting technology. Additionally, the vast diversity of available interactive learning content makes it impossible for any single platform to support it all natively. The OLI platform is designed to build weak links and strong bonds to externally provided learning tools. The platform does not place many software constraints on the technologies it integrates with, but through usage of APIs, it creates a strong bond to its student performance analysis system (Dashboard, Logging, DataShop, etc.).

One benefit of the OLI platform's integration mechanism is that it easily allows the inclusion of research prototypes. The platform currently supports approximately 40 custom integrations, the majority of which are research-oriented. Some of these research projects are small add-ons which, for example, log specific user interactions (such as page interaction behaviors), but many of the

Creative Commons License, Attribution - NonCommercial-NoDerivs 3.0 Unported (CC BY-NC-ND 3.0)

integrations are full-scale research platforms in and of themselves (e.g. VLab; Aleven et al, 2016; Blink et al, 2014).

4 OLI DATA

Broadly, OLI data is classified into three categories of learning interaction analysis, each with its own client, log service, and processing components.

- **Student page interaction data** is captured as a log stream that records students' basic interactions with learning content. Questions such as, "How are students navigating the course materials?" and, "How much time is spent on learning activities?" can be answered from this data.

- **Student learning in activities** is captured such that feedback can be provided, student responses can be graded, and skill data can be updated. Within OLI, log data can capture additional learning behaviors, and can capture arbitrary additional data elements (specific to individual activity types); this approach provides rich source of information from which many different views of a learner's performance can be extracted.

- **Student problem navigation data** is produced by learning activities that require students to engage in more expanded and involved interactions with problem materials. For example, a math problem might require multiple simplification steps. For such multi-step problems, OLI logs data in DataShop/Tutor messaging format, which is a transaction-based format that can capture the precise way in which a student arrives at an answer.

4.1 Emerging Data Trends

Learning analytics and algorithms continue to provide a deeper view into student learning. In order to support new tools and techniques, OLI provides an extendable and interoperable method of logging data. See the figure 2, below, which outlines how the OLI architecture supports semantic data analysis.

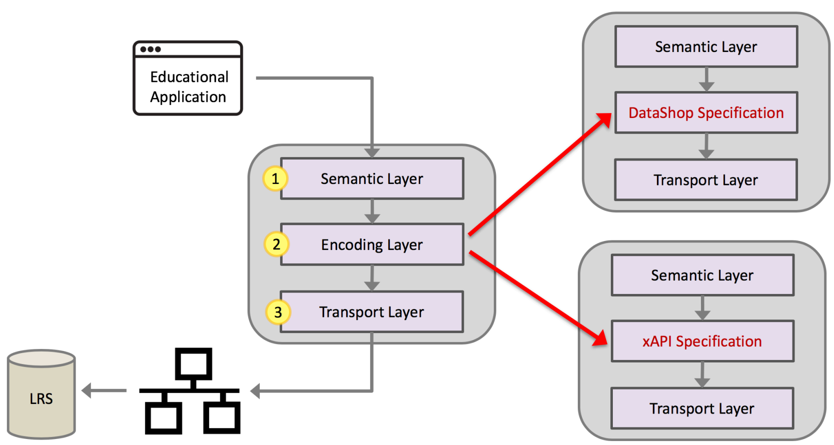

Figure 2 Data Logging Architecture

Learning content software developers have access to three layers of data encoding (note that in Figure 2, the term *Educational Application* denotes integrated learning activities as well as content pages and native OLI activities).

1) At the highest level, developers are concerned with ensuring that the meanings of student interactions are preserved. For example, it is important that analysis can discern what the state of a learning activity is when a student asks for a hint. Knowing the state defines what

Creative Commons License, Attribution - NonCommercial-NoDerivs 3.0 Unported (CC BY-NC-ND 3.0)

feedback is given, and paired with student request and tutor response, allows analytics to understand where and why a student is struggling.

2) Each message needs to be encoded such that the receiving end can determine the intent of both the student and responses by an automated feedback system. A number of different formats are available, each with slightly different goals and specifications. Since the OLI platform possesses no *a priori* knowledge of which analysis system will be used and which data format it uses, it ensures that an abstraction within its logging code can switch the data encoder. It is our intention to make these tools public and accessible to the greater technology-enhanced learning community, and we are therefore in the process of making these tools open source[2].

