This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Hosseini, Roya; Akhuseyinoglu, Kamil; Brusilovsky, Peter; Malmi, Lauri; Pollari-Malmi, Kerttu; Schunn, Christian; Sirkiä, Teemu

Improving Engagement in Program Construction Examples for Learning Python Programming

Published in:
INTERNATIONAL JOURNAL OF ARTIFICIAL INTELLIGENCE IN EDUCATION

DOI:
[10.1007/s40593-020-00197-0](https://doi.org/10.1007/s40593-020-00197-0)

Published: 01/06/2020

Document Version
Peer reviewed version

Please cite the original version:
Hosseini, R., Akhuseyinoglu, K., Brusilovsky, P., Malmi, L., Pollari-Malmi, K., Schunn, C., & Sirkiä, T. (2020). Improving Engagement in Program Construction Examples for Learning Python Programming. *INTERNATIONAL JOURNAL OF ARTIFICIAL INTELLIGENCE IN EDUCATION*, 30(2), 299-336.
<https://doi.org/10.1007/s40593-020-00197-0>

Improving Engagement in Program Construction Examples for Learning Python Programming

Roya Hosseini · Kamil Akhuseyinoglu ·
Peter Brusilovsky · Lauri Malmi · Kerttu
Pollari-Malmi · Christian Schunn · Teemu
Sirkiä

Received: date / Accepted: date

Abstract This research is focused on how to support students’ acquisition of program construction skills through worked examples. Although examples have been consistently proven to be valuable for student’s learning, the learning technology for computer science education lacks program construction examples with interactive elements that could engage students. The goal of this work is to investigate the value of the “engaging” features in programming examples. We introduce PCEX, an online tool developed to present program construction examples in an engaging fashion. We also present the results of a controlled study with a between-subject design that was conducted in a large introductory Python programming class to compare PCEX with non-interactive worked examples focused on program construction. The results of our study show the positive impact of interactive program construction examples on student’s engagement, problem-solving performance, and learning.

Keywords Introductory programming education, CS1, Python, Program construction, Worked examples, Classroom study

1 INTRODUCTION

Programming code examples play a crucial role in learning how to program. Instructors use examples extensively to demonstrate the semantics of the programming language being taught and to highlight fundamental coding patterns. Programming textbooks also place a heavy emphasis on examples, with a large proportion of textbook space being devoted to program examples and associated comments. Moreover, the code of all presented examples is typically provided on an accompanying CD or website to encourage students to explore, run, and modify the examples. Finally, recent studies have also shown that, on many occasions, students prefer seeing examples to receiving hints on how to build the code and fix their programs [60].

To increase the overall pedagogical value of code examples, textbooks and e-learning tools frequently augment the code with additional information that focuses on building

R. Hosseini, K. Akhuseyinoglu, P. Brusilovsky, C. Schunn
University of Pittsburgh; E-mail: {roh38, kaa108, peterb, schunn}@pitt.edu

L. Malmi · K. Pollari-Malmi · T. Sirkiä
Aalto University; E-mail: {lauri.malmi, kerttu.pollari-malmi, teemu.sirkia}@aalto.fi

either *program comprehension* or *program construction* skills. Program comprehension is related to reading, comprehending, and tracing the code, while program construction is related to writing or building the code. However, e-learning tools tend to support these two types of programming skills quite unevenly. While there is an active stream of work that has focused on augmenting code examples to support program comprehension (i.e., dynamic code animations and interactive program visualizations [68]), there are comparatively few e-learning tools focused on augmenting code examples with the goal of supporting program construction skills (i.e., a step-by-step demonstration on how to solve a specific programming problem [65]).

The most popular approach to present program construction examples is text-based *worked examples*. A typical worked example, as presented in the majority of printed and online textbooks, focuses on a meaningful programming problem to solve, presents code of a sample solution, and augments this code with a textual explanation of the code design and implementation process shown in steps with various granularity. While these explanations include all necessary information, it is presented in a static and non-engaging form. Some recent work has attempted to address the static textual nature of program construction examples by offering tools to construct interactive explanations [9] and codecasts [65]; yet the technology for presenting program construction examples is lagging behind the state-of-the-art in presented program comprehension examples that actively engage students in the work with examples [28, 46]. In contrast, the research on program-construction problems has recently explored a number of interactive approaches to engage students in program construction activities, such as the modification of code elements or adding missing parts in the code using, for instance, the Activecode elements [20, 41] or Parsons problems (puzzles) approach [29, 53]. Such activities can range from smaller-scale program completion to constructing fully working programs. Yet, unlike worked examples, program-construction problems offer no explanations and have focused more on knowledge assessment than knowledge acquisition.

In our work, we attempted to combine the explanatory power of traditional worked examples with the elements of interactivity and engagement offered by interactive examples and programming problems. We believe that our work bridges the gap between these two categories of learning activities and forms a continuum from static textbook examples into actual full-coding exercises. We introduce PCEX, an online learning tool that focuses on presenting program construction examples to students. PCEX offers examples in an interactive, engaging form in order to increase students' motivation to work with examples and improve their overall learning.

We formulated the following research questions in order to build a better understanding of the impact of PCEX on learning programming and to compare the impact of this interactive style to that of non-interactive examples. In this study, the non-interactive examples we presented were textbook-style worked examples that focused on program construction skills. By textbook-style worked examples, we mean the static worked examples typically presented in textbooks, which often lack interactivity and a challenge component. Specific hypotheses related to the research questions are presented in Section 4.1:

[RQ1.] Will working with PCEX examples engage students more than working with textbook-style worked examples? [**Hypothesis 1**]

[RQ2.] Will working with PCEX examples lead to better performance in solving program construction problems than working with textbook-style worked examples? [**Hypothesis 2**]

[RQ3.] Will working with PCEX examples lead to better learning outcomes than working with textbook-style worked examples? [**Hypothesis 3**]

In this article, we address these research questions by presenting results from a semester-long classroom evaluation that we conducted in a large introductory programming class to compare PCEX with textbook-style worked examples that focused on program construction. Our results demonstrated the benefits of PCEX examples over textbook-style worked examples on student’s engagement, performance, and overall learning outcomes.

The remainder of this article is organized as follows. After reviewing the related work in the next section, we describe our proposed tool, PCEX, in Section 3. Section 4 describes our study design, while Section 5 presents the results of the study. Section 6 concludes with a summary and discussion of this work and direction for future research.

2 BACKGROUND AND RELATED WORK

2.1 Worked examples in problem-solving

Over the last 30 years, worked examples, also referred to as worked-out examples [3], have gradually emerged as an important instructional approach that is supported by learning technology. A sizable body of research on instructional practices that support the use of worked examples for acquiring cognitive skills has been accumulated in various domains, such as mathematics and physics [3, 15, 69]. Worked examples are comprised of the presentation of a problem, the solution steps, and the final solution. Students use these examples as models of how to solve certain types of problems. Research on studying worked examples has consistently shown that in the early stages of skill acquisition, when students typically have little or no domain knowledge, instruction that relies more heavily on studying worked examples is more effective for learning than the traditional approach of being focused on only problem-solving. It has been shown that early example-based instruction leads to better learning outcomes, which are reached in less time and with less effort. This is usually referred to as the “worked example effect” [70].

The examples appear to be most valuable when presentation of the examples is combined with problem-solving. Specifically, past studies found that the pairing of worked examples with practice problems is more effective than providing learners with only practice problems [69] or only examples [71]. During the later stages of skill acquisition, however, the positive effect of worked examples gradually declines. In fact, example-based learning is inferior to problem-solving when learners have gained a reasonable degree of domain knowledge [30, 31, 32]. While learning from examples is superior to problem-solving for learners with little domain knowledge, this advantage disappears over time as the learners develop more content expertise. This phenomenon is referred to as the “expertise reversal effect” [32]. On one hand, this research points out the importance of examples for less prepared students, while on the other hand, it stresses the importance of carefully adapting to the current level of the student’s knowledge by decreasing the number of worked examples as the student gains expertise. A successful approach to implement this is using “faded worked examples” where the supportive guidance for students is gradually reduced as the student moves towards fully independent problem-solving [58]. A further improvement is *adaptive fading*, where the level of fading and/or the choice between a regular worked example, a faded example, or a problem is determined based on the students’ understanding of the examples [62] or their performance [45]. More recent research explored an adaptive alternation of worked examples, problem-solving, and *erroneous examples* [13, 14], a promising alternative to the faded examples [40].

According to [49], worked examples are hypothesized to be effective for several reasons. First, they provide constraints to the solution space; i.e., they highlight the correct solution

path so students do not need to waste time on incorrect or unfruitful searches. As a result, novices obtain the information needed to gain generalized knowledge more quickly [63]. Second, they reduce irrelevant cognitive load by highlighting the important elements of the problem and solution for the learner to focus on, encode, and reason about [51, 73]. Third, they encourage constructive cognitive processes such as self-explanation, in which the learner explains to herself the underlying conceptual logic and justification behind each step [12, 57]. This may be particularly important for helping students link the features of an example to abstract domain concepts or underlying principles.

It is worth noting that while the example-based learning approach often fosters learning for students with low levels of prior knowledge, this outcome is not guaranteed [2]. The positive impact of examples for the student’s learning can only be observed when (a) examples are designed effectively, and (b) students study the examples thoroughly.

2.2 Worked examples in programming

While worked examples are generally less explored in the domain of programming, there is a reasonable body of research that has guided our work. The earliest successful research showing the value of examples for learning programming dates back to the early 1980s, when Pirolli and Anderson [54] reported that examples are helpful for guiding students to solutions for novel and difficult problems. Since then, a considerable amount of research has been devoted to the development of example-based learning environments to support students in learning programming in various programming languages, such as LISP [25, 36, 74, 75, 76], Prolog [7], C/Java [9, 22, 39, 66], Javascript [19], SQL [13, 14, 45], and mini-languages [8, 78].

We classify program examples that have been used in teaching and learning to program into two groups, according to their primary instructional goal: *program behavior examples* and *program construction examples*. Program behavior examples are used to demonstrate the semantics (i.e., behavior) of various programming constructs (i.e., what is happening inside a program or an algorithm when it is executed). Program construction examples attempt to communicate important programming patterns and practices by demonstrating the construction of a program that achieves various meaningful purposes (e.g., summing an array). This distinction might not be clear-cut for code-only examples, since the same code could be used for both purposes. However, attempts to augment examples with learning technologies to increase their instructional value (i.e., adding code animation or explanations) usually focus on one of these goals.

