

# Integrating Diverse Learning Tools using the PrairieLearn Platform

**Matthew West**  
University of Illinois  
Urbana, IL, USA  
mwest@illinois.edu

**Nathan Walters**  
San Francisco, CA, USA  
nwalter2@illinois.edu

**Mariana Silva**  
University of Illinois  
Urbana, IL, USA  
mfsilva@illinois.edu

**Timothy Bretl**  
University of Illinois  
Urbana, IL, USA  
tbretl@illinois.edu

**Craig Zilles**  
University of Illinois  
Urbana, IL, USA  
zilles@illinois.edu

## ABSTRACT

In this article, we describe PrairieLearn, a flexible open-source platform for asking questions to students that is in broad use for both homework and exams. We demonstrate PrairieLearn’s flexibility and ability to integrate existing code and questions into a single platform using three case studies: Parson’s problems, designing finite-state machines, and auto-grading “Explain in plain English” questions. We highlight aspects of PrairieLearn’s structure that enable this flexibility, in particular PrairieLearn’s ability to execute arbitrary code both during the generation of a question instance and during grading student answers.

## CCS Concepts

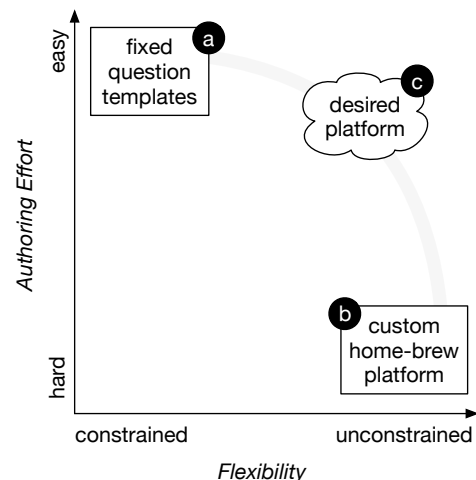
•Software and its engineering → Software organization and properties; •Social and professional topics → Student assessment;

## Author Keywords

assessment, software architecture, flexible, extensible

## INTRODUCTION

In the past two decades, we’ve seen significant growth in the number and usage of web-based platforms for asking questions related to course content. In addition to traditional Learning Management Systems (LMS, e.g., Canvas, Moodle), there have been publisher solutions that provide content associated with textbooks (e.g., WileyPlus, Connect), start-ups focused on providing domain-specific commercial platforms (e.g., Codio, Vocareum), as well as a wide range of solutions that instructors have built for themselves. Because they can potentially be used for both homework and exams, we refer to these as *question-asking platforms* (QAPs).



**Figure 1.** The current question-asking platform (QAP) landscape largely presents two options to instructors: (a) QAPs that enable easy authoring of a relatively small number of question types, or (b) the freedom to implement questions however they want by writing a new platform from scratch. This paper focuses on a software architecture designed to (c) provide significant flexibility while retaining the ability to easily implement common question features.

In most STEM classes, a large amount of content is objectively gradeable, enabling QAPs to perform auto-grading. Auto-grading provides the pedagogical benefit of immediate feedback to students, while simultaneously reducing the workload of instructors. Using auto-grading—where it is warranted—enables instructors and other course staff to focus their efforts on higher value tasks, like tutoring and grading subjective content (e.g., coding style).

For all that QAPs offer, we find that the current QAP landscape presents a practical tension between the ease of authoring questions and the flexibility in the kinds of questions that can be asked and auto-graded. Figure 1 illustrates this trade-off. Most commercial platforms provide rather constrained authoring interfaces that make it easy to write just a few specific kinds of questions [13] (e.g., multiple-choice questions, which require only a prompt, a correct answer, and a set of distractors), represented by Figure 1(a). Some instructors and researchers

**a** Return a formatted HTML list

The function below takes a single input, containing list of strings. This function should return string containing an HTML unordered list with each of `in_list` strings as a list item. For example, if provided the list ['CS 257', 'CS 281', 'CS 399'], your function should return a string containing the following. The only whitespace that is important is those contained in the given strings (e.g., the space between `CS` and `105`).

```
<ul>
<li>CS 257</li>
<li>CS 281</li>
<li>CS 399</li>
</ul>
```

```
print.py
1 def make_html_unordered_list(in_list):
```

**b** Path Weight

Consider the following graph:

If we consider the graph to be a network flow, what is the capacity of path  $AB \rightarrow BE \rightarrow EF \rightarrow FD$ ?