3) The message delivery from web browser to log service needs to be robust and efficient. This layer of the architecture supports message bundling to ensure that the browser has as few connections to the log service as possible, and the transport layer also supports retries and a local queue in the event of a log service becoming unresponsive due to networking issues.

4.2 Data Formats

A core aspect of the OLI approach to learning is our data-driven student model; the OLI platform captures exhaustive, real-time data on student interactions with learning materials and instructor interactions, and with learning materials and analytics tools. This data is used to drive feedback loops for learners and instructors (often in real time), as well as for Learning Engineers (for iterative course improvement) and Learning Scientists (for ongoing research and evaluation).

This exhaustive approach to data capture means that, in theory, any researcher, designer, or engineer can assemble the necessary data components for their current tasks or inquiries; in practice, however, the components are captured at a grain size fine enough that significant amounts of aggregation and pre-analysis are necessary to provide information in a useful form. To that end, OLI has a number of standard reports that capture the most frequently used approaches to our data.

4.2.1 Course Design and Improvement

OLI offers a number of reports that provide insight into the performance of learning materials and highlight potential areas for improvement. These reports range from raw numeric data to more carefully processed and designed spreadsheets which have been refined over multiple iterations. These more heavily processed reports exist to make the data visualization easier and more accessible to course authors and learning engineers, with color coding used to provide a first pass at interpretation and identification of potential problem spots in the course. Data for design and improvement includes:

- Number of students
- Average number of attempts
- Average help need

[2] https://github.com/Simon-Initiative/DataShopLogger

Creative Commons License, Attribution - NonCommercial-NoDerivs 3.0 Unported (CC BY-NC-ND 3.0)

- Eventually correct
- First try correct
- Utilization, completion, and accuracy rates
- Chart of low- and high-stakes performance per skill
- Chart of low- and high-stakes performance per learning objective
- Aggregate skill view showing potentially problematic skills where:
    - Assessments are missing
    - Practice is inadequate
    - Assessment and practice may be misaligned
- Aggregate learning objective view showing potentially problematic objectives where:
    - Assessments are missing
    - Practice is inadequate
    - Assessment and practice may be misaligned

Current plans for the improvement of course design and improvement analytics include embedding information from these reports directly into the course authoring platform.

4.2.2 Research and Evaluation

DataShop: OLI course data can be loaded into LearnLab's DataShop[3], providing an extensive range of analytic and reporting tools. DataShop spans the gap between research and improvement, with capabilities and methods that can be used for research and evidence-based course improvement. Tools include:

- Knowledge Component Modeling
- Learning Curve Analysis
- Problem Breakdown
- Performance Profiler
- Error Report

See: https://pslcdatashop.web.cmu.edu/Project?id=122

Evaluation Dataset: This data set is used when conducting more formal evaluation studies; it's an exceptionally large set normally accessed via a database, though export to CSV is possible. It contains aggregated information that can be used to analyze and answer questions including:

- To what extent did students access OLI content?
- To what extent did students complete the high-stakes assessments?
- To what extent did students' use of the course go beyond simply accessing/completing activities and assessments in a way that could have led to gains in their learning?
- What were student success, assistance, and help-seeking behaviors for low-stakes activities?
- How well did students perform on their initial attempt at any high-stakes assessment?
- How well did students perform on their last attempt at any high-stakes assessment?
- What were faculty access and use patterns for tools, analytics, and content?

[3] http://pslcdatashop.web.cmu.edu

Creative Commons License, Attribution - NonCommercial-NoDerivs 3.0 Unported (CC BY-NC-ND 3.0)