Program behavior examples have been extensively studied. While textbooks and tutorials still explain program behavior by using textual comments attached to lines of program code, a more advanced method for this purpose — *program visualization*, which visually illustrates the runtime behavior of computer programs — is becoming increasingly popular. Over the past three decades, a number of specialized educational tools for observing and exploring program execution in a visual form have been built and assessed [68]. Despite their visual and dynamic nature, the majority of tools for presenting animated examples, especially the first generations, could be considered non-engaging, in that they limit the student’s role to passively watching the animation (e.g., [42, 61, 66]). Following several studies that have demonstrated the low effectiveness of “passively-watched” animated examples, several researchers have experimented with interactive animations that are explorable and challenging, for example, allowing the student to change input data [35], asking students to predict the result of a specific step [11], or asking strategic questions about the visualization [26, 47]. A meta-study or first-generation research on program visualization [28] helped to

stress the value and the importance of interactive engagement in the context of animated examples. With the help of this analysis, an international working group attempted to classify several known types of interactive engagement (such as *viewing*, *responding*, *changing*, or *constructing*) by its cognitive level or depth associating deeper engagements with higher levels of taxonomy of educational objectives [6]. Since that time, a considerable volume of empirical evidence confirmed that that more active and deeper engagement of students into working with visualized program examples improved their learning [11, 28, 35, 44, 46, 64] and had a positive influence on their problem-solving abilities in that domain [23].

Advanced technologies for presenting program construction examples are much less developed. In contrast to interactive and engaging worked examples in math and physics, for many years, the dominant approach for presenting worked code examples was simply text with comments [19, 37, 43]. More recently, this technology has been enhanced by adding audio narrations to explain the code [20] or by showing video fragments of code screencasts with the instructor’s narration being heard while watching slides or an editor window [65]. Both approaches, however, can still be considered as a passive presentation that does not allow for exploration and engagement. An attempt to add interactivity and exploration to worked program construction examples in the form of commented code was made in the WebEx system, which allows students to interactively explore line-by-line comments for program examples via a web-based interface [10]. More recently, [34, 52] attempted to make the screencast-based examples more engaging by allowing inline code editing or embedding programming exercises into the videos. However, with the notable exception of building faded examples for SQL [45], existing approaches for presenting program construction examples lack the engagement power of modern technologies for presenting program behavior examples. In this work, we attempt to use research findings in the area of program behavior examples to produce interactive program construction examples, called PCEX, that better engage students and improve their learning.

2.3 Block-based programming and Parsons problems

Block-based programming environments adopt a drag-and-drop interface that allows users to construct and execute computer programs by composing atomic blocks of code together. Common examples of block-based programming environments are Scratch [59]; Snap, a web-based environment based on Scratch [27]; and Alice. Scratch programs allow users to build the program in small chunks, execute any piece of program code individually, and see its effects immediately. Alice 2 [17] and its successor Alice 3 [18] apply an event-based object-oriented model. In Alice, users can program within a 3D environment, add objects to the scene, manipulate the objects’ attributes, and call the object’s methods.

Parsons problems are another related research to block-based programming. Parsons and Haden [53] originally created Parsons problems as an easy way for novices to solve programming assignments without having to type code or think about the exact syntax of the programming language. The idea is clever: a small number of code fragments are presented in a random order and the novice is asked to construct the described function or a small program by placing the fragments in the correct order. Each fragment may contain one or more lines of code and all of the fragments may not be required in the solution. In [29], the authors introduced a new family of Parsons problems inspired by the Python programming language. They proposed a two-dimensional variant of Parsons problems where lines of code are not only sorted, but also placed on a two-dimensional surface. The vertical dimension is used for ordering the code, as in traditional Parsons problems. The horizontal dimension

is used to define code blocks, based on indentation. In another attempt, Ericson et al. [21] proposed a 2D Parsons problem with paired distractors where each distractor was shown paired with the matching correct code block so that the learner only had to choose the distractor or the correct code. More recently, [24] implemented and explored a version of Parsons problem with internal re-shuffling for mobile phones.

3 ENGAGING PROGRAM CONSTRUCTION EXAMPLES

This section introduces PCEX, an online tool developed to present program construction examples in an engaging fashion. First, it defines exactly what the student’s engagement is, what are its different constructs, and which of the engagement constructs are used in the PCEX examples. Then, it presents the interface and discusses which interactivity elements are used in the PCEX examples to engage students.

3.1 Targeted aspects of engagement

Engagement is a construct that has many different definitions in education, ranging from activity completion to particular cognitive and affective forms of activity completion. Therefore, we need to define our conception of student’s engagement to ground our approach to creating examples that support engagement. We define engagement as the extent of a student’s active involvement in a learning activity [16]. It is often considered to be a multi-dimensional construct of possibly four distinct, yet inter-correlated and mutually supportive aspects: *behavioral engagement*, *emotional engagement*, *cognitive engagement*, and *agentic engagement* [16, 55, 56]. Each of these dimensions offers different ways to measure student’s engagement; although, as mentioned above, these measures are not fully independent.

Behavioral engagement refers to how effortfully the student is involved in the learning activity, in terms of attention, effort, and persistence [16, 67]. Higher behavioral engagement is associated with larger amount of work that the students are willing to perform, higher levels of completion and retention, and higher persistence in problem-solving. These aspects could usually be measured by processing students’ logs. Emotional engagement refers to the presence of positive emotions during task involvement, such as interest, and to the absence of negative emotions, such as anxiety [67]. Higher levels of emotional engagement could be recognized by higher levels of student’s satisfaction, bringing a more positive attitude to the learning process, and a willingness to recommend a novel learning tool to peers. These aspects could be measured through both surveys and questionnaires. Cognitive engagement refers to how strategically the student attempts to learn in terms of using sophisticated rather than superficial learning strategies, such as elaboration rather than memorization [72]. Cognitive engagement is defined to a considerable extent by the design of learning activities. In particular, the taxonomy of educational objectives [6] was designed to help instructors in creating educational activities with higher levels of cognitive engagement (such as activities based on responding, changing, or constructing in the case of animated program examples). Agentic engagement is a fourth and newly proposed aspect of student’s engagement that refers to the extent of the student’s constructive contribution to the flow of the instruction in terms of asking questions, expressing preferences, and letting the teacher know what the student wants and needs [55].

The work on engaging examples presented in this paper focuses predominantly on behavioral and cognitive engagement, which are the aspects that are most consistently linked with learning outcomes [5]. Our design focuses on cognitive engagement: we sought to design

examples that would more actively involve students in working through the full examples and encourage them to think more deeply. In accordance with past research on engagement [67], we expect that higher cognitive engagement will be associated with higher behavior engagement and better learning outcomes. To better assess the increased behavior engagement, we evaluate our technology in a *voluntary practice* context. In our study we separate *compulsory* learning content, which the students are required to work with to receive their course grade and *voluntary* learning content, which the students are encouraged, but not required to work with. The presence of this voluntary content allows us to more reliably assess the content-influenced behavior engagement (working with compulsory content is not a reliable measure since many students are working only to “get the points”). In our study, this voluntary content allows students to explore interactive examples, as well as work with additional challenges where they can apply their knowledge through related problem-solving tasks. In this context, behavioral engagement could be measured by counting how many tasks or steps they carry out with the voluntary content, as well as how much time they spend working with the content. The details of measures are explained in Section 4.6. To more reliably measure student’s behavior engagement with voluntary content, we performed our experiment as a semester-long study in a real (and large) class. This design choice limited our ability to measure overall levels of emotional and agentic engagement, since neither long surveys nor active student-teacher communication were feasible in this context. While we expect that our technology could also lead to higher emotional and agentic engagement, we will need different kinds of experiments to confirm these expectations.

3.2 Characteristics and design

PCEX (Program Construction EXamples) is an interactive tool to support mastering program construction skills through examples. The innovative idea behind PCEX is to create “rich examples” that support free exploration and challenge the student. Figure 1 illustrates a PCEX example. Each PCEX example includes a “goal” (Figure 1, A) and worked program steps (Figure 1, B). The goal states the function that the example program performs. The worked steps begin with a subgoal label (Figure 1, C) and are represented in the form of the sequence of short fragments of code (no more than a few lines of code) that illustrate how the program is constructed. Labeling subgoals in worked examples is known to increase student’s performance by leading students to group a set of steps and encouraging them to self-explain the reason for clustering those steps [12]. The example is enriched with instructional explanations that are indicated by question mark icons next to either all or a subset of example lines (Figure 1, D). Once a student clicks on a question mark, an explanation is shown on the right side (Figure 1, E). The student can request additional details for the selected line by clicking on the “Additional Details” button (Figure 1, G) or can navigate to the previous or next line to read an explanation (Figure 1, F).

In addition to being *explorable*, PCEX examples *challenge* students by engaging them with a problem-solving activity. When a student clicks on the “Challenge me” button (Figure 1, H), an interactive challenge activity is presented to the student, as shown in Figure 2. The goal of a challenge is to encourage students to apply the program construction knowledge presented in the original example to self-assess whether their understanding is correct. In essence, a challenge is a programming problem that is similar to the original example, in both the goal to achieve and the code. A challenge has both a problem statement (Figure 2, I) and code. However, the code has no explanation and is not complete — one or more of the code lines are missing. The student’s goal is to complete the code by dragging and

Example: Finding the Smallest Divisor of a Positive Number

A Construct a program that finds the smallest divisor (other than 1) of a positive number. For example, the smallest divisor of 4 is 2.

B

```

1 #Step 1: Assign initial values to the variables which we need for this program
2 num = 15
3 divisor = 2
4
5 #Step 2: Find the smallest divisor of the number
6 while num % divisor != 0:
7     divisor += 1
8 print("The smallest divisor of", num, "is", divisor)

```

C #Step 2: Find the smallest divisor of the number

D

E

We need to increment the divisor repeatedly as long as the divisor is not a factor of the number. Therefore, we need to use a loop structure. Since we don't know ahead of time how many times the loop will be repeated, we need to use a while loop. The condition in the while loop tests whether the body of the loop should be repeated, so it should test whether the divisor is not a factor of the number.

We could check whether the divisor is not a factor of the number by computing the remainder of the division of the number by the divisor.

F PREVIOUS NEXT

G PREVIOUS ADDITIONAL DETAILS

H Challenge Me!

Fig. 1 A Python programming worked example in the PCEX activity. The example includes the goal (A), interactive worked code (B), the subgoal label presented as a comment (C), the link to instructional explanations (question mark symbols) (D), explanations (E), a navigation link to the explanation for the previous/next line (F), additional details for the highlighted line (G), and a challenge navigation link (H).

dropping lines from the set of options (Figure 2, J) into each of the missing fields. This drag-and-drop interaction approach is similar to Parsons problems [53]. The student can decrease/increase the indentation of the line by using the indentation buttons (Figure 2, K/L), and the student can check whether the challenge is solved correctly by clicking on the “Check” button (Figure 2, M). The feedback is presented to the student by highlighting correctly (in green) and incorrectly (in red) placed lines. The student can also request a hint (Figure 2, N). If the student cannot solve the challenge within three attempts, he or she can request the solution.

At any moment, the student can navigate to the core explained example by pressing the “Back” button (Figure 2, O). Also, if an example has several challenges, the student may navigate between them by pressing the “Challenge Me” button (Figure 2, P). The student can navigate to the next challenge only when the current challenge is solved or the solution is seen after the third incorrect attempt.