Path Capacity:

**c** Arithmetic Machine Datapath

In this problem, you are given a MIPS instruction and the values that are stored in registers 1-7 (below). Please enter the values that are on each labeled bus of the datapath in hexadecimal in the highlighted boxes for the instruction: `xor $2, $3, $7`

Format hexadecimal numbers without any prefixes (correct: 1A, incorrect: 0x1A).

**d** Parson's Problem: Temperature Scan

An airport staff member is measuring the body temperature of every passenger to detect which (if any) are having a high fever. Please construct a function `detect_fever`, which takes one argument `temperatures`, a list of floats containing the passengers' temperatures (in °C) and returns a list of indexes of passengers who have a body temperature higher than 38.5°C. If no passenger's temperature exceeds 38.5°C, return an empty list. For example, the following function call

```
temperatures = [39, 37.5, 37.3, 37.9, 38.4, 38, 38.6, 37.5, 37.6, 37.8, 37.4, 40]
print(detect_fever(temperatures))
```

should have the output below

```
[0, 6, 11]
```

```
fever_indexes.insert(i)
else:
return []
fever_indexes = index(temperatures)
return fever_indexes
if i > 38.5:
```

```
def detect_fever(temperatures):
fever_indexes = []
for i, t in enumerate(temperatures):
if t > 38.5:
fever_indexes.append(i)
return fever_indexes
```

**e** Finite State Machine (Least Recently Used)

Design a finite state machine that determines the least recently used number in the set {0, 1, 2} and outputs it as the 2-bit unsigned number (`xy`) each cycle. Your FSM receives a 2-bit unsigned binary number (`ab`). Your start state should assume that numbers 0, 1, 2 were previously seen in that order, making 0 the least recently used number. If an illegal number that is not part of the set is received as an input, return to the same state.

**f** Code Reading Problem

Write a short, high-level English language description of the code in the highlighted region. Do not give a line-by-line description.

Assume that the variable `x` is a list of numbers (either `int` or `float`) and the variable `y` is a number. You can assume that the code compiles and runs without error.

```
def f(x, y):
z = 0
for val in x:
if val < y:
z += 1
return z
```

**Figure 2. Examples showing flexibility of PrairieLearn, described in more detail in the text: (a) Code writing using external autograder, (b) randomly parameterized graph, (c) label the datapath value for a randomly generated instruction, (d) Parson's problem, (e) an FSM drawing mini-CAD tool, (f) autograded Explain-in-plain-English question using NLP.**

unsatisfied with these commercial offerings build their own specialized tools, represented by Figure 1(b). While interoperability standards like LTI [16] and SCORM [31, 4] are useful for integrating a collection of specialized tools into a single interface, they don't decrease the effort of implementing specialized tools.

This paper describes PrairieLearn [39, 38], an open-source question asking platform that provides both flexibility to author novel question types (approaching what could be obtained from a custom platform) and ease to author common question types (approaching what could be obtained from a commercial QAP). PrairieLearn is actively being used by the instructors of over 100 courses across multiple universities in a broad range of subjects (e.g., CS, Mechanical Eng., Chemistry, Statistics, Nutrition) and typically grades over 100,000 student answers daily. In addition, PrairieLearn has been used in a number of research studies [1, 2, 3, 5, 7, 9, 10, 8, 6, 12, 14, 23, 24, 27, 30, 32, 33, 28, 34, 35] and was a critical to the development of the Computer-Based Testing Facility (CBTF) [40, 41, 42, 43].

In Section 2, we attempt to demonstrate PrairieLearn's flexibility by showing a range of example questions that it supports. In Section 3, we highlight the three principles of PrairieLearn's design to which we attribute this flexibility, specifically:

1. Use of standard web technologies (Section 3.1)
2. Encapsulation of common functionality as *Elements* (Section 3.2)
3. Support for external auto-graders (Section 3.3)

### DEMONSTRATION OF FLEXIBILITY

We include images and descriptions of six problems that exemplify the kind of flexibility that PrairieLearn provides. For this audience, we've chosen to focus on CS-related questions, but there are similar questions in other domains including 3D-manipulation of robots, drawing forces to complete free-body diagrams, and randomly parameterized truss structures. Importantly, none of these question types are directly supported by the tool; these questions have been implemented on top of the interfaces that the tool provides.