5 BUILDING COURSES WITH OLI

While the traditional OLI design process has proven successful in creating online learning experiences that demonstrably enact learning and support instruction (Thille & Smith, 2011), the process for developing OLI courses has continued to be time- and resource-intensive. Traditionally, this course design and implementation process hinged upon the role of the OLI Learning Engineer, whose task was to work closely with faculty, domain experts, and a larger course development team (potentially including learning scientists, instructional designers, assessment specialists, technologists, and other experts, as appropriate) to collaboratively design the learning experience, and then implement that design in OLI's XML structure. In addition to the XML authoring requirements, the system's build, deploy, and publishing process required additional expertise with subversion control systems and Linux-based command line tools. Furthermore, the deployment process itself often created extended intervals between design and implementation and finally publishing the completed, rendered courseware. Beyond the challenge of finding sufficient numbers of learning engineers possessing the requisite talents in learning design, project management, and technology, the process also made ongoing editing and revision challenging, and created a barrier for many educators who were interested in participating more directly in the authoring and improvement process. These hurdles have limited OLI in its ability to fully engage in the reuse/revision/remix approaches that are such an essential part of Open Education; developing and expanding the number of participants who use OLI as a community-based research activity is a core part of the Initiative's mission (Thille, 2012), and these barriers to authoring courses have slowed participation in the project by potential authors, hindering this part of the mission.

To address these challenges, OLI has invested heavily in developing an accessible, WYSIWYG authoring platform. This set of tools provides a better architecture to scaffold the design and development process, walking course developers and faculty through the process of articulating student-centered, measurable learning objectives and sub-skills; developing aligned practice and assessment opportunities with targeted hints and feedback; authoring expository learning elements; tagging course elements with the knowledge components represented by learning objectives and skills; and organizing these elements into a structured, coherent learning experience. The authoring tool is publicly available at http://echo.oli.cmu.edu. Beyond engaging with a larger community of authors and educators, the tool should also serve the learning analytics community by streamlining the process of developing well-instrumented learning experiences. Current development efforts are focused on expanding the design and improvement analytics that are embedded in the authoring tool, and on developing more thoughtful scaffolding for the authoring and developing process.

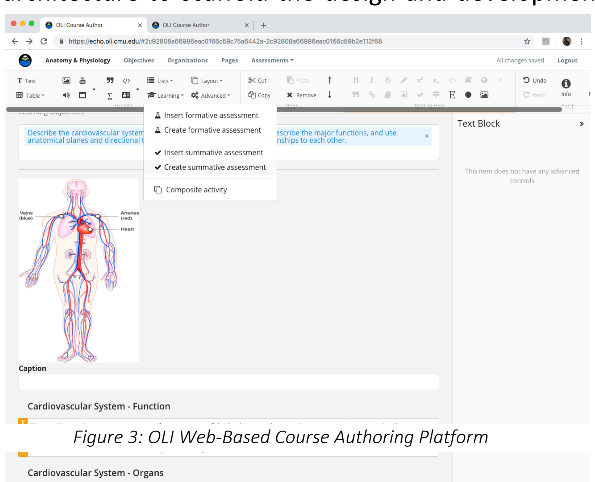

Figure 3: OLI Web-Based Course Authoring Platform

Creative Commons License, Attribution - NonCommercial-NoDerivs 3.0 Unported (CC BY-NC-ND 3.0)

REFERENCES

Aleven, V., Sewall, J., Popescu, O., Ringenberg, M., van Velsen, M. and Demi, S., Embedding Intelligent Tutoring Systems in MOOCs and e-Learning Platforms, In Proceedings of the Thirteenth International Conference on Intelligent Tutoring Systems (ITS) (2016)

Bier, N. and Jerome, W. (2012). Learning Data Visualization. Open Education Annual Conference. Vancouver, Canada, October 17, 2012.

Bowen, W. G., Chingos, M. M., Lack, K. A. and Nygren, T. I. (2014), Interactive Learning Online at Public Universities: Evidence from a Six-Campus Randomized Trial. J. Pol. Anal. Manage., 33: 94111. doi: 10.1002/pam.21728

Blink M.J., Stamper J.C., Carmichael T. (2014) SCALE: Student Centered Adaptive Learning Engine. In: Trausan-Matu S., Boyer K.E., Crosby M., Panourgia K. (eds) Intelligent Tutoring Systems. ITS 2014. Lecture Notes in Computer Science, vol 8474. Springer, Cham

Griffiths, R., Chingos, M., Mulhern, C., & Spies, R.. (2014). Interactive Online Learning on Campus: Testing MOOCs and Other Platforms in Hybrid Formats in the University System of Maryland. New York: ITHAKA S+R. Available at: http://sr.ithaka.org/research-publications/Interactive-Online-Learning-on-Campus