3.3 Implementation details

In this work, we use the term *PCEX activity* to refer to the PCEX worked example and its associated challenges. PCEX activities are created by annotating the example and challenge code with a set of predefined tags (see Figure 3). These tags are used to define the goal or the problem statement, the instructional explanations, the lines to be faded, the set of line options, and the user input to the program. We used *docstrings* for Python or *doc comments* for Java to add these annotations to the program code.

The annotated code is parsed to generate a corresponding JSON file for each of the PCEX activities. The JSON file is then used by a single-page JavaScript application to support interactive work with the examples and challenges as shown in Figures 1 and 2.

Fig. 2 A Python programming challenge in the PCEX activity that follows the worked example shown in Figure 1. The challenge includes the goal (I), has one or more missing lines, and asks the student to drag and drop a line from the given options (J) to each missing line to construct the program. The student can decrease/increase the indentation of the line by using the indentation buttons (K)/(L). A student can request feedback by pressing the “Check” button (M), and the feedback message is shown in part (N). The student can go back to the example by pressing the “Back” button (O). If there are more challenges available, the student can go to the next challenge by pressing the “Challenge Me” button (P). This button is shown only when the current challenge is solved or after the student checks the solution after the third incorrect attempt.

The JavaScript application only needs the parsed JSON file to operate functionally, and it communicates back to the student modeling server [77] to log students’ interactions. The current version of the PCEX supports any executable code in both Java and Python.

PCEX employs an output-based evaluation of students’ solutions. As a result, the program code of a challenge is required to produce an output. The JSON file generated by the parsing process contains the correct program output (i.e. ideal solution) and the program outputs of all possible students’ solutions. This is achieved by first completing the code with all permutations of the line options (Figure 2, J) and then by running each completed code. This way, we can identify the cases where the student’s solution is correct, even if the solution consists of a different option order than the ideal solution. Moreover, this approach helps us to promptly provide feedback to the student, as opposed to waiting for the program execution to finish each time the student clicks on the “Check” button (Figure 2, M).

4 THE STUDY

To evaluate the impact on student’s engagement and learning from adding explorability and challenges to program construction examples, we conducted a classroom study followed by a survey to collect students’ feedback on the overall usefulness of the new program construction examples. Using a classroom study format was essential for our evaluation, since assessing engagement issues requires a realistic context. While the less controlled nature of a classroom study makes it harder to assess the overall impact of learning technology on learning gain, this type of study enables us to more reliably measure the engagement and motivation side

```

1  """
2  @goalDescription(Construct a program that finds the largest divisor of a positive number, excluding the number itself. For example, the largest divisor of 24 is 12.)
3  @name(Finding the Largest Divisor of a Positive Number)
4  @distractor(code( divisor = 1 ))
5  @distractor(code( divisor = num ))
6  @distractor(code( divisor += 1 ))
7  @distractor(code( divisor = num % divisor ))
8  """
9  #Step 1: Assign initial values to the variables which we need for this program
10 """@helpDescription(We define variable num to store the number that we want to find its largest divisor. We could initialize it to any positive integer greater than 1
11 num = 15
12 """@blank(We define variable divisor to store the largest divisor of the number. We initialize variable divisor to num-1 because we want to find the largest divisor of
13 divisor = num-1
14 """
15 #Step 2: Find the largest divisor of the number
16 """@helpDescription(We need to decrement the divisor repeatedly as long as the divisor is not a factor of the number. Therefore, we need to use a loop structure. Since
17 """@helpDescription(The body of the while loop is repeated as long as the divisor is not a factor of the number. The loop terminates when the loop condition evaluates
18 """@helpDescription(Also note that the % operator returns the remainder of the division.)"""
19 while num % divisor != 0 :
20     """@blank(When the divisor is not a factor of the number, we decrement the variable divisor by 1.)"""
21     """@blank(The -= operator subtract 1 from the value of variable divisor and the result is stored back into the variable divisor. Therefore, it is functionally equ
22     divisor -= 1
23     """@helpDescription(This statement prints to the default standard output stream the largest divisor of the number.<br>The printed text shows the arguments of the prin
24     print("The largest divisor of", num, "is", divisor)

```

Fig. 3 An annotated Python programming code for the PCEX challenge shown in Figure 2. The annotated code includes the problem statement (@goalDescription, Q), the title of the challenge (@name, R), the set of option lines (@distractor, S), the lines to be faded (@blank, T), and the instructional descriptions (@helpDescription, U). The annotation tags are inserted using Python *docstrings*.

of student's work. In this section, we present details of the study, beginning by explaining the design of the study, then presenting the information about participants and the study procedure. We conclude this section with information about the objective and subjective measures that were collected in the study and are later used in the data analysis.

4.1 Hypotheses

This study tests the following hypotheses regarding the benefits of PCEX compared to textbook-style, static, worked examples:

- H1.* Work with PCEX examples would engage students more than the work with textbook-style worked examples;
- H2.* Work with PCEX examples would lead to better performance in solving program construction problems than the work with textbook-style worked examples; and
- H3.* Work with PCEX examples would lead to better learning outcomes than the work with textbook-style worked examples.

4.2 Study design

To test our hypotheses, we designed a classroom study with two groups. In both groups, students could practice with voluntary examples and problems that were accessible through an online practice system. In order to make the groups equivalent in terms of the quantity and quality of learning material, both groups also received the same problems, program examples, and explanations. The difference between the groups was only in the way that examples were presented to the students. In the Experimental group, students could practice with PCEX activities that presented the sequence of similar examples with the first example fully worked (Figure 1) and the rest of the examples presented as tasks that challenged students (Figure 2). In the Control group, in contrast, students could practice with textbook-style example activities (see Section 4.3) that presented similar worked examples in the same order, all in the same form as shown in Figure 4.

4.3 Control condition interface: Textbook-style worked examples

The Control condition interface presents the examples in a traditional textbook-style form using a simple technology that shows examples statically. Figure 4 illustrates a worked example that was presented in this static interface. Each example includes a “goal” (Figure 4, A) and worked program steps (Figure 4, B). To make the presentation of the worked examples similar to those in programming textbooks, code segments and explanations are interleaved by default. The interface also enables the student to have a view of the code, similar to how the code is presented in an integrated development environment (IDE). The student can click on the “Hide Explanations” button (Figure 4, C) to switch to the mode that presents the code with no explanations. Figure 5 illustrates the code-only mode, with explanations hidden. The student can switch back to the mode where the explanations are shown by clicking on the “Show Explanations” link (Figure 5, E). The student can go to the next similar example by clicking on the “Next Example” link (Figure 4, D). The next example will be presented in the same style as the previous example. The student can navigate between similar examples by using the navigation links.

Fig. 4 The default mode of a worked example in the Control interface that presents static worked examples for program construction skills. The example is presented in a textbook-style form and includes the goal (A), worked code with subgoal labels and explanations (B), the link to hide explanations (C), and a navigation link to the next similar example (D). When appropriate, a previous example arrow will appear on the left side.

4.4 Study procedure

The study was carried out in an introductory Python programming course. This course was a CS1 service course targeted to students of bachelor's degree programs in various engineering fields at a large research university. The students were not Computer Science majors, and

Example: Finding the Smallest Divisor of a Positive Number

Construct a program that finds the smallest divisor (other than 1) of a positive number. For example, the smallest divisor of 4 is 2.

[Show Explanations](#) **E**

```

1 num = 15
2 divisor = 2
3 while num % divisor != 0 :
4     divisor += 1
5 print("The smallest divisor of", num, "is", divisor)

```

Fig. 5 The code-only mode of a worked example in the Control interface that presents worked examples focused on program construction skill. The student can click on link (E) to view the explanations.

most of them were likely to complete only one or two programming courses in their bachelor-level studies. The course consisted of lectures, nine exercise rounds, and an exam. To pass the course, the student had to solve enough problems (small Python coding exercises where the student mostly wrote the whole program alone) in each of the exercise Rounds 1–8 and pass the final test. To obtain a good grade, the student was also required to solve enough problems in Round 9, which contained coding exercises about object-oriented programming.

Potential participants included 723 undergraduate students who were enrolled in the course in the Fall semester of 2017. Students were randomly assigned to one of the Control and Experiment groups, explained in Section 4.2. They were given one week to take an online pre-test (see Section 4.5.2). In the second week of the class, the practice system with all practice content (see Section 4.5.1) was introduced and all students were provided with a link and an individual account to access the system. The use of the practice system was voluntary; however, the students were offered a few extra points to encourage them to explore the system and its practice content. More specifically, the students who explored at least eight examples and solved at least seven Parsons problems (see Figure 6) were given 200 extra exercise points. The maximum number of exercise points the students could receive in this course by solving mandatory exercises (assignments) was 5820. In addition, the students who answered to the course feedback survey organized by the university were given 200 exercise points. Thus, the total number of exercise points a student could gather was 6220, from which 200 (about 3 %) could be earned by using the practice system. Because the minimum number of the exercise points which the student needed for the best exercise grade was 5520, using the practice system was genuinely voluntary. The required minimum number of the examples and problems to receive those extra points was also much lower than the number of examples and problems available in the practice system (see Section 4.5.1). Doing more than the required minimum was not incentivized by the awarding of any additional points.

A student could practice with the system up until the end of the semester, and at that point was given one week to respond to an online example evaluation survey (described in

Section 4.7) and take the online post-test, which was isomorphic to the pre-test. Completion of the pre-test, post-test, and survey was voluntary, too. Therefore, to increase students' participation, the pre-test was included in the first assignment, and students were offered a total of 4 extra final exam points for answering the post-test and the survey (2 points for each). The maximum number of exam points a student could achieve without those extra points was 96.

4.5 Materials

4.5.1 Practice content

Practice content in the system included 52 example activities and 31 2D Parsons problems [29]. In this type of problem, the student was asked to construct the described program by putting the fragments in the correct order with the correct indentation (see Figure 6). Feedback was shown upon student's request, which highlighted the correct and incorrect lines in the code.

Fig. 6 An instance of a Parsons problem in the practice system. The student assembles the solution to the question (written at the bottom) by dragging lines of code to the right side.

Problem presentation was the same across the Control and Experimental groups, but examples were presented differently. In the Control group, examples were presented as textbook-style worked examples (as in Figure 4), while in the Experimental group, examples were presented as PCEX activities (as in Figures 1 and 2).

All practice content was organized into 15 topics that were ordered by increasing difficulty and presented through the Mastery Grids portal [38] that used visual personal progress tracking (known as Open Student Modeling (OSM)) to engage students to work with the content. Each example activity started with a worked example and was followed by one to three similar examples (in the Control group) or challenges (in the Experimental group), and the median was one (1). In total, the example activities included 52 worked examples, collectively containing 531 line explanations, and 71 similar examples/challenges, which were presented after the worked examples in the Control/Experimental group.