In addition to multiple-choice questions, one bread-and-butter type of question is a code writing question (Figure 2(a)). This kind of question builds on the text editing element (Section 3.2) and the external auto-grader (Section 3.3) to allow an instructor-defined set of tests and scoring rubric to be applied to the student code.

A QAP is much more useful if it supports the development of *item generators* [17], which generate random parameters in order to create one of a large collection of question instances. Item generators enable re-use of the same question each semester and mitigate cheating because each student receives a different version of the question. Figure 2(b) shows an example of such a generator that renders a random graph for the student and asks the student to determine a feature of the graph. Figure 2(c) shows another item generator that asks students to compute the values on datapath and control wires of a MIPS processor for a randomly generated instruction.

In the sub-sections that follow, we provide detailed discussion of three problems that demonstrate PrairieLearn's capabilities.

### Parson's Problems

Parson's problems (shown in Figure 2(d)) present students with a collection of tiles containing lines of code; students are asked to select and order tiles to complete a program to accomplish a particular task. Parson's problems were proposed to provide students an opportunity to practice problem solving that involved lower cognitive load and focused less on syntax [29].

The first Parson's problem implementation in PrairieLearn was ported from the MIT licensed `js-parsons` library [18, 15] in one course's repository. As PrairieLearn uses standard web technologies—see Section 3.1—the front end portion of code was largely untouched, with the tile randomization and grading restructured to fit PrairieLearn's interfaces.

The code was implemented as a PrairieLearn *element* (Section 3.2). As an element, the Parson's problem implementation is separated from the specific text associated with a given question. After a year's use in one course, the element was significantly re-written to provide functionality beyond the original code; it now support both order-based correctness checking and running the constructed code against unit tests using the external autograder (Section 3.3) and released to all PrairieLearn users as the `pl-order-block` element.

### Auto-grading Finite State Machines

Figure 2(e) shows a question that asks students to design a finite-state machine (FSM) with certain properties. Its implementation is similar to the Parson's problem described in Section 2.1 in that its user interface is largely derived from existing open source code. Specifically, the mini CAD tool for drawing FSMs is also derived from MIT licensed software [37]. The interface allows creating new states, associating outputs with those states (only Moore machines are supported), connecting states with edges, labelling those edges with predicates, moving states and edges, and naming states (to help students keep track of their design).

When saved or graded, a textual representation of the student design is generated. This textual representation is auto-graded via simulation by running a collection of test vectors against the model and comparing the output to a "golden" implementation. Students are provided feedback by providing an example test input that failed and the series of outputs on their implementation and the expected output for that input.

### Auto-grading 'Explain in plain English' questions

Figure 2(f) shows an implementation of an auto-grading Explain in plain English (EipE) question, where the student is shown a piece of code and asked to provide a natural language description of the code. The code *reading* skill, which EipE questions assess and provide practice for, is thought to be an important developmental skill in learning to program [19, 22, 20, 36, 26, 25]. Lister et al. state that, while their data doesn't support the idea of a strict hierarchy, "We found that students who cannot trace code usually cannot explain code, and also that students who tend to perform reasonably well at code writing tasks have also usually acquired the ability to both trace code and explain code." [21]

The implementation shown in PrairieLearn supports auto-grading the natural language descriptions of code that students provide using natural language processing (NLP) [11]. While PrairieLearn has no built-in support for NLP, it allows arbitrary libraries to be used by question code. In this case, we're using Python's natural language toolkit `nltk` to tokenize, `unicode` to force students answers to ASCII, and `pyspellchecker` to do spelling correction. Then `numpy` is used to compute a score from the resulting bag-of-words and bigrams using a pre-trained model [11]. Because PrairieLearn allows authoring autograders in Python, our current implementation is built inside the questions themselves; if we wanted to use other languages, we could have implemented the autograder using the external autograder support.

### FLEXIBILITY/EASE PRINCIPLES

#### Question Flexibility through an Authoring Interface using Standard Web Technologies

Many LMS's provide highly structured authoring environments, often using a web-based GUI to fill in the content for an existing question format. While such an authoring environment meets its goal of simplicity and prevents bugs by construction, it also highly constrains the kinds of questions that can be asked.

