

# An Approach to Knowledge Component/Skill Modeling in Online Courses

by Norman Bier<sup>1</sup>, Sean Lip<sup>2</sup>, Ross Strader<sup>3</sup>, Candace Thille<sup>3</sup>, and Dawn Zimmaro<sup>3</sup>,

30 April 2014

<sup>1</sup> Open Learning Initiative, Carnegie Mellon University, <http://oli.cmu.edu>

<sup>2</sup> Google Inc., <http://www.google.com>

<sup>3</sup> Open Learning Initiative, Stanford University, <http://oli.stanford.edu>

## I. INTRODUCTION

The Open Learning Initiative (OLI) is an open educational resources initiative that began in 2002 with a grant from The William and Flora Hewlett Foundation<sup>4</sup>. Since its inception, over 100,000 students have participated in an OLI course at an academic institution and over one million students have engaged in one of the free and open versions of an OLI course. Experts in learning science, human computer interaction, instructional design, assessment, and software engineering work with disciplinary experts to create open online learning environments based on the integration of technology and the science of learning with teaching. These course materials are designed to *enact* instruction online in support of an individual learner. The OLI design process is founded on these principles:

- Apply learning science research and the scientific method to course development, implementation, and evaluation.
- Develop interactive learning environments collaboratively.
- Create feedback loops for continuous improvement.
- Engage communities of use for evaluation and improvement.

Improving learner productivity depends on designing learning environments based on principles derived from learning science research into how people learn. One of the main principles is that goal-directed practice and targeted feedback are critical to learning. Goal-directed practice involves working toward a specific level of performance and continually monitoring performance relative to clearly defined goals. When these goals are explicitly communicated to students, they guide and support students' purposeful practice and help them monitor their progress. In addition, students' practice is more effective when instructors (a) provide feedback that explicitly relates students' performance to the criteria, (b) ensure that the feedback is timely, frequent, and constructive, and (c) provide opportunities for these students to incorporate that feedback into further practice. Many learning studies have shown that students' learning improves and their understanding deepens when they are given timely and targeted feedback on their work

---

<sup>4</sup> The Open Learning Initiative funders include The William and Flora Hewlett Foundation, Bill & Melinda Gates Foundation, Lumina Foundation, Spencer Foundation, The Kresge Foundation, the National Science Foundation, The Walter S. Johnson Foundation, Carnegie Mellon University, and Stanford University.

(Wiggins, 1998; Butler & Winne, 1995; Corbett & Anderson, 2001; National Research Council, 2000).

OLI learning environments are built on decades of learning science research using methods proven to address instructional challenges and create conditions that enable robust learning and effective instruction. There is strong evidence that the OLI design approach supports and enhances student learning. John Hattie (2009), in a meta-analysis of over 800 studies on factors that influence achievement, reports the effect sizes for various factors that are the foundation of OLI course designs:

- Formative evaluation (0.90)
- Effective feedback (0.73)
- Meta-cognition (0.69)
- Mastery based learning (0.58)
- Interactive content (0.52)

Additionally, specific studies have shown the effectiveness of OLI courses, including:

- OLI statistics course: OLI students completed the course in half the time with half the number of in-person course meetings. OLI students showed significantly greater learning gains on the national standard “CAOS” test for statistics knowledge and similar exam scores. There was no significant difference between OLI and traditional students in the amount of time spent studying statistics outside of class. There was no significant difference between OLI and traditional students in follow-up measures given 1+ semesters later (Lovett, Meyer, & Thille, 2008).
- Trial of OLI Statistics Course at Public Institutions: This trial compared a hybrid version of the OLI statistics course with a traditional face-to-face statistics course with randomly assigned students at six institutions. Students in the hybrid format had comparable or better learning gains and took 25% less time to learn the same outcomes (Bowen, Chingos, Lack, & Nygren, 2012).
- Students using an OLI course in the fully online mode at a large public university with a high proportion of English-as-a-second-language-students achieved the same learning outcomes as students in traditional classes, and many more successfully completed the course (Schunn & Patchan, 2009).
- OLI stoichiometry course: The number of interactions with the virtual lab outweighed all other factors including gender and SAT score as the predictor of positive learning outcomes (Evans, Leinhardt, & Yaron, 2008).
- Engineering Statics: OLI instruction was comparable to traditional instructor-led coverage (Dollar & Steif, 2008).
- Community College OLI Study (Psychology, A&P, Biology, Statistics): Faculty use of and experience with the OLI course was associated with higher student achievement gains and may help smooth out expected negative outcomes associated with race (Kaufman, Ryan, Thille, & Bier, 2013).