4.5.2 Pre- and post-tests

The pre-test consisted of 10 questions that evaluated a student's prior knowledge of a subset of programming constructs that were covered in the course. The questions were designed to cover both simple and complex concepts in Python programming. The first five questions were multiple-choice questions that asked students to select the correct code snippet for each given task. The remaining five questions asked students to determine the correct order of the provided lines of code, in order to achieve a certain purpose. The lines were shuffled and included distractors. A post-test was isomorphic to the pre-test, employing the same concepts but using different surface features. Cronbach's α between pre-test and post-test was .58 (indicating acceptable reliability). The maximum possible score on the pre/post-test was 10; one point was awarded for each question.

4.6 Metrics

We employed a variety of measures to compare the impact of the PCEX and textbook-style program construction examples on student's engagement and learning. Table 1 provides an overview of the metrics we used in our study. We grouped these into *engagement*, *performance*, and *learning* metrics.

4.6.1 Engagement metrics

Behavioral engagement metrics focused on 1) work done on the examples and problems, and 2) overall system usage. The specific measures include the number of attempts, as well as the time on task for both examples and problems.

Each example activity consisted of subtasks that were similar code examples presented in different formats, depending on the group. In the Experimental group, the first subtask was a worked example and the next-to-last subtask was a challenge. In the Control group, all subtasks were worked examples presented in a static, textbook-style form. To measure the amount of work done on examples, we looked at the student's work on the first and next-to-last subtasks individually. Therefore, the collected data included the following measures for representing the work done on examples: number of first subtasks viewed, total time on first subtasks, number of next-to-last subtasks viewed, total time on next-to-last subtasks, and interactions with explanations. In the Control group, interactions with explanations indicated the number of times the student switched between modes with and without explanations. In the Experimental group, on the other hand, the interactions with explanations indicated the number of times the student clicked on the question mark symbol next to a line to view the explanation.

The data also included several measures related to overall system usage such as total time spent on the practice system, total attempts on activities, and number of sessions that the student practiced with the system.

4.6.2 Performance metrics

We looked into coding performance within and outside of the practice system. The total correct attempts on Parsons problems and the number of distinct Parsons problems solved were measures of coding performance within the practice system. The measures of coding performance outside of the practice system included the total number of points that the

Table 1 Overview of the metrics used to evaluate the PCEX examples.

Aspect	Goal	Metric
Engagement	Work on examples	Number of first subtasks viewed Total time on first subtasks Number of next-to-last subtasks viewed Total time on next-to-last subtasks Interactions with explanations
	Work on problems	Total correct attempts on Parsons problems Distinct Parsons problems solved Total time on Parsons problems
Performance	Inside-system problem-solving performance	Total correct attempts on Parsons problems Distinct Parsons problems solved
	Outside-system problem-solving performance	Assignment points Early submission
Learning	Near transfer	Pre- and post-test
	Far transfer	Final exam <ul style="list-style-type: none"> • program comprehension questions • program construction questions (basic, complex)

student earned from coding assignments in the class and how early students submitted their coding assignments. It is not evident that earliness in assignment submission could be used as a standard measure of student’s performance, as it could be a sign of better motivation or preparation. However, the study in [4] has found some correlation between student’s performance and how early a student submits the assignment. As a result, we thought it would be interesting to see the impact of work on examples on early submission.

The assignment points were obtained from nine rounds of coding exercises during the course, producing a total with a minimum of 0 and maximum of 5,820 points. Each assignment could be completed during a certain time window during the term. Sometimes students spent a large amount of time completing the assignment and other times they were quicker (apparently better prepared from prior instruction). As a proxy measure for the degree of students’ preparation, enabling them to complete their assignment sooner, we calculated the inverse of the median number of days that assignment was submitted after the assignment had been introduced¹.

4.6.3 Learning metrics

We measured student’s learning by using the pre- and post-test (near transfer measure) and final exam score (far transfer measure). The final exam consisted of questions that assessed program comprehension and program construction skills. The program construction questions were further grouped into two basic questions and one complex question. The first basic question asked about basic concepts in the course. The second basic question asked the student to complete a relatively simple task: writing a simple class and a main program that used that class. The complex question required a deeper understanding of loops, opening files,

¹ We chose start date over due date as the time reference because some students submitted assignments after the due date was passed. Therefore, using the difference between the due date and the submission date would have made the interpretation of results more difficult, as some differences would be negative.

splitting strings, and exception handling, and was more difficult than the basic questions. We separated the grades for each group of exam questions. The maximum number of points that students could obtain was 21 on the program comprehension questions, 55 on the basic program construction questions, and 20 on the complex program question. The Cronbach's α between the three subcomponents in the exam was .67, which indicated an acceptable level of reliability.

4.7 Example evaluation survey

To collect students' feedback on the example activities, we administered a two-part survey related to system use and system impact. In the first part of the survey, students responded to questions about the amount of system use: Yes-more than 10 times; Yes-between 5 and 10 times; Yes-less than 5 times; and No. Those who chose one of the last two options were asked to provide their opinion on six follow-up items that focused on why the system was not used. Two of the items referred to a *bad system experience*, two emphasized *no help needed*, and two addressed *other reasons*, especially a poor introduction to the system and a lack of time to use the system. For this section of the survey and all questions in the second part of the survey, students were asked to respond using a 5-point Likert scale ranging from *Strongly Disagree* (1) to *Strongly Agree* (5).

The second part of the survey aimed to evaluate the impact of the example activities, focusing on only students who used the system. Following the suggestion in [33] that identified key constructs required to evaluate a learning objective, we included three constructs: learning, quality, and engagement. Each construct had four items, with two negatively worded and two positively worded. For the learning construct, the items referred to a student's perception of how much they learned from the examples. For the quality construct, the items referred to the quality of the example activities. Finally, for the engagement construct, the items examined the level of student's involvement in the example activities.

5 RESULTS

5.1 Students' participation and collected data

Out of the 723 students enrolled in the class, only 202 used the system (i.e., had at least one attempt on an example or problem), 696 took the pre-test, 457 took the post-test, 447 took both the pre-test and post-test, and 456 answered the questionnaire. Among the students who used the system, 6 students had an extremely high pre-test score (i.e., 90 percent or above). We discarded the data of those 6 students since they had little to learn and were likely only participating for extra credit. The final dataset included the data from 196 students: 118 students in the Control group and 78 students in the Experimental group. We used this data to perform our analysis of engagement and performance.

According to the learning gain data, we observed that some students had negative learning gains (minimum value was -0.5), which likely reflects that a few students did not take the post-test seriously. More precisely, among 196 students who used the system, 49 earned fewer points in the post-test than in the pre-test, of which 32 were in the Control group and 17 were in the Experimental group. For only the learning analysis section, we excluded the group of students who had negative learning gains². After discarding these students, we

² Using all students' data showed the same pattern of results for this analysis.

were left with data on 147 students (86 in the Control group, 61 in the Experimental) for the learning analysis. Note that there were no significant differences in the mean pre-test scores of both the Experimental ($M = 3.5, SD = 1.8$) and Control group ($M = 3.2, SD = 1.5$), $F(1, 116.08) = .77, p = 0.38$.

Prior to data analysis, we identified outliers in the collected data by using Tukey’s box-plot method, which defines outliers as being outside the interval $[Q1 - 1.5 \times IQR, Q3 + 1.5 \times IQR]$, where Q stands for “quartile” and IQR stands for “interquartile range”. We used winsorization to replace the outliers with less extreme values in the same direction (i.e., $Q1 - 1.5 \times IQR$ or $Q3 + 1.5 \times IQR$).

Finally, we examined the collected data to remove highly correlated measures. We observed that two of the measures related to the overall system usage, namely, total practice time and total attempts on activities, were highly correlated with each other ($\rho = 0.83$). Also, both measures were found to be highly correlated with the number of first subtasks viewed, and ρ was .80 for the total practice time and .99 for the total attempts on activities. As a result, we chose to use only the number of sessions as a measure of overall system usage. Similarly, the total correct attempts on Parsons problems was highly correlated with distinct Parsons problems solved ($\rho = 0.99$); therefore, we excluded the total number of Parsons problems solved from our analysis.

The following subsections 5.2, 5.3, and 5.4 present the results from the data analyses. Table 2 summarizes our main results by showing the hypotheses of classroom study (described in Section 4.1), the corresponding data analyses, and whether the hypotheses were confirmed by the data analyses. See the corresponding subsections for further details.

Table 2 Summary of hypotheses and results of the classroom study (✓: confirmed hypothesis; ✗: hypothesis that is not confirmed).

Hypotheses	Data Analyses	Measures	Hypotheses Confirmed? Group effect	Hypotheses Confirmed? Interaction effect
H1 – PCEX vs. textbook-style examples: engagement	Section 5.2	time-on-task interactions with explanations	✓ ✗	N/A N/A
H2 – PCEX vs. textbook-style examples: problem-solving	Section 5.3	dist. Parsons problems solved assignment points early submission	✗ ✓ ✓	✗ ✗ ✓
H3 – PCEX vs. textbook-style examples: learning	Section 5.4	post-test exam: code comprehension exam: basic code construction exam: complex code construction	✗ ✗ ✗ ✓	✓ ✗ ✗ ✗

5.2 Engagement analysis

Our first hypothesis ($H1$) was that PCEX would cause students to be more actively involved in PCEX worked examples than in the textbook-style worked examples. To test hypothesis $H1$, we used a one-way ANOVA analysis to compare the group means; for statistically significant cases, effect sizes using eta-squared (η^2) are presented in (Table 3). The metrics in Table 3 are ordered with a decreasing order of effect size. We used Cohen’s rules of thumb for interpreting this data, with an effect size of 0.02 being considered “small” in magnitude, 0.06 being “medium”, and 0.14 being “large”.

We can see that among all the measures, the ones that are related to the worked examples display the largest differences between the two groups. The largest effect comes from working

Table 3 Means (and SD) for engagement metrics in the Control and Experimental groups, along with inferential statistics and effect sizes, contrasting the groups.