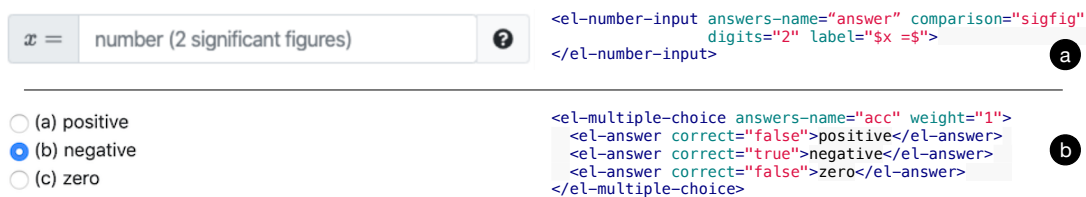
In contrast, PrairieLearn enables questions to be authored using standard web technologies: HTML5 (including client-side Javascript), server-side Python, JSON for question metadata, and Docker for external autograders (see Section 3.3). This minimally constraining authoring interface has three important implications:

1. It facilitates instructors having all content under one tool.
2. It permits using existing software libraries for question construction (e.g., the questions shown in Figure 2(d-f)).
3. All content is stored in the `git` version control system to allow flexible authoring as source code.

In practice, if you could build a stand-alone web site to ask your kind of question, you can implement your question using PrairieLearn. We've found this flexibility to be essential in our effort to move from paper exams to computer-based exams without dumbing down our courses (e.g., by only using multiple-choice questions). In fact, PrairieLearn facilitates novel question development because it provides authors with standard services (e.g., authentication, storing grades, logging) which allows authors to focus just on issues salient to the particular question.

#### Elements as Reusable Components

There are significant similarities among groups of questions. For some kinds of questions, the similarity is in the whole structure of the question (e.g., multiple-choice questions, Parson's problems). In other questions, it is smaller components (e.g., accepting a numeric value with a particular precision, displaying a graph). It would be ridiculous to provide an authoring environment which requires question authors to replicate the functionality for these common components and structures.



**Figure 3. Other example built-in elements: (a) a real number entry element that checks student’s answers using specified numerical precision algorithm and (b) a multiple-choice element that handles randomization of answer order and gives feedback. Note that the elements are processed so that the HTML rendered to students doesn’t indicate the correct answer.**

For this reason, PrairieLearn provides the abstraction of *elements*, as shown in Figure 3. Elements act like classes/objects in a programming language, in that they encapsulate common coherent bits of functionality (and presentation), which can be instantiated by questions. Importantly, elements are **not** question templates. Multiple elements can be composed to create a single question. In fact, this is very common, with questions requesting multiple numeric inputs, questions using elements to render figures and requesting numeric input (e.g., Figure 2(b)), or questions using an element to render code to the student and requesting a string input (e.g., the autograding EipE question in Figure 2(f)).

### External Autograders

While many questions can be auto-graded by elements or by an optional Python script that can be included in a question, some questions (e.g., Java programming questions) would be difficult to auto-grade in those ways. To ensure flexibility, PrairieLearn provides an *external autograder* system that allows arbitrary code to be executed in a sandboxed environment in order to grade a student submission. These questions typically involve either an in-browser IDE (e.g., Figure 2(a)) or a file-upload element for questions where students develop code/produce files locally on their machine.

There are two key ideas to how PrairieLearn implements a flexible external autograder interface; the first of which is containers. Containers are a bit like low-overhead virtual machines that enable question authors to specify exactly what software is available to the autograder (including which versions of the software). This configurability has enabled users to create code writing autograders for a wide variety of languages, including Java, C++, Python, R, OCaml, Haskell, MIPS assembly, and Verilog. PrairieLearn uses Docker to specify containers.

The second key idea is the notion of a clean interface. This interface is implemented through passing files in PrairieLearn. Question authors specify the set of files copied into the container; in most implementations, this consists of: 1) the submitted student work, 2) question-specific unit tests, 3) a question-agnostic (but language-specific) testing framework, including a script that manages the grading, and 4) the data dictionary generated by the question. PrairieLearn expects the container to produce a JSON file in a certain format that indicates whether grading was successful, any overall feedback, and a collection of rubric elements, each including the points earned and points available from that rubric item along with specific feedback on that rubric item (e.g., error from the failing test case).