Herckis, Lauren and Smith, Joel. (2018). Understanding and Overcoming Institutional Roadblocks to the Adoption and Use of Technology-Enhanced Learning Resources in Higher Education. Report to the Carnegie Corporation of New York. https://www.cmu.edu/simon/news/docs/ccny-report.pdf (Accessed 11 December, 2018)

Kaufman, J., Ryan, R., Thille, C., & Bier, N. (2013). Open Learning Initiative Courses in Community Colleges: Evidence on Use and Effectiveness. Carnegie Mellon University, Pittsburgh, PA. Available online: http://www.hewlett.org/sites/default/files/CCOLI_Report_Final_1.pdf

Kernohan, David and Thomas, Amber. (2018). *OER: A Historical Perspective* (Jisc, 2012): http://repository.jisc.ac.uk/4915/ (accessed 27 August, 2018)

Koedinger, K. R., Corbett, A. T., & Perfetti, C.(2012). The knowledge-learning-instruction framework: Bridging the science-practice chasm to enhance robust student learning. *Cognitive Science*, 36, 757–798.

Koedinger, K., Kim, J., Zhuxin, J., McLaughlin, E., and Bier, N. (2015). Learning is Not a Spectator Sport: Doing is Better than Watching for Learning from a MOOC, *Proceedings of the ACM Conference on Learning @ Scale.*

Koedinger, K. R., McLaughlin, E. A., Jia, J. Z., & Bier, N. L. (2016). Is the Doer Effect a Causal Relationship? How Can We Tell and Why It's Important. In *Proceedings of the Sixth International Conference on Learning, Analytics and Knowledge,* pp.388-397. Edinburgh, UK.

Creative Commons License, Attribution - NonCommercial-NoDerivs 3.0 Unported (CC BY-NC-ND 3.0)

Leinhardt, G., Cuadros, J., & Yaron, D. (2007). "One Firm Spot": The Role of Homework as Lever in Acquiring Conceptual and Performance Competence in College Chemistry. *Journal of Chemical Education*, *84*(6), 1047.

Lovett, M., Meyer, O., & Thille, C. (2008). The Open Learning Initiative: Measuring the effectiveness of the OLI statistics course in accelerating student learning. Journal of Interactive Media in Education.

Lovett, M. (2012, July). Cognitively informed analytics to improve teaching and learning. Presentation at EDUCAUSE Sprint, July 25, 2012: http://www.educause.edu/ir/library/powerpoint/ESPNT122Lovett.ppt

Ritchie, Ann Lyon. (2018). Educators Return to LearnLab Summer School. https://www.cmu.edu/simon/news/stories/2018-learnlab.html (accessed 27 August, 2018)

Ryan, S., Kaufman, J.H., Greenhouse, J.B., She, R., & Shi, J. (2016). The effectiveness of blended online learning courses at the community college level. Community College Journal of Research and Practice, 40 (1) pp. 285-289

Shestak, Elizabeth. (2017). Simon Summer Institute Helps Pittsburgh-area Professors Build Ed Tech Tools for Their Classrooms. https://www.cmu.edu/simon/news/stories/2017-pche-summer-institute.html (accessed 27 August, 2018)

Stamper, J., Koedinger, K.R. (2011). Human-machine Student Model Discovery and Improvement Using DataShop. In Kay, J., Bull, S. and Biswas, G. eds. Proceeding of the 15th International Conference on Artificial Intelligence in Education (AIED2011). pp. 353-360. Berlin Germany:Springer.

Thille, C. & Smith, J. (2011) Cold Rolled Steel and Knowledge: What Can Higher Education Learn About Productivity?, Change: The Magazine of Higher Learning, 43:2, 21-27, DOI: 10.1080/00091383.2011.556988

Thille, C. (2012). Changing the Production Function of Higher Education. Part of the Series: Making Productivity Real: Essential Readings for Campus Leaders. Washington, DC: American Council on Education.

Wilson, M., & De Boeck, P. (2004). Descriptive and explanatory item response models. In *Explanatory item response models* (pp. 43-74). Springer New York.

Creative Commons License, Attribution - NonCommercial-NoDerivs 3.0 Unported (CC BY-NC-ND 3.0)