	Control (N=118) Mean $\pm$ SD	Experimental (N=78) Mean $\pm$ SD	One-way ANOVA	Effect size (η^2)
WORK ON EXAMPLES				
Total time on next-to-last subtasks (mins.)	8.3 $\pm$ 14.9	36.7 $\pm$ 29.5	$F(1, 103.12) = 61.91, p < .001^{***}$	.29 <i>L</i>
Interactions with explanations	30.0 $\pm$ 23.9	14.9 $\pm$ 24.0	$F(1, 164.18) = 18.74, p < .001^{***}$	.09 <i>M</i>
Total time on first subtasks (mins.)	2.1 $\pm$ 1.9	3.1 $\pm$ 2.3	$F(1, 144.79) = 10.54, p = .001^{**}$	.06 <i>M</i>
#first subtasks viewed	23.6 $\pm$ 18.9	19.0 $\pm$ 16.9	$F(1, 177.19) = 3.24, p = .073$	.02 <i>S</i>
#next-to-last subtasks viewed	31.6 $\pm$ 27.2	24.5 $\pm$ 24.0	$F(1, 178.56) = 3.66, p = .057$	.02 <i>S</i>
OVERALL SYSTEM USAGE				
Number of sessions	3.8 $\pm$ 3.0	5.1 $\pm$ 4.1	$F(1, 194) = 6.63, p = .011^*$	.03 <i>S</i>
WORK ON PROBLEMS				
Distinct Parsons problems solved	14.5 $\pm$ 11.9	12.3 $\pm$ 11.5	$F(1, 168.82) = 1.56, p = .213$	
Total time on Parsons problems (mins.)	31.0 $\pm$ 29.4	24.9 $\pm$ 26.5	$F(1, 176.56) = 2.29, p = .132$	

*** $p < .001$; ** $p < .01$; * $p < .05$; . $p < .1$
 Effect size: *S*=small; *M*=medium; *L*=large

on the next-to-last subtasks. The mean of total time spent on the next-to-last subtasks was 4 times greater in the Experimental group, as compared to the Control group (36.7 vs. 8.3 mins.). This is expected because the next-to-last subtasks in the Experimental group required the student to solve the challenge, while the next-to-last subtasks in the Control group did not. The second largest effect was seen for working on the first subtasks. The total amount of time that students spent on the first subtasks was about 1.5 times greater in the Experimental group (3.1 minutes) than in the Control group (2.1 minutes).

We also observed that students had more interactions with explanations in the Control group than in the Experimental group. On average, the students in the Control group clicked on the show/hide explanations links 30 times, which was twice as high as in the Experimental group (14.9 times). We also found differences between the two groups in terms of the number of first subtasks and next-to-last subtasks that students viewed. Students in the Control group viewed, on average, more first subtasks and next-to-last subtasks; yet the differences reached only marginal significance and the effect size was small.

In addition to the differences between usage of activities, we also observed that the two groups were different in terms of number of practice sessions. The mean number of sessions differed by one across the two groups, with the Experimental group having more sessions ($M = 5.1, SD = 4.1$) than the Control group ($M = 3.8, SD = 3$). However, both the number of Parsons problems solved and the time spent on Parsons problems were not statistically different between the two groups.

In total, the ANOVA analyses of cognitive-behavioral engagement data partially supported hypothesis *H1*, favoring greater time-on-task but not more interactions with explanations. While students in the Control group viewed marginally more examples, they spent less time on the example subtasks. Students in the Experimental group, on the other hand, spent more time working with examples — about 1.5 more time on the first subtasks and 4 times more time on the next-to-last subtasks. The increase on time-on-task can be attributed to students becoming more involved when working with the PCEX examples than when working with the textbook-style worked examples.

Contrary to our hypothesis for greater interactions with the PCEX examples, we found that the Control group students used the “show/hide explanation” button more frequently than the Experimental group students clicked on example lines to view explanations. This difference, however, could be due to the way in which the explanations were presented in the textbook-style worked examples and PCEX examples. As mentioned earlier in Section 4.3, explanations were shown by default in the textbook-style examples, to make the presentation

of the worked examples similar to the programming textbooks. On the other hand, explanations were only available by taking action in the PCEX examples. Therefore, we hypothesize that the students viewing the textbook-style examples were influenced by needing a greater number of clicks to hide the explanations compared to viewing the code alone, similar to the way in which it was presented to the Experimental group. Our data supports this hypothesis by showing that, on average, the median of clicks on the “hide explanation” button was twice that of the median of clicks on the “show explanation” button in an example (0.2 vs. 0.1).

5.3 Performance analysis

Our second hypothesis (*H2*) was that working with PCEX would improve student’s performance on program construction tasks more than that of the textbook-style worked examples. To test hypothesis *H2*, we ran a series of regression analyses—to examine the effect of the group, the amount of work on examples, and the interaction between the group and the amount of work on examples—for the student’s coding performance, while controlling for prior learning, as indicated by the pre-test. That is, the independent variables were:

- *Group*: a dummy variable representing the group that student belonged to, with the Control group serving as the reference group factor;
- *WOE*: Work On Examples (*WOE*) is a continuous variable representing the combined work on the first subtasks and next-to-last subtasks in the example activities;
- *Pre-test*: a continuous variable representing the student’s pre-test score; and
- The *Group* $\times$ *WOE* interaction: *Group* $\times$ *WOE* is the interaction between *Group* and *WOE* variables. This interaction means that the effect of *Group* on the performance measure is different for different values of work on examples. For example, a positive value for the effect of the interaction term would imply that the more work is done on examples, the more positive becomes the effect of the Experimental group on student’s performance. All numeric independent variables (*Pre-test* and *WOE*) were mean-centered to reduce potential multicollinearity problems [1].

The dependent variables were measures of coding performance, including:

- Distinct Parsons problems that the student solved;
- Total points that the student earned in the coding assignments; and
- Earliness of the student’s submissions in the coding assignments.

The data that was used in the performance analysis was limited to the 194 students who used the system and who had taken the pre-test. Table 4 shows the results of the fitted models. The results revealed that *WOE* and *Pre-test* were positive predictors of distinct Parsons problems that student solved ($F(4, 189) = 217.2$, $p < .001$, $R^2 = .82$). Importantly, the effect of *WOE* on predicting distinct Parsons problems solved ($\beta = .88$) was about 14.7 times larger than the effect of the pre-test score ($\beta = .06$). We found no significant influence of the *Group* ($p = .371$) or the *Group* $\times$ *WOE* interaction ($p = .299$). This indicates that working on examples was associated with solving more Parsons problems correctly, regardless of the treatment group.

Group, *WOE*, and *Pre-test* were found to be predictors of assignments points, ($F(4, 189) = 5.54$, $p < .001$, $R^2 = .10$), since all were positively associated with assignment points. Among them, *Group* was the most influential ($\beta = .28$) and *WOE* was the second most influential predictor ($\beta = .26$). We found no significant effect for *Group* $\times$ *WOE* interaction ($p = .551$).

Table 4 Regression results of Group, amount of work on examples (WOE), and the interaction of amount of work on examples with group ($\text{Group} \times \text{WOE}$), predicting distinct Parsons problems solved, assignment points, and earliness of submission of coding assignments.

Predictors	Dist. Parsons problems solved				Assignments points				Early submission			
	<i>B</i>	<i>SE</i>	β	R^2	<i>B</i>	<i>SE</i>	β	R^2	<i>B</i>	<i>SE</i>	β	R^2
Group	.67	.74	.06	.82	290.91	146.78	.28*	.10	.13	.04	.39**	.22
WOE	.23	.01	.88***		6.07	2.00	.26**		.00	.00	−0.02	
Pre-test	.38	.21	.06		105.17	40.73	.18*		.05	.01	.29***	
Group $\times$ WOE	.02	.02	.07		−2.05	3.43	−0.09		.004	.001	.59***	

*** $p < .001$; ** $p < .01$; * $p < .05$; . $p < .1$

These results suggest that although work on examples was generally helpful for getting more points in the coding assignments, it was overall more helpful to be in the Experimental group and practice with engaging examples.

The results of the regression analyses also showed that although *WOE* was not a significant overall predictor of earliness of submission ($p = .82$), *Group* and *Group $\times$ WOE* interaction were a predictor of early submission (Figure 7). As the interaction plot shows, more activity with examples was associated with submitting assignments earlier in only the Experimental group. Furthermore, the interaction was found to be the most important predictor for earliness of submission ($\beta = .59$), with its effect being about two times greater than the effect of the pre-test score ($\beta = 0.29$).

Fig. 7 Interaction between work on examples (WOE) and Group factor (Control/Experimental) for predicting the earliness of submissions for coding assignments. Notches indicate 95% confidence interval ranges.

In total, the multiple regression analyses supported hypothesis *H2* for coding performance outside of the practice system, but not within the practice system. We found an overall positive effect in favor of the Experimental group on assignment points and earliness of submission. Being in the Experimental group and practicing with PCEX examples was associated with obtaining more points on coding assignments and also with submitting the assignments earlier. Furthermore, the interaction effect of the work with examples and the treatment group on earliness of submission was significant, which indicates that only more work on PCEX examples (and not textbook-style worked examples) led to earlier submission of assignments. Our analysis did not show any difference between the treatment groups in terms of performance within the practice system, though. More work with examples was associated with solving more Parsons problems, regardless of the treatment group; that

is, textbook-style worked examples and PCEX examples were both helpful for improving student's performance on Parsons problems.

5.4 Learning analysis

Our third hypothesis ($H3$) was that working with PCEX would improve learning outcomes more than working with the textbook-style worked examples. To test hypothesis $H3$, multiple regression analyses were performed to check the effect of the work on examples, the group, and the interaction between the work on examples and group on the post-test score and exam grade. The independent variables in all the regression models were similar to the independent variables in Section 5.3.

For predicting the post-test score, we used data of 194 students who used the system and who had taken the pre-test. For predicting the exam grade, this data was further limited to the 170 students who had taken the final exam. Table 5 shows the results of the regression model that tested the effect of *Group*, *WOE*, and *Group* $\times$ *WOE* interaction on the post-test score, controlling for the effect of the pre-test score. As expected, *Pre-test* was the most important predictor of students' post-test scores ($\beta = .42$). Neither *Group* ($p = .128$) nor *WOE* ($p = .677$) were accurate predictors of post-test scores, but *Group* $\times$ *WOE* was ($\beta = .29$); this finding suggests that the effect of the group depended on the amount of work on example (Figure 8). As the figure shows, work with examples in the Experimental group increased the student's post-test scores more than work with examples in the Control group.

Table 5 Regression results of predicted post-test score while controlling for the pre-test score.

Predictors	<i>B</i>	<i>SE</i>	β	R^2
Group	.32	.21	.22	.22
WOE	.0	.0	.04	
Pre-test	.35	.06	.42***	
Group $\times$ WOE	.01	.0	.29*	

*** $p < .001$; ** $p < .01$; . $p < .1$

Fig. 8 Interaction between work on examples (WOE) and Group factor (Control/Experimental) for predicting the post-test score. Notches indicate 95% confidence interval ranges.

The results of the fitted models for predicting the exam grade are shown in Table 6. As the table shows, *WOE* positively predicts exam grade across all categories of questions; namely, program comprehension ($F(4, 165) = 3.98, p = .004, R^2 = .09$), basic program construction ($F(4, 165) = 2.56, p = .04, R^2 = .06$), and complex program construction ($F(4, 165) = 3.07, p = .018, R^2 = .07$). Also, the *Group* factor is a significant predictor for the exam grade on the complex question. The interaction *Group* $\times$ *WOE* is not statistically significant for any exam category. For the complex program construction questions, *Group* ($\beta = .28$) and *WOE* ($\beta = .23$) are the first and second most important predictors, respectively, and the effect of each was about two times greater than the effect of the *Pre-test* ($\beta = .12$).