The PrairieLearn server supports high throughput and low latency for this service by maintaining a pool of virtual processors, which PrairieLearn scales based on demand. At currently levels of usage it has sustained 20 problems/second with maximum time-to-grade below 5 seconds. Its overall median time to launch a grader (student submission click to grading code running inside a newly-launched container) is 1.2 seconds.

### CONCLUSION

As we shift more and more learning and assessment to digital systems there is a need for platforms that permit the integration of a broad range of questions and activities. Perhaps more importantly, we need platforms that reduce the effort of building questions and activities by managing all of the details unrelated to the particular questions being authored, without needlessly limiting the kinds of questions that can be written. We believe that PrairieLearn effectively walks this tightrope, by 1) allowing questions to be authored using standard web technologies, 2) encapsulating common functionality as elements that can be used to implement questions, and 3) by supporting external auto-graders so that arbitrary environments, languages, and libraries can be used to grade student work.

PrairieLearn is an open-source platform with an active user community that spans a broad range of universities. We would love to have you join our community.

### ACKNOWLEDGMENTS

Many people have helped to make PrairieLearn a success over the past decade. First, we would like to thank everyone who has contributed to its open-source development. Second, we would like to thank the instructors who use PrairieLearn in their own courses, particularly those who chose to try it in the early years and made suggestions for improvement. Third, we would like to thank four colleagues in particular who supported the growth of this project: Laura Hahn, Dave Mussulman, Carleen Sacris, and Tim Stelzer. Funding for this project was provided by the Strategic Instructional Innovations Program (SIIP) of the Grainger College of Engineering at the University of Illinois, and it was partially supported by the NSF under grants DUE-1347722, CMMI-1150490, and DUE-1915257.

### REFERENCES

- [1] Sushmita Azad, Binglin Chen, Maxwell Fowler, Matthew West, and Craig Zilles. 2020. Strategies for deploying unreliable AI graders in high-transparency high-stakes exams. In *International Conference on Artificial Intelligence in Education*. Springer, 16–28.