- An OLI course at Carnegie Mellon University, Computing@Carnegie Mellon, has seen enrollments of 6,400 students since 2010. Data from the course is used as part of an annual course improvement and redesign cycle. The iterative course improvements have been relatively successful, demonstrating an increase in successful student completion and the use of support mechanisms by at-risk students (Lakhavani & Blair, 2011; Lakhavani & Rupp, 2013).

In order to create online learning environments that provide goal-directed practice based on learning science research, OLI course design begins by identifying and refining student-centered, measureable learning objectives. These learning objectives specify what students will be able to do at the end of a section of instructional content, typically defined as a module; developing these learning outcomes is a key part of the course development team's process and is usually accomplished by close collaboration between learning engineers and domain experts. A skills map is then created by the domain expert and/or learning engineer to identify the skills that compose the objectives. A skill identifies the discrete concept, or knowledge component, embedded within the learning objective. Each learning objective comprises one or more skills—skills are identified at a level of granularity such that they will be actionable for instructors (or course designers). Interactive activities and quizzes are created to assess students' learning related to the various skills identified in the skills map. OLI's learning model links to the skills map to predict student mastery of each skill and related learning objectives. In this paper, we will provide an overview of the skills mapping process, theoretical learning model approach used to predict mastery, and analytics reports currently embedded in the designs for OLI courses, and suggest recommendations for future improvements.

## II. TERMINOLOGY

### Course Structure


Courses created on the OLI platform are made up of one or more **units**, each consisting of one or more **modules**. Each module is a sequence of **sections**, and each section is made up of **pages**, each of which consists of a sequence of text, videos and/or **activities**. The activities are the focus of this paper.

An activity (sometimes called a resource) is made up of one or more **questions**, and these questions may each consist of multiple **parts**. The following screenshot shows an example of an activity with a single question that is made up of three parts. It is also possible for a question to have no parts; in that case, the question ends up being a 'leaf node' in the content tree.

## learn by doing

You need to determine the location of the center of gravity of the body shown. Which coordinate can be determined by inspection, using the symmetry?

Hint



- $x_G$   yes  no  
 $y_G$   yes  no  
 $z_G$   yes  no

### Skill Mapping

**Learning objectives** identify what a student should be able to do by the end of a module. Each learning objective comprises one or more knowledge components (KCs)—more commonly described as skills in the OLI interface. These **skills** (or sub-objectives) break down the learning objective into more specific cognitive processes, such as ‘can calculate the median of a given set of numbers’. Typical courses comprise about 30 -- 50 learning objectives and 100 – 1,000 skills. Note that individual skills can be associated with more than one objective. Each activity, question, or part of a question, can be associated with one or more skills.

### III. THE DESIGN PROCESS

Using the learning objectives for a course, domain experts and cognitive scientists specify the skills that comprise these learning objectives and are used in the activities of the course. This process builds a map between activities and objectives through these component skills:



Together this specification of Objectives, Skills and Activities constitutes a Knowledge Component Model or Skills Map. The specific process for determining the preliminary skills map has varied between OLI courses, with many older courses being retrofitted with preliminary models, and newer courses having skill maps created in conjunction with the course development process. This preliminary specification can be guided by a variety of approaches,

ranging from a more intuitive process that leverages instructors, learning engineers' and cognitive scientists' expertise to more rigorous approaches that resemble more formal cognitive task analysis (CTA). In almost all cases, objectives are associated directly with activities, questions or parts in a one-to-one way; as a sufficient number of questions are developed or reviewed, these objectives are broken out in more granular skills as necessary. The appropriate level of granularity continues to be a matter of research (and much debate) within the OLI community. Absent more formal evidence, OLI design teams tend to be guided by pragmatism:

- Objectives are decomposed as little as possible during first development of the model; reductionism for its own sake is to be avoided.
- Objectives should be decomposed only to the point of being actionable. The final purpose of developing a KC/Skills model is not to perfectly model the student learning process, but instead is to provide guidance to instructors and course improvement teams. Skills are specified in way that assists these instructors and developers in identifying areas of student struggle and then working to remediate.
- Generally, objectives are written to speak to a student audience; skills are written to speak to instructors.

The specification consists of three tabs of a spreadsheet, which are described below.

### Learning Objectives

The Learning Objective tab consists of a column for specifying learning objective ID's (matched to the same objectives that exist in the published course) and a series of skills columns (skill1, skill2....) listing the individual skills that comprise this objective. These skill IDs must match the skills listed in the Skills Tab.

<b>Learning Objective</b>	<b>Skill1</b>	<b>Skill2</b>	<b>Skill3</b>
apply_concepts	confintinterp	pvalueinterp	samplesizevprecis
apply_probability_rules	probdef	rule5	rule3
apply_rules_means_variances	randvarmean	randvarsd	
apply_sampling_distribution_of_sample_mean	samplesizerep	CLT	
apply_sampling_distribution_of_sample_proportion	samplesizerep		
apply_standard_deviation_rule	sdrule		

### Skills

The Skills tab consists of a key that is used to refer to the given skill in the other spreadsheets along with a human-readable full description of the skill.

<b>Skill</b>	<b>Title</b>
applyrandvar	Application of random variables
binparams	Binomial parameters
catvsquant	Categorical vs quantitative
CLT	Central Limit Theorem
compare2	Comparing 2 boxplots
computechisq	Compute chi-square statistic
computecondprob	Compute conditional probability

## Problems

The Problems tab details three columns that identify specific learning activities and skills columns (skill1, skill2....) listing the individual skills that are associated with this activity, part or question:

- Resource: the name of the activity
- Problem: an identifier for a question in this activity (if applicable)
- Step: the part of the question (if applicable)
- Skill1, Skill 2, etc.: A list of skills that this question or part is associated with.

Resource	Problem	Step	Skill1	Skill2
u4_m3_matchedpairs6_tutor1	q1		onevstvosided	
u4_m3_twosamples2_tutor1	q1	p1	specifyhyp	onevstvosided
u4_m3_twosamples2_tutor1	q1	p2	specifyhyp	
u4_m3_twosamples2_tutor1	q1	p3	specifyhyp	onevstvosided
u4_m3_twosamples2_tutor1	q1	p4	specifyhyp	
u4_m3_twosamples3_tutor1	q1		interpptest	
u4_m3_twosamples7_lbd_tutor1	q1		hyptest-logic	

Note that skills are almost always mapped to questions, but they can also be mapped to individual parts of a question or to an entire activity (resource).

Preliminary testing of this model can be done using scripts that identify common areas (typos, unspecified skills, etc.), however these scripts identify only those errors that would prevent a skill model from being refined or loaded into the OLI system—these scripts do not test the model for validity. Because the development of these spreadsheets is a complicated process that relies heavily on humans, it does tend to be error prone; to that end, these scripts have proven invaluable as part of the development process.<sup>5</sup> This testing has proven especially important given the current process for deploying the skills model into the OLI system for use in the prediction engine. During this process, the OLI system administrator takes these spreadsheets, converts them to CSV files, and runs a shell script that transfers the contents of these files into a database in the OLI system. This process is very fragile: errors are displayed one at a time, and after one error is displayed the process must be re-run.

## **IV. VALIDATING AND REFINING THE MODEL**


Although the initial development of this skills model relies heavily on human intuition, when the learning activities have been used by students, the data gathered from these interactions are used to validate and refine the model. After an OLI course has been offered, the dataset from the course is loaded into the PSLC's DataShop.<sup>5</sup> Once the dataset is available in DataShop, the current skills spreadsheet can be loaded (either manually or using the scripts<sup>5</sup> mentioned above).