Table 6 Regression results of group, amount of work on examples (WOE), and the interaction of amount of work on examples with group predicting exam grade on the program comprehension and construction questions.

Predictors	Program Comprehension				Program Construction							
					Basic				Complex			
	<i>B</i>	<i>SE</i>	β	R^2	<i>B</i>	<i>SE</i>	β	R^2	<i>B</i>	<i>SE</i>	β	R^2
Group	-0.66	.75	-0.13	.09	1.03	1.68	.09	.06	1.22	.66	.28 .	.07
WOE	.02	.01	.20*		.06	.02	.23*		.02	.01	.23*	
Pre-test	.43	.21	.16*		.94	.47	.15*		.31	.18	.12 .	
Group $\times$ WOE	.01	.02	.10		-0.04	.04	-0.15		-0.01	.02	-0.12	

* $p < .05$; . $p < .1$

In total, the multiple regression analyses of the learning outcomes supports hypothesis *H3*, which demonstrates a significant positive interaction between work with examples and treatment group on post-test (near transfer test) and a marginally positive effect of the treatment group on the complex program construction question in the final exam (far transfer test). More specifically, more work with only the PCEX examples (and not the textbook-style worked examples) was associated with a higher post-test score. Additionally, being in the Experimental group and practicing with PCEX examples was marginally associated with obtaining a higher grade in the complex program construction question of the final exam. We found no differences between the PCEX and textbook-style worked examples on other questions in the final exam. Work with examples in both groups was found to be positively associated with obtaining a higher grade in the program comprehension questions and basic program construction questions of the final exam.

5.5 Survey analysis

Before analyzing the survey group differences, we assessed each construct's reliability using Cronbach's α . We dropped two items from the engagement construct and one item from the quality construct because their item-construct correlations were lower than the recommended value of .30. Additionally, we checked whether the internal consistency could improve if any of the items within a construct were deleted. We discarded one item in the quality construct because the removal of that item increased the internal consistency among the items of the construct and improved the α from .7 to .73. No item was discarded from the learning and engagement construct, as all items had acceptable internal consistency with the other items within that construct. The α was .74 for the learning construct and .7 for the engagement construct. After this step, all three constructs appeared to be sufficiently reliable for assessing

the value of the examples, with α values exceeding the suggested minimum acceptable α coefficient of .50 [50].

The survey was available for all students who enrolled in the course. As mentioned in Section 4.7, only students who used the system were able to respond to the second part of the survey that sought to evaluate the examples. Out of the 456 students who completed the survey, 15% ($N = 67$) used the system more than 10 times, 14% ($N = 62$) used the system between 5 and 10 times, 26% ($N = 120$) used the system less than 5 times, and 45% ($N = 207$) did not use the system at all. The mean and standard deviation for the survey items (except the items that were discarded during the reliability analysis) are shown in Table 7. Overall, students mostly agreed that they did not use the system due to preferring other resources and materials, not feeling the need for additional help, and a lack of time. Students disagreed with items that suggested other reasons for low/zero of the practice system, including the items that referred to having a bad system experience.

Table 7 Mean and standard deviation of the responses for the classroom study survey. The response options represented a 5-point Likert scale, ranging from 1 (Strongly Disagree) to 5 (Strongly Agree).

Item	Mean (SD)
<i>REASONS FOR LOW/ZERO USAGE</i>	
I preferred to use other resources and material to learn Python	4.0 (0.9)
I was doing well in class without the system and did not need any extra help	3.9 (1.0)
I did not have enough time to use the system	3.4 (1.1)
The system was not introduced properly in class	3.0 (1.0)
I didn't think the system can help me to better master Python	3.0 (0.9)
The user interface was too confusing to use	2.9 (0.8)
<i>EXAMPLE EVALUATION</i>	
<i>Quality</i>	
The explanations in the examples were easy to follow	3.6 (0.8)
The explanations in the examples were not hard to understand ^a	3.5 (0.9)
<i>Learning</i>	
The explanations in the examples helped me to better understand the Python programming concepts ^a	3.3 (0.8)
Working with the examples helped me learn Python	3.3 (0.8)
Exploring similar examples helped me learn Python	3.2 (0.7)
The examples helped me in solving Python exercises in this class ^a	3.0 (0.9)
<i>Engagement</i>	
I tried hard to understand the examples	2.9 (1.0)
I did not skim-read the examples ^a	2.6 (1.0)

^a A reverse-coded item.

Overall, about 60% of the students agreed with each of the items in the quality construct, while less than half of the students agreed with each of the items related to the learning construct (the agreement level ranged from 30% to 45%). Overall agreement with each of the items in the engagement construct was low (below 30%). As a result, many students did not perceive the examples to be engaging. In both groups, students were most positive about the quality of examples and the least positive about overall engagement with the examples. The Kruskal-Wallis test showed no significant differences in the the mean ratings of the study groups in the quality construct ($\chi^2(1) = .67, p = .41$), learning construct ($\chi^2(1) = .26, p = .61$), or the engagement construct ($\chi^2(1) = .32, p = .57$). The pattern of results did not change when we excluded data of students who used the system fewer than 10 times or fewer than 5 times.

6 SUMMARY AND DISCUSSION

6.1 Summary

This work describes the classroom study that was conducted to address RQ1, RQ2, and RQ3 (described in Section 1), in order to measure the effect of PCEX examples on student's engagement, problem-solving performance, and learning, as compared to non-interactive textbook-style examples. This section summarizes and discusses our findings from the study.

6.1.1 Overall effects on engagement

Students were more engaged in the work with PCEX than with the textbook-style worked examples. The largest effect comes from the work on the next-to-last subtasks. The mean of total time spent on the next-to-last subtasks was 4 times greater in the Experimental group compared to the Control group. The second largest effect was for the work on the first subtasks. The total time that a student spent on the first subtasks was about 1.5 times higher in the Experimental group than in the Control group. We also observed that the mean interaction with explanations was low in both groups, which suggests that the interactivity element implemented in the first subtasks in the Experimental group was not as engaging as that in the next-to-last subtasks (i.e., challenges). We also observed that students in Experimental group had, on average, more practice sessions. Although the effect size was small, it suggests that students in the Experimental group were more interested in returning to the system for practice. We could attribute this to PCEX, since the only difference between the groups was in how examples were presented in the system.

6.1.2 Overall effects on problem-solving performance

We found that working on examples and having higher pre-test scores were associated with solving more Parsons problems correctly, regardless of the treatment group. Among these two predictors, the effect of the work on examples on predicting distinct Parsons problems solved was about 14.7 times larger than the effect of the pre-test score. We also found that work on examples, group, and pre-test were all positive predictors of assignment points. Notably, group and work on examples were found to be the most influential predictors, with higher predictive power than the pre-test scores. We also found that the effect of the group on earliness of submission depended on the amount of work on examples. More work on the examples in the Experimental group was associated with submitting assignments earlier. Furthermore, the interaction was found to be the most important predictor for the earliness of submission; its effect was about two times greater than that of the pre-test score.

6.1.3 Overall effects on learning outcomes

For predicting post-test scores, as expected, pre-test scores had the most predictive power. Yet, we observed that work with examples in the Experimental group increased the students' post-test scores more than the work with examples in the Control group. This difference increased as the amount of work increased in the Experimental group. Furthermore, work on both examples and pre-test positively predicted exam grade across all categories of questions. Work on examples became a more important predictor than pre-test scores as our analysis moved from program comprehension to basic program construction and then to

complex program construction questions in the exam. For the complex program construction questions, group was also a positive predictor, and the effect of the group and work on examples was about two times greater than that of the pre-test scores.

Past research on engagement [16, 55, 56, 67] hints that the observed increase in learning outcomes is caused by both a higher level of cognitive engagement with the new examples (which was our design goal) and a considerably higher level of behavioral engagement (which was directly observed). Our study does not differentiate between these two factors. However, even if the larger share of performance increase was caused by the observed behavioral engagement increase (i.e., time and efforts on task), it is still a highly desired and important outcome in our explored context: the work with voluntary content. In the context of modern e-learning, more and more learning content is released to the students in a voluntary form; i.e., the students are encouraged, but not required, to work with such content and no grade points are offered for doing so. Moreover, a considerable fraction of modern interactive content (simulations, animations, worked examples) are not designed to assess student's performance and are difficult to grade objectively. In this context, not the effectiveness, but the simple amount of work becomes the most critical factor in ensuring the impact of this content on student's learning. While novel types of content, such as animated program examples, are usually proven to be effective in lab studies where the use of this content is required, a number of classroom studies confirmed that this content has low to zero usage when it is offered in a voluntary form [48]. Thus, the very increase of voluntary work with learning content when it is offered in a more engaging and interactive form is an important desired outcome.

6.1.4 Students' feedback

Survey results showed that students were positive toward the quality of explanation, almost neutral toward the helpfulness of the examples for learning, and did not perceive examples to be engaging.

6.2 Discussion

Our learning technology tool, PCEX, was designed to replicate and expand (to another domain) prior research on worked examples in the domains of math and science (e.g., [3, 15, 69]). It was also designed to help students acquire program construction skills by presenting examples in an engaging fashion. Our findings from the classroom study support the positive impact of PCEX examples on student's engagement, problem-solving performance, and learning. Our findings reconfirm the results of previous work in the domains of math and science [2] and program behaviour examples [23, 28, 46] which showed that interactive worked examples help achieve better learning outcomes. However, some aspects of our findings need to be discussed further:

First, while the measure of earliness of submission could be an indicator of how students were prepared better or worse to submit the code, we also acknowledge that submitting coding assignments earlier could have been due to better time management skills or other factors related to a student's self-regulation skills.

Second, on one hand, hypothesis *H1* states that engaging examples increase student's engagement (e.g., time on task) while working with the examples. On the other hand, Hypotheses *H2* and *H3* state that engaging examples lead to better performance in coding and learning outcomes, respectively. We acknowledge that *H2* and *H3* are likely results of *H1*. In

general, time on task is positively correlated with learning and better performance, but the main challenge is how to get the student motivated to spend more time. The use of engaging examples, which we proposed in this work, is a solution to this challenge. Our results showed that students spent more time on the PCEX examples because they were engaged more and, as a result, learned more and obtained better problem-solving performance.

Third, we observed from the survey responses that students in both groups were neutral toward learning from examples and had a negative opinion about being engaged by the examples. These observations contradict what we have found from our engagement, learning, and performance analyses, because results showed that PCEX examples involved students in working more with examples; thus, increasing the time spent on task. In a similar fashion, work with PCEX examples improved both student’s learning and coding skills.