- [2] Jacob Bailey and Craig Zilles. 2019. uAssign: Scalable interactive activities for teaching the Unix terminal. In *Proceedings of the 50th ACM Technical Symposium on Computer Science Education*. 70–76.
- [3] Liia Butler, Geoffrey Challen, and Tao Xie. 2020. Data-Driven Investigation into Variants of Code Writing Questions. In *2020 IEEE 32nd Conference on Software Engineering Education and Training (CSEE&T)*. IEEE, 1–10.
- [4] Nicolas Casel, Marc El Alami, Damien Garot, and Denis Zampunieris. 2007. A new software architecture for learning managements systems with SCORM support. In *Proceedings of "IADIS-International Conference on e-Learning"*. Eds. M. Baptista Nunes & M. McPherson, 8–11.
- [5] Binglin Chen, Sushmita Azad, Max Fowler, Matthew West, and Craig Zilles. 2020. Learning to Cheat: Quantifying Changes in Score Advantage of Unproctored Assessments Over Time. In *Proceedings of the Seventh ACM Conference on Learning@ Scale*. 197–206.
- [6] Binglin Chen, Matthew West, and Craig Zilles. 2017. Do performance trends suggest wide-spread collaborative cheating on asynchronous exams?. In *Proceedings of the Fourth ACM Conference on Learning at Scale (L@S 2017)*.
- [7] Binglin Chen, Matthew West, and Craig Zilles. 2018. How much randomization is needed to deter collaborative cheating on asynchronous exams?. In *Learning at Scale*.
- [8] Binglin Chen, Matthew West, and Craig Zilles. 2019a. Analyzing the decline of student scores over time in self-scheduled asynchronous exams. *Journal of Engineering Education* 108, 4 (2019), 574–594.
- [9] Binglin Chen, Matthew West, and Craig Zilles. 2019b. Predicting the difficulty of automatic item generators on exams from their difficulty on homeworks. In *Proceedings of the Sixth (2019) ACM Conference on Learning@ Scale*. 1–4.
- [10] Binglin Chen, Craig Zilles, Matthew West, and Timothy Bretl. 2019c. Effect of discrete and continuous parameter variation on difficulty in automatic item generation. In *International Conference on Artificial Intelligence in Education*. Springer, 71–83.
- [11] Max Fowler, Binglin Chen, Sushmita Azad, Matthew West, and Craig Zilles. 2021. Autograding “Explain in Plain English” questions using NLP. In *The Proceedings of the 52nd SIGCSE Technical Symposium on Computer Science Education (SIGCSE)*.
- [12] Max Fowler and Craig Zilles. 2021. Superficial Code-guise: Investigating the Impact of Surface Feature Changes on Students’ Programming Question Scores. (2021).
- [13] M. Á. C. González, F. J. G. Peñalvo, M. J. C. Guerrero, and M. A. Forment. 2009. Adapting LMS Architecture to the SOA: An Architectural Approach. In *2009 Fourth International Conference on Internet and Web Applications and Services*. 322–327.
- [14] Geoffrey L Herman, Zhouxiang Cai, Timothy Bretl, Craig Zilles, and Matthew West. 2020. Comparison of Grade Replacement and Weighted Averages for Second-Chance Exams. In *Proceedings of the 2020 ACM Conference on International Computing Education Research*. 56–66.
- [15] Petri Ihantola and Ville Karavirta. 2011. Two-dimensional parson’s puzzles: The concept, tools, and first observations. *Journal of Information Technology Education* 10, 2 (2011), 119–132.
- [16] IMS GLOBAL. 2021. Basic Overview of how LTI Works. (2021). <https://www.imsglobal.org/basic-overview-how-lti-works>
- [17] S.H. Irvine and P.C. Kyllonen. 2002. *Item Generation for Test Development*. Lawrence Erlbaum Associates.
- [18] Ville Karavirta, Petri Ihantola, Juha Helminen, and Mike Hewner. 2018. js-parsons: a JavaScript library for Parsons Problems. (2018). <https://js-parsons.github.io/>
- [19] Raymond Lister, Elizabeth S Adams, Sue Fitzgerald, William Fone, John Hamer, Morten Lindholm, Robert McCartney, Jan Erik Moström, Kate Sanders, Otto Seppälä, Beth Simon, and Lynda Thomas. 2004. A multi-national study of reading and tracing skills in novice programmers. *ACM SIGCSE Bulletin* 36, 4 (2004), 119–150.
- [20] Raymond Lister, Colin Fidge, and Donna Teague. 2009a. Further Evidence of a Relationship Between Explaining, Tracing and Writing Skills in Introductory Programming. In *Proceedings of the 14th Annual ACM SIGCSE Conference on Innovation and Technology in Computer Science Education (ITiCSE '09)*. ACM, New York, NY, USA, 161–165. DOI: <http://dx.doi.org/10.1145/1562877.1562930>
- [21] Raymond Lister, Colin Fidge, and Donna Teague. 2009b. Further Evidence of a Relationship Between Explaining, Tracing and Writing Skills in Introductory Programming. In *Proceedings of the 14th Annual ACM SIGCSE Conference on Innovation and Technology in Computer Science Education (ITiCSE '09)*. ACM, New York, NY, USA, 161–165. DOI: <http://dx.doi.org/10.1145/1562877.1562930>
- [22] Mike Lopez, Jacqueline Whalley, Phil Robbins, and Raymond Lister. 2008. Relationships between reading, tracing and writing skills in introductory programming. In *International Workshop on Computing Education Research*. ACM, 101–112.
- [23] Suleman Mahmood, Mingjie Zhao, Omar Khan, and Geoffrey L Herman. 2020. Caches as an Example of Machine-gradable Exam Questions for Complex Engineering Systems. In *2020 IEEE Frontiers in Education Conference (FIE)*. IEEE, 1–9.