---

<sup>5</sup> pslcdatashop.web.cmu.edu

DataShop offers a variety of tools and approaches for considering, testing, and refining Knowledge Component models. Of these, the most commonly used is Learning Curve analysis, which can help to quickly identify areas where skills may need to be decomposed, merged or otherwise may need to be changed based on the differences between expected and actual student performance. Much more detail on this process can be found at the DataShop website: [https://pslcdatashop.web.cmu.edu/ResearchGoals#rg\\_11](https://pslcdatashop.web.cmu.edu/ResearchGoals#rg_11).

Subsequent offerings of the course provide more data, leading to an ongoing cycle of model validation and refinement:



Data from the 'open and free' version of the course is not included in the calculation. This is because there is more variability among the students taking the 'open and free' version, and a larger percentage of these students do not engage meaningfully with the course.


The model is also validated by a post hoc comparison of the predictions produced by the skill mapping with actual student performance on high-stakes assessments that assess the skill, such as an offline final exam or a standardized assessment like the Advanced Placement (AP) test. The specific assessment used depends on the course. This provides a systematic way to test the skill mapping against the attempts of students whose knowledge state with respect to the specific skill is known.

## V. THEORETICAL MODEL

The creation of skills models serves a number of purposes, including assisting in the iterative course improvement process; measuring, validating and improving the model of student learning that underlies each course; and offering information necessary to support learning scientists in making use of OLI datasets for continued research. One of the best known uses of these skills maps, however, is to support learning analytics for instructors. In this context, the objective of skill mapping in OLI is to determine the probability that a given student has learned a given skill. Individual skills are treated as mathematically independent variables. In addition, it is assumed that learning a skill is a one-way process: once a skill is learned, it never gets unlearned.

In the current version of the model used in OLI courses, student learning is modeled using a Bayesian hierarchical statistical model with the latent variables of interest, students' learning states, becoming more accurate as more data is accrued about performance on a given skill (Lovett, 2012). The specific details of the theoretical model are not discussed here; however, the key idea is that the model takes as inputs data about student performance on questions and activities related to a certain skill, and estimates the probability that the student has mastered that skill. To accomplish that, the data stream sent into the learning model for a student's first action taken on a problem for each skill embedded in that problem includes the following: Resource ID, Problem ID, Problem Step ID, Skill ID, Student Response (1 = correct, 0 = incorrect/hint). The data stream sent out of the learning model to the dashboard for each learning objective to which a given skill is aligned includes the following: Skill ID, Learning Objective ID, Student ID, Mastery level (Green = Mastered, Yellow = Moderately Mastered, Red = Not Mastered, Gray = Not enough data).

There are many statistical models that could be used for accomplishing this task. Indeed, an emerging area of research in this field is to compare performance of various models in predicting student learning.




Note that the model only considers the first attempt the student makes on each question, and ignores subsequent attempts -- so, if the student were to retry the second question, and get it correct, this information is **not** incorporated into the model. This is because many questions in OLI courses give formative feedback immediately when a student gives an incorrect answer. After answering such a question for the first time, students tend to click on the remaining choices as well and look at the feedback for those. Thus, only a student's first attempt is considered an honest one, since it is made without the benefit of feedback from the system. The model also assumes an honest attempt by the student to learn the material; this assumption indicates that even incorrect answers can assist in supporting the learning when appropriate feedback is provided.

## VI. DISPLAY OF ANALYTICS

### Instructor Dashboard<sup>6</sup>

It is infeasible to display mastery data for hundreds or thousands of skills to the instructor. Skills are therefore aggregated into learning objectives, and this allows the instructor to see, at a glance, where students are succeeding and where they are struggling (Bier et al., 2011). This enables instructors to spend their time with students in a way that better utilizes their expertise. Instead of spending valuable class time going over concepts that students were able to learn outside of class, they can instead use the time to address problem areas (Strader et al., 2012).