The difference in the quantitative and survey analysis implies that, first, PCEX examples engaged students and improved their learning — but students did not realize that they were engaged and improving their skills. Second, lower ratings might mean that students had higher expectations for the learning tool in the context of our study. Third, it might be that students perceived engagement items in the survey differently than we expected they would. Another reason might be that when students were questioned about how engaging the PCEX examples were (survey), they were comparing it to game apps they play, but when they began to study programming, their actual engagement with the system (log-based data) was reflecting how enthusiastic they felt about interactive PCEX examples, as compared to studying a flat, static book (or the tutoring equivalent of a flat, static book).

Finally, we acknowledge that there is a blurred line between PCEX examples and problems due to the dual nature of “engaging” features within tutoring systems. In general, any engagement asking for additional student’s action will cause the student to move from examples to problem-solving. Yet, we find *examples* to be a more appropriate category for our developed learning tool rather than *problems*, because the “engaging” features that we are using in the PCEX examples have been used in previous learning tools that are known as examples. In particular, in the domain of CS education, animated examples engage the student by asking a question or allowing the student to change the input used by the example. Other aspects that shaped our study include similarities to previous example research in the domains of math and science, such as missing steps that the student has to fill in, or self-explanation prompts that the student has to answer.

6.3 Limitations and future work

This work can be extended in several ways. The first direction for future work would be to improve the PCEX interface features, while the second direction for future work would be to extend the study conducted in this work that assessed the impact of the examples. The last direction for future work would be to share the PCEX examples with a broader audience, as well as to conduct more studies that connect our work with previous work on program behavior examples. The following subsections discuss these and other directions for future work.

6.3.1 Example design

In the current design, PCEX combines explorability and challenge in the same interface, which makes it impossible to separately assess the value of explorability and challenges. In the future, it would be good to conduct a study to evaluate the combined and separate impact

of explorability and challenge in the PCEX activities. Also, the interface of worked examples does not have any features to highlight which concepts or code segments are important to be studied for each block of program code. This is not an issue for small program coding, but the student may easily get lost when the program code has many lines of code. Future work may explore possible ways to emphasize important code segments or motivate the student to think more deeply about the code. One approach would be to use self-explanation prompts in the worked examples. Future research should study how and when to present the self-explanation prompts to enhance learning from examples. Finally, our studies showed that students rarely accessed available hints and explanations, and that it was not clear why these features were not used extensively. Thus, future research needs to run more usability studies for understanding how to improve the design of these features.

6.3.2 The study assessing examples

First, the usage of the practice system was voluntary and, as a result, many of the students did not use the system. For example, only 200 (28%) students used the system in the classroom study. It's not clear how this bias affected this group of students. This limitation stems from the voluntary nature of practice, which might appeal to certain types of students. Although self-selection limits us from having a reliable, fully controlled study, it also helps us to conduct studies that are closer to the natural context of learning to program. Specifically, students who seek help for learning to program have access to abundant resources for learning programming, including online tutors such as the Python tutor, programming platforms such as CodeWorkOut, CloudCoder, and CodingBat, and programming MOOCs and video tutorials that are available on YouTube. The non-mandatory design of the classroom study in this work enabled us to investigate the impact of the proposed learning tool in natural context. Future work may investigate what happens when the system is offered in a mandatory way to see if mandatory work within the practice system would have positive effects on students who would not be using it otherwise.

Second, due to the loosely controlled nature of the classroom study, students may have learned the skills that we controlled for by using resources outside the practice system. This is a factor that we can't control for in our analyses. Future work may investigate the impact of the system in a more controlled way, perhaps by distributing the practice across multiple lab sessions and assessing changes in student's knowledge both before and after each practice session.

Although the classroom study was conducted in a rather large course, it was limited to a specific population. The target course in the study was presented as a mandatory course for engineering students. The participants were mainly studying electrical or civil engineering with a small number of students from other engineering programs. For most students, except for electrical engineering students, this was the only compulsory programming course in their curriculum. Our experience has been that there is a considerable portion of such non-CS students who are not highly motivated to learn to program, especially when compared with computer science (CS) students. Thus, our results from the classroom study may not generalize well for CS major students elsewhere, but may better generalize for CS minors or non-CS majors. Another limitation in the classroom study was that the Control and Experimental groups had different numbers of students using the system. This is a limitation of our analysis and our results could have been influenced by these groups having unmatched numbers. Therefore, a similar study should be conducted in the future with CS majors to validate our observations in the classroom study.

Future studies should also conduct individual interviews with students to understand the possible reasons for reporting lower engagement in the survey, and also to discuss options for enhancing engaging features in PCEX.

6.3.3 Other directions

We plan to connect our work to previous work on program behavior tools. In particular, we plan to study the impact of integrating program behavior and program construction learning activities when they are offered together in a practice system. We also plan to build a repository of PCEX examples to make them available to the public. Finally, we plan to develop an open-source authoring tool to allow instructors and researchers to create and share their own PCEX examples.

References

1. Leona S Aiken, Stephen G West, and Raymond R Reno. *Multiple regression: Testing and interpreting interactions*. Sage, 1991.
2. Robert K Atkinson and Alexander Renkl. Interactive example-based learning environments: Using interactive elements to encourage effective processing of worked examples. *Educational Psychology Review*, 19(3):375–386, 2007.
3. Robert K Atkinson, Sharon J Derry, Alexander Renkl, and Donald Wortham. Learning from examples: Instructional principles from the worked examples research. *Review of educational research*, 70(2):181–214, 2000.
4. Tapio Auvinen, Lasse Hakulinen, and Lauri Malmi. Increasing students awareness of their behavior in online learning environments with visualizations and achievement badges. *IEEE Transactions on Learning Technologies*, 8(3):261273, 2018.
5. Meghan Bathgate and Christian Schunn. The psychological characteristics of experiences that influence science motivation and content knowledge. *International Journal of Science Education*, 39(17):2402–2432, 2017.
6. B. S. Bloom. *Taxonomy of Educational Objectives, Handbook I: The Cognitive Domain*. David McKay Co Inc., New York, 1956.
7. Paul Brna. Searching for examples with a programming techniques editor. *Journal of computing and information technology*, 6(1):13–26, 1998.
8. Peter Brusilovsky. Explanatory visualization in an educational programming environment: connecting examples with general knowledge. In *4th International Conference on Human-Computer Interaction, EWHCI'94*, volume 876 of *Lecture Notes in Computer Science*, pages 202–212. Springer-Verlag, 1994.
9. Peter Brusilovsky and Michael V Yudelson. From webex to navex: Interactive access to annotated program examples. *Proceedings of the IEEE*, 96(6):990–999, 2008.
10. Peter Brusilovsky, Michael Yudelson, and I-Han Hsiao. Problem solving examples as first class objects in educational digital libraries: Three obstacles to overcome. *Journal of Educational Multimedia and Hypermedia*, 18(3):267–288, 2009.
11. Michael D Byrne, Richard Catrambone, and John T Stasko. Evaluating animations as student aids in learning computer algorithms. *Computers & education*, 33(4):253–278, 1999.
12. Richard Catrambone. The subgoal learning model: Creating better examples so that students can solve novel problems. *Journal of Experimental Psychology: General*, 127(4):355–376, 1998.

13. Xiingliang Chen, Antonija Mitrovic, and Moffat Matthews. Learning from worked examples, erroneous examples and problem solving: Towards adaptive selection of learning activities. *IEEE Transactions on Learning Technologies*, pages 1–1, 2019.
14. Xingliang Chen, Antonija Mitrovic, and Moffat Mathews. Investigating the effect of agency on learning from worked examples, erroneous examples and problem solving. *International Journal of Artificial Intelligence in Education*, 2019.
15. Michelene TH Chi, Miriam Bassok, Matthew W Lewis, Peter Reimann, and Robert Glaser. Self-explanations: How students study and use examples in learning to solve problems. *Cognitive science*, 13(2):145–182, 1989.
16. Sandra L Christenson, Amy L Reschly, and Cathy Wylie. *Handbook of research on student engagement*. Springer Science & Business Media, 2012.
17. Stephen Cooper, Wanda Dann, and Randy Pausch. Teaching objects-first in introductory computer science. In *ACM SIGCSE Bulletin*, volume 35, pages 191–195. ACM, 2003.
18. Wanda Dann, Dennis Cosgrove, Don Slater, Dave Culyba, and Steve Cooper. Mediated transfer: Alice 3 to java. In *Proceedings of the 43rd ACM technical symposium on Computer Science Education*, pages 141–146. ACM, 2012.
19. Aleksandar Davidovic, Jim Warren, and Elena Trichina. Learning benefits of structural example-based adaptive tutoring systems. *IEEE Transactions on Education*, 46(2):241–251, 2003.
20. Barbara J Ericson, Mark J Guzdial, and Briana B Morrison. Analysis of interactive features designed to enhance learning in an ebook. In *Proceedings of the eleventh annual International Conference on International Computing Education Research*, pages 169–178. ACM, 2015.
21. Barbara J Ericson, Lauren E Margulieux, and Jochen Rick. Solving parsons problems versus fixing and writing code. In *Proceedings of the 17th Koli Calling Conference on Computing Education Research*, pages 20–29. ACM, 2017.
22. Micaela Esteves and António Mendes. Oop-anim, a system to support learning of basic object oriented programming concepts. In *Proceedings of CompSysTech’2003-International Conference on Computer Systems and Technologies. Sofia, Bulgaria*, 2003.
23. Chris Evans and Nicola J Gibbons. The interactivity effect in multimedia learning. *Computers & Education*, 49(4):1147–1160, 2007.
24. Geela Venise Firmalo Fabic, Antonija Mitrovic, and Kourosh Neshatian. Evaluation of parsons problems with menu-based self-explanation prompts in a mobile python tutor. *International Journal of Artificial Intelligence in Education*, 2019. ISSN 1560-4306.
25. Katherine Wanjiru Getao. An environment to support the use of program examples while learning to program in lisp. In *Proceedings of the IFIP TC13 Third International Conference on Human-Computer Interaction*, pages 1015–1016. North-Holland Publishing Co., 1990.
26. Steven R Hansen, N Hari Narayanan, and Dan Schrimpscher. Helping learners visualize and comprehend algorithms. *Interactive Multimedia Electronic Journal of Computer-Enhanced Learning*, 2(1), 2000.
27. Brian Harvey and Jens Möning. Bringing “no ceiling” to scratch: Can one language serve kids and computer scientists. *Proc. Constructionism*, 2010.
28. Christopher D Hundhausen, Sarah A Douglas, and John T Stasko. A meta-study of algorithm visualization effectiveness. *Journal of Visual Languages & Computing*, 13(3): 259–290, 2002.
29. Petri Ihantola and Ville Karavirta. Two-Dimensional Parson’s Puzzles: The Concept, Tools, and First Observations . *Journal of Information Technology Education: Innova-*