- [24] Jason W. Morphey, Mariana Silva, Geoffrey Herman, and Matthew West. 2020. Frequent mastery testing with second-chance exams leads to enhanced student learning in undergraduate engineering. *Applied Cognitive Psychology* 34, 1 (2020), 168–181.
- [25] Laurie Murphy, Sue Fitzgerald, Raymond Lister, and Renée McCauley. 2012a. Ability to 'Explain in Plain English' Linked to Proficiency in Computer-based Programming. In *Proceedings of the Ninth Annual International Conference on International Computing Education Research (ICER '12)*. ACM, New York, NY, USA, 111–118. DOI: <http://dx.doi.org/10.1145/2361276.2361299>
- [26] Laurie Murphy, Renée McCauley, and Sue Fitzgerald. 2012b. 'Explain in Plain English' Questions: Implications for Teaching. In *Proceedings of the 43rd ACM Technical Symposium on Computer Science Education (SIGCSE '12)*. ACM, New York, NY, USA, 385–390. DOI: <http://dx.doi.org/10.1145/2157136.2157249>
- [27] Terence Nip, Elsa L. Gunter, Geoffrey L. Herman, Jason W. Morphey, and Matthew West. 2018. Using a Computer-based Testing Facility to Improve Student Learning in a Programming Languages and Compilers Course. In *Proceedings of the 49th ACM Technical Symposium on Computer Science Education (SIGCSE '18)*. ACM, New York, NY, USA, 568–573. DOI: <http://dx.doi.org/10.1145/3159450.3159500>
- [28] Nicolas Nytko, Matthew West, and Mariana Silva. 2020. A simple and efficient markup tool to generate drawing-based online assessments. In *ASEE Annual Conference and Exposition, Conference Proceedings*, Vol. 2020.
- [29] Dale Parsons and Patricia Haden. 2006. Parson's programming puzzles: a fun and effective learning tool for first programming courses. In *Proceedings of the 8th Australasian Conference on Computing Education-Volume 52*. 157–163.
- [30] Seth Poulsen, Liia Butler, Abdussalam Alawini, and Geoffrey L Herman. 2020. Insights from Student Solutions to SQL Homework Problems. In *Proceedings of the 2020 ACM Conference on Innovation and Technology in Computer Science Education*. 404–410.
- [31] Rustici Software. 2021. SCORM solved and explained. <https://scorm.com/>. (2021).
- [32] Alan T Sherman, Geoffrey L Herman, Linda Oliva, Peter AH Peterson, Enis Golaszewski, Seth Poulsen, Travis Scheponik, and Akshita Gorti. 2020. Experiences and Lessons Learned Creating and Validating Concept Inventories for Cybersecurity. In *National Cyber Summit*. Springer, 3–34.
- [33] Mariana Silva. 2017. Algorithmic grading strategies for computerized drawing assessments. In *ASEE Annual Conference and Exposition, Conference Proceedings*, Vol. 2017.
- [34] Mariana Silva, Matthew West, and Craig Zilles. 2020. Measuring the Score Advantage on Asynchronous Exams in an Undergraduate CS Course. In *Proceedings of the 51st ACM Technical Symposium on Computer Science Education*. 873–879.
- [35] Paras Sud, Matthew West, and Craig Zilles. 2019. Reducing Difficulty Variance in Randomized Assessments. In *ASEE Annual Conference and Exposition, Conference Proceedings*.
- [36] Anne Venables, Grace Tan, and Raymond Lister. 2009. A Closer Look at Tracing, Explaining and Code Writing Skills in the Novice Programmer. In *Proceedings of the Fifth International Workshop on Computing Education Research*. ACM, 117–128.
- [37] Evan Wallace. 2010. Finite State Machine Designer (JavaScript). (2010). <http://madebyevan.com/fsm>
- [38] M. West. 2021. PrairieLearn. (2021). <https://github.com/PrairieLearn/PrairieLearn>
- [39] Matthew West, Geoffrey L. Herman, and Craig Zilles. 2015. PrairieLearn: Mastery-based Online Problem Solving with Adaptive Scoring and Recommendations Driven by Machine Learning. In *2015 ASEE Annual Conference & Exposition*. ASEE Conferences, Seattle, Washington.
- [40] C. Zilles, R. T. Deloatch, J. Bailey, B. B. Khattar, W. Fagen, C. Heeren, D Mussulman, and M. West. 2015. Computerized Testing: A Vision and Initial Experiences. In *American Society for Engineering Education (ASEE) Annual Conference*.
- [41] Craig Zilles, Matthew West, Geoffrey Herman, and Timothy Bretl. 2019. Every university should have a computer-based testing facility. In *Proceedings of the 11th International Conference on Computer Supported Education (CSEDU)*.
- [42] Craig Zilles, Matthew West, David Mussulman, and Timothy Bretl. 2018a. Making testing less trying: Lessons learned from operating a Computer-Based Testing Facility. In *2018 IEEE Frontiers in Education (FIE) Conference*. San Jose, California.
- [43] C Zilles, M West, D Mussulman, and C Sacris. 2018b. Student and instructor experiences with a computer-based testing facility. In *10th annual International Conference on Education and New Learning Technologies (EDULEARN)*.