The following screenshot shows a view of the instructor dashboard:



6

The learning levels are color-coded as follows:


- **Green:** students who have mastered the learning objective (as predicted by the model). Mastery is calculated based on rolling up the mastery scores for the skills that make up the learning objective, and averaging these.
- **Yellow:** students who are struggling with the learning objective
- **Red:** students who have not mastered the learning objective
- **Grey:** students for whom the model does not have enough data to predict their mastery of the learning objective. This happens because the student has not answered enough problems. The model defaults to a minimum of two questions for the skill before it dares to make a prediction; this minimum is customizable.

---

<sup>6</sup> The Carnegie Mellon Learning Dashboard was originally developed in a project led by Marsha Lovett and funded by the Spencer Foundation. The OLI implementation of the CMU Learning Dashboard was created by Judy Brooks, Bill Jerome, Marsha Lovett, John Rinderle, Ross Strader, and Candace Thille.

When a learning objective is clicked on, two additional charts are displayed:

1. The 'estimated learning by students' section allows a finer-grained drill-down into the extent to which students are mastering this learning objective. The instructor can click on a color-coded category and email all students in that category; he/she can also specify one or more particular students to email.<sup>6</sup>



2. The 'class accuracy by sub-objective' section shows the number of students that attempted at least one question associated with the sub-objective (i.e., skill). For those students, the histogram displays the average percentage of correct first-attempt answers; this average is taken across all questions associated with the corresponding sub-objective.



Clicking on a particular sub-objective will display the associated activities and questions related to that sub-objective as well as student performance on individual questions:

Activities in this module for *Predicting distribution shape* 

ACTIVITY	STUDENTS	% CORRECT
<a href="#">Histogram (3 of 3): Extra Problems &gt; Did I Get This?</a> <a href="#">View in Course</a> 	6	57%
<a href="#">Histogram (3 of 3): Extra Problems &gt; Did I Get This?</a> <a href="#">View in Course</a> 	6	50%
<a href="#">Histogram (3 of 3) &gt; Did I Get This?</a> <a href="#">View in Course</a> 	8	38%
<a href="#">Checkpoint: Examining Distributions Checkpoint 1</a> <a href="#">View in Course</a> 	8	100%
<a href="#">Checkpoint: Examining Distributions Checkpoint 2</a> <a href="#">View in Course</a> 	8	63%

6

» Question 1 «

6 responses, 83% correct

**A**  
1


**B**  
5

**C**

**D**  
0

**E**  
0

Choose the best description of the data used to generate this histogram (note that the horizontal axis has no scale, so you will make your choice based solely upon the histogram's shape):



**A.**  SAT Math scores of 1,000 future engineers and scientists .

**B.**  Results of rolling a six-sided die 1,000 times .

**C.**  Cholesterol levels of 1,000 adults .

**D.**  Shoe sizes of 1,000 men and women .

**E.**  Prices of 1,000 California homes .

## Student Dashboard<sup>6</sup>

In addition, some preliminary work has been done in creating a limited learning dashboard to be made available to the student, though very little research has been done to indicate what data should be displayed to students and whether this feedback is effective to support learning. This dashboard is very basic: it displays the student's mastery of each learning objective, and also allows students to view which activities they have and haven't done.

## VIII. GENERAL FINDINGS

The course design process, where skills are identified and listed, has proved to be a valuable source of information for course authors. Although the model and statistical algorithms can be improved, they are already better than most for evaluating student learning. The OLI approach represents a new method to teaching for many course authors and for instructors using the course, but with sufficient onboarding they can see the value of this approach and understand where their role is in this process.