- tions in Practice, 10:1–14, 2011.
30. Slava Kalyuga, Paul Chandler, and John Sweller. Incorporating learner experience into the design of multimedia instruction. *Journal of educational psychology*, 92(1):126–136, 2000.
 31. Slava Kalyuga, Paul Chandler, Juhani Tuovinen, and John Sweller. When problem solving is superior to studying worked examples. *Journal of educational psychology*, 93(3):579–588, 2001.
 32. Slava Kalyuga, Paul Ayres, Paul Chandler, and John Sweller. The expertise reversal effect. *Educational psychologist*, 38(1):23–31, 2003.
 33. Robin H Kay and Liesel Knaack. Assessing learning, quality and engagement in learning objects: the learning object evaluation scale for students (loes-s). *Educational Technology Research and Development*, 57(2):147–168, 2009.
 34. Kandarp Khandwala and Philip J Guo. Codemotion: expanding the design space of learner interactions with computer programming tutorial videos. In *Proceedings of the Fifth Annual ACM Conference on Learning at Scale*, pages 57:1–57:10, 2018.
 35. Andrea W Lawrence. *Empirical studies of the value of algorithm animation in algorithm understanding*. PhD thesis, Georgia Institute of Technology, 1993.
 36. Henry Lieberman. An example-based environment for beginning programmers. In *Artificial intelligence and education*, pages 135–151. Ablex Publishing, Norwood, NJ, 1987.
 37. Marcia C. Linn and Michael J. Clancey. The case for case studies of programming problems. *Communications of the ACM*, 35(3):121–132, 1992.
 38. Tomasz Loboda, Julio Guerra, Roya Hosseini, and Peter Brusilovsky. Mastery grids: An open source social educational progress visualization. In Sara de Freitas, Christoph Rensing, Pedro J. Muoz Merino, and Tobias Ley, editors, *9th European Conference on Technology Enhanced Learning (EC-TEL 2014)*, volume 8719 of *Lecture Notes in Computer Science*, pages 235–248, 2014.
 39. Tomasz D Loboda and Peter Brusilovsky. User-adaptive explanatory program visualization: evaluation and insights from eye movements. *User Modeling and User-Adapted Interaction*, 20(3):191–226, 2010.
 40. Bruce M. McLaren, Deanne Adams, Kelley Durkin, George Goguadze, Richard E. Mayer, Bethany Rittle-Johnson, Sergey Sosnovsky, Seiji Isotani, and Martin van Velsen. To err is human, to explain and correct is divine: A study of interactive erroneous examples with middle school math students. In *7th European Conference on Technology Enhanced Learning (EC-TEL 2012)*, volume 7563 of *Lecture Notes in Computer Science*, pages 222–235, 2012.
 41. Bradley N Miller and David L Ranum. Beyond pdf and epub: toward an interactive textbook. In *Proceedings of the 17th ACM annual conference on Innovation and technology in computer science education*, pages 150–155. ACM, 2012.
 42. Youzou Miyadera, Kunimi Kurasawa, Shoichi Nakamura, Nobuyoshi Yonezawa, and Setsuo Yokoyama. A real-time monitoring system for programming education using a generator of program animation systems. *JCP*, 2(3):12–20, 2007.
 43. Briana B Morrison, Lauren E Margulieux, Barbara Ericson, and Mark Guzdial. Sub-goals help students solve parsons problems. In *Proceedings of the 47th ACM Technical Symposium on Computing Science Education*, pages 42–47. ACM, 2016.
 44. Niko Myller. Automatic prediction question generation during program visualization. In *Proceedings of the Fourth Program Visualization Workshop*, 2006.
 45. Amir Shareghi Najar, Antonija Mitrovic, and Bruce M. McLaren. Learning with intelligent tutors and worked examples: selecting learning activities adaptively leads to better learning outcomes than a fixed curriculum. *User Modeling and User-Adapted*

- Interaction*, 26(5):459–491, 2016. ISSN 1573-1391.
46. Thomas L Naps. Jhavé: Supporting algorithm visualization. *IEEE Computer Graphics and Applications*, 25(5):49–55, 2005.
 47. Thomas L Naps, James R Eagan, and Laura L Norton. Jhavé – an environment to actively engage students in web-based algorithm visualizations. In *ACM SIGCSE Bulletin*, volume 32, pages 109–113. ACM, 2000.
 48. Thomas L. Naps, Guido Rling, Vicki Almstrum, Wanda Dann, Rudolf Fleischer, Chris Hundhausen, Ari Korhonen, Lauri Malmi, Myles McNally, Susan Rodger, and J. ngel Velzquez-Iturbide. Exploring the role of visualization and engagement in computer science education. *ACM SIGCSE bulletin*, 35:131152, 2002.
 49. Timothy J Nokes-Malach, Kurt VanLehn, Daniel M Belenky, Max Lichtenstein, and Gregory Cox. Coordinating principles and examples through analogy and self-explanation. *European Journal of Psychology of Education*, 28(4):1237–1263, 2013.
 50. Jum C Nunnally. *Psychometric Theory: 2d Ed.* McGraw-Hill, 1978.
 51. Fred GWC Paas and Jeroen JG Van Merriënboer. Variability of worked examples and transfer of geometrical problem-solving skills: A cognitive-load approach. *Journal of educational psychology*, 86(1):122–133, 1994.
 52. Jungkook Park, Yeong Hoon Park, Jinhan Kim, Jeongmin Cha, Suin Kim, and Alice Oh. Elicast: embedding interactive exercises in instructional programming screencasts. In *Proceedings of the Fifth Annual ACM Conference on Learning at Scale*, pages 58:1–58:10, 2018.
 53. Dale Parsons and Patricia Haden. 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*, pages 157–163. Australian Computer Society, Inc., 2006.
 54. Peter L Pirolli and John R Anderson. The role of learning from examples in the acquisition of recursive programming skills. *Canadian Journal of Psychology/Revue canadienne de psychologie*, 39(2):240–272, 1985.
 55. Johnmarshall Reeve. How students create motivationally supportive learning environments for themselves: The concept of agentic engagement. *Journal of Educational Psychology*, 105(3):579–595, 2013.
 56. Johnmarshall Reeve and Ching-Mei Tseng. Agency as a fourth aspect of students’ engagement during learning activities. *Contemporary Educational Psychology*, 36(4): 257–267, 2011.
 57. Alexander Renkl. Learning from worked-out examples: A study on individual differences. *Cognitive science*, 21(1):1–29, 1997.
 58. Alexander Renkl and Robert Atkinson. *An Example Order for Cognitive Skill Acquisition*. Oxford University Press, 2007.
 59. Mitchel Resnick, John Maloney, Andrés Monroy-Hernández, Natalie Rusk, Evelyn Eastmond, Karen Brennan, Amon Millner, Eric Rosenbaum, Jay Silver, Brian Silverman, et al. Scratch: programming for all. *Communications of the ACM*, 52(11):60–67, 2009.
 60. Kelly Rivers. *Automated Data-Driven Hint Generation for Learning Programming*. PhD thesis, Carnegie Mellon University, 2017.
 61. Jorma Sajaniemi and Marja Kuittinen. Program animation based on the roles of variables. In *Proceedings of the 2003 ACM symposium on Software visualization*, pages 7–ff. ACM, 2003.
 62. Ron J. C. M. Salden, Vincent Aleven, Rolf Schwonke, and Alexander Renkl. The expertise reversal effect and worked examples in tutored problem solving. *Instructional Science*, 38(3):289–307, May 2010. ISSN 1573-1952.

63. Ron JCM Salden, Kenneth R Koedinger, Alexander Renkl, Vincent Aleven, and Bruce M McLaren. Accounting for beneficial effects of worked examples in tutored problem solving. *Educational Psychology Review*, 22(4):379–392, 2010.
64. Andrew Sears and Rosalee Wolfe. Visual analysis: adding breadth to a computer graphics course. In *ACM SIGCSE Bulletin*, volume 27, pages 195–198. ACM, 1995.
65. Rémi Sharrock, Ella Hamonic, Mathias Hiron, and Sebastien Carlier. Codecast: An innovative technology to facilitate teaching and learning computer programming in a c language online course. In *Proceedings of the Fourth ACM Conference on Learning at Scale*, pages 147–148. ACM, 2017.
66. Teemu Sirkiä. A javascript library for visualizing program execution. In *Proceedings of the 13th Koli Calling International Conference on Computing Education Research*, pages 189–190. ACM, 2013.
67. Ellen A Skinner, Thomas A Kindermann, and Carrie J Furrer. A motivational perspective on engagement and disaffection: Conceptualization and assessment of children’s behavioral and emotional participation in academic activities in the classroom. *Educational and Psychological Measurement*, 69(3):493–525, 2009.
68. Juha Sorva, Ville Karavirta, and Lauri Malmi. A review of generic program visualization systems for introductory programming education. *ACM Transactions on Computing Education (TOCE)*, 13(4):15:1–15:64, 2013.
69. John Sweller and Graham A Cooper. The use of worked examples as a substitute for problem solving in learning algebra. *Cognition and instruction*, 2(1):59–89, 1985.
70. John Sweller, Jeroen JG Van Merriënboer, and Fred GWC Paas. Cognitive architecture and instructional design. *Educational psychology review*, 10(3):251–296, 1998.
71. John Gregory Trafton and Brian J Reiser. The contributions of studying examples and solving problems to skill acquisition. In *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society*, pages 1017–1022. ACM, 1993.
72. Christopher O Walker, Barbara A Greene, and Robert A Mansell. Identification with academics, intrinsic/extrinsic motivation, and self-efficacy as predictors of cognitive engagement. *Learning and individual differences*, 16(1):1–12, 2006.
73. Mark Ward and John Sweller. Structuring effective worked examples. *Cognition and instruction*, 7(1):1–39, 1990.
74. Gerhard Weber. Individual selection of examples in an intelligent learning environment. *Journal of Interactive Learning Research*, 7(1):3–31, 1996.
75. Gerhard Weber and Peter Brusilovsky. Elm-art: An adaptive versatile system for web-based instruction. *International Journal of Artificial Intelligence in Education (IJAIED)*, 12:351–384, 2001.
76. Gerhard Weber and Antje Mollenberg. Elm-pe: A knowledge-based programming environment for learning lisp. In *Proceedings of ED-MEDIA 1994*, pages 557–562. ERIC, 1994.
77. Michael Yudelson, Peter Brusilovsky, and Vladimir Zadorozhny. A user modeling server for contemporary adaptive hypermedia: An evaluation of push approach to evidence propagation. In Cristina Conati, Kathleen McCoy, and Georgios Paliouras, editors, *11th International Conference on User Modeling, UM 2007*, volume 4511 of *Lecture Notes in Computer Science*, pages 27–36. Springer Verlag, 2007.
78. Rui Zhi, Thomas W. Price, Samiha Marwan, Alexandra Milliken, Tiffany Barnes, and Min Chi. Exploring the impact of worked examples in a novice programming environment. In *Proceedings of the 50th ACM Technical Symposium on Computer Science Education, SIGCSE ’19*, pages 98–104. ACM, 2019.