The model described, however, has a number of limitations. Here are some of them:

- It would be nice to formulate more formal guidelines for course teams on how to come up with the list of skills and the mappings, so that the process is less *ad hoc*.
- The present system is rather rigid, in several ways:
  - It does not allow for pluggable predictive models of mastery.
  - When the skills map is changed, all pre-existing student mastery data is deleted. (Note that this only happens between courses.) This obsolescence of data was a poor design decision because a skills map cannot be updated in the middle of a course.
  - There is no decay model for ‘unlearning’.
  - The current two-tiered model of “learning objectives + skills” is somewhat baked in, and it might be nice to have enough flexibility to allow for additional levels of granularity.
  - If a student gets a question wrong, this affects all skills associated with that question. However, particular wrong answers may correspond to lack of mastery of only one of those skills.
- Although the process of creating and mapping skills is useful for creating a course, the current system of managing a bunch of spreadsheets is painful, error-prone, and inaccessible to many people.
- The visual display of information is attractive but could benefit from additional design attention now that it has seen extensive use.
- Only the first attempt at a question is used to determine mastery. Since there are a finite number of questions, further student attempts at questions for a particular skill do not change the estimate of mastery after they have exhausted the pool of questions. So, for example, once you have attained the green level, continued mistake-making after that may not reduce your learning estimate.

## IX. REFERENCES

Bier, N., Lovett, M., & Seacord, R. (2011). An online learning approach to information systems security education. *Proceedings of the 15th Colloquium for Information Systems Security Education*. Fairborn, OH: CISSE.

Bowen, W.G., Chingos, M.M., Lack, K.L., & Nygren, T.I. (2012). *Interactive learning online at public universities: Evidence from randomized trials*. ITHAKA.

Butler, D. L., & Winne, P. H. (1995). Feedback and self-regulated learning: A theoretical synthesis. *Review of Educational Research*, 65(3), 245 -281.

Corbett, A.T. & Anderson, J.R. (2001). Locus of feedback control in computer-based tutoring: Impact on learning rate, achievement and attitudes. *Proceedings of ACM CHI 2001 Conference on Human Factors in Computing Systems*, 245-252.

Dollár, A. & Steif, P. S. (2008). An interactive, cognitively informed, web-based statics course. *International Journal of Engineering Education*, 24, 1229-1241.

Evans, K. L., Leinhardt, G., & Yaron, D. (2008). Learning stoichiometry: A comparison of text and multimedia formats. *Chemistry Education: Research and Practice*, 9, 208-218.

Hattie, J. C. (2009). *Visible learning: A synthesis of over 800 meta-analyses relating to achievement*. New York: Routledge.

Kaufman, J., Ryan, C., Thille, C., & Bier, N. (2013). *Open Learning Initiative courses in community colleges: Evidence on use and effectiveness*. Retrieved from: [http://www.hewlett.org/sites/default/files/CCOLI\\_Report\\_Final\\_1.pdf](http://www.hewlett.org/sites/default/files/CCOLI_Report_Final_1.pdf).

Lakhavani, F. M., & Blair, M. A. (2011). Using the Open Learning Initiative (OLI) to deliver digital citizenship instruction to university freshmen at Carnegie Mellon University. In T. B. M. Ebner (Ed.), *Proceedings of World Conference on Educational Multimedia, Hypermedia and Telecommunications*, 2744-2749.

Lakhavani, F. M., & Rupp, A. (2013). Reengineering an introductory computer education course for undergraduate students. In R. Lansiquot (Ed.), *Cases on Interdisciplinary Research Trends in Science, Technology, Engineering, and Mathematics: Studies on Urban Classrooms*, pp. 75-90. Hershey, PA: Information Science Reference.

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

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*.

National Research Council. (2000). *How people learn: Brain, mind, experience, and school: Expanded edition*. Washington, DC: The National Academies Press.

Rabiner, L.W. (1989). A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2), 257 - 286.

Schunn, C. D. and Patchan, M. (2009) An evaluation of accelerated learning in the CMU Open Learning Initiative course Logic & Proofs. *Technical report by Learning Research and Development Center*, University of Pittsburgh., Pittsburgh, PA.

Strader, R., & Thille, C. (2012). The Open Learning Initiative: Enacting Instruction Online. In D. Oblinger, *Game Changers: Education and Information Technologies*, 201-213. Washington, D.C.: EDUCAUSE.

Wiggins, G. (1998). *Educative assessment: Designing assessments to inform and improve student performance*. San Francisco, CA: Jossey-Bass.