ARTICLE

Investigating the Effect of Agency on Learning from Worked Examples, Erroneous Examples and Problem Solving

Xingliang Chen[1] · **Antonija Mitrovic**[1] (ID) · **Moffat Mathews**[1]

Abstract

Agency refers to the level of control the student has over learning. Most studies on agency in computer-based learning environments have been conducted in the context of educational games and multimedia learning, while there is little research done in the context of learning with Intelligent Tutoring Systems (ITSs). We conducted a study in the context of SQL-Tutor, an ITS that teaches database querying, with students solving a fixed set of ten problems. Before each problem, students worked on a preparatory task, which could be presented as a worked example, erroneous example, or another isomorphic problem. There were two conditions in the study. In the High-Agency condition, students could select the type of preparatory task freely or skip it altogether. In the Low-Agency condition, an adaptive strategy selected preparatory tasks for students on the basis of their performance. The participants were classified as High Prior Knowledge (HPK) or Low Prior Knowledge (LPK), based on their scores on the pre-test. Due to the timing of the study, we had 40 participants who completed all elements of the study. The participants in both Low- and High-Agency conditions improved significantly from the pre- to post-test, and there was no difference between the LPK and HPK students on post-test scores. Therefore, we have not found an effect of agency on learning. The Low Agency condition was beneficial for both HPK and LPK students, while in the High Agency condition there was significant improvement between the pre- and post-test only for the LPK students. In the High-Agency group, the HPK students selected more challenging learning activities, but did not outperform LPK students on the post-test scores. The limitation of our study is the small sample size.

Keywords Problem solving · Worked examples · Erroneous examples · Adaptive selection of learning activities · Agency · Intelligent tutoring systems

✉ Antonija Mitrovic
    tanja.mitrovic@canterbury.ac.nz

1   Intelligent Computer Tutoring Group, Department of Computer Science and Software Engineering, University of Canterbury, Christchurch, New Zealand

Introduction

Agency, which is closely related to self-regulated learning (Zimmerman 2008), refers to the capacity of students to make choices during learning. Self-regulation includes monitoring one's own behavior and its effects, judging it according to personal standards, and affecting self-reaction (Bandura 1991). In order for a student to self-regulate, he/she uses personal agency to make choices on future actions. Although there are attempts to investigate how we can best leverage student agency, it is not clear from literature in which circumstances agency may or may not be beneficial for learning. For instance, advanced students are often good self-regulated learners (Schunk and Zimmerman 2007; Zimmerman 2008), but novices are generally not good at regulating their learning, and hence benefit from instructional choices being made for them (Zimmerman 2000). Mitrovic (2001) and Mitrovic and Martin (2002) also demonstrated that advanced students were better at evaluating their knowledge, while novice students were worse at selecting problems to work on.

Several studies investigated the effect of agency on learning, and reported conflicting findings. Some studies found that increased student agency is associated with higher levels of motivation and involvement, and resulted in better learning outcomes (Snow et al. 2015; Rowe et al. 2011). Tabbers and de Koeijer (2010) demonstrated that giving students control over the pace and order of instructions in an animated multimedia presentation led to higher learning outcomes. Similarly, letting students customize game components has also shown to be positive for learning (Cordova and Lepper 1996; Snow et al. 2015).

On the other hand, Sawyer et al. (2017) focused on the variations in agency within the game Crystal Island. The students in the high-agency condition could control how they obtain knowledge by interacting with the environment and game characters, while the students in the low-agency condition had to follow a prescribed order of actions. The low-agency condition students acquired significantly higher learning gains compared to their peers in the high-agency condition. Nguyen et al. (2018) compared learners in two versions (low agency vs. high agency) in a mathematics educational game. In the low-agency condition, learners were guided to play games in a prescribed sequence, while their peers in the high-agency version could choose the games and the order in which to play them. Unlike the study conducted by Sawyer et al. (2017), they did not find any significant difference in learning between the low and high-agency conditions.

Although there have been many studies on the benefits of learning from worked examples and learning with Intelligent Tutoring Systems (ITSs), most of them represent settings with limited student control. In studies with worked examples, most often examples are presented in the fixed order. On the other hand, ITSs typically select the best problems for students to attempt and students typically have control over asking for help. Agency is not often studied in those kinds of experiments, which motivated us to conduct an experiment reported in this paper.

In a previous study (Chen et al. 2017), we added an adaptive strategy to SQL-Tutor, an ITS that teaches database querying, which selected learning activities to present to the student as preparation for problem solving. The strategy selected either a Worked Example (WE), an Erroneous Example (ErrEx), or a problem to be solved (PS), based on the student's performance, or skipped the preparation task completely in case the

student had shown high performance on previous problems. We used PS, WEs and ErrExs, as these types of learning activities have been shown to be effective learning strategies across a broad range of domains (Kalyuga et al. 2001; van Gog 2011; McLaren and Isotani 2011; Chen et al. 2016a; Durkin and Rittle-Johnson 2012; Stark et al. 2011). Two low-agency conditions in that study were 1) the adaptive condition, and 2) the fixed order condition, which restricted students to learn with a fixed sequence of worked example/problem-solving pairs and erroneous example/problem-solving pairs. The results showed that the adaptive condition was more beneficial for learning: the students who received learning activities adaptively achieved the same learning outcomes as their peers in the fixed order condition, but with fewer learning activities.

However, researchers also warn about negative consequences of too much adaptive support, which can be detrimental to students because it frees them from thinking (Hübscher and Puntambekar 2001). The capability to select learning activities is important for learning; a learner should be able to reflect on what is important to them and what they ought to consider learning about next (Mitrovic and Martin 2003).

Consequently, in the study reported in this paper, we investigated the effects of learning using variations of agency within SQL-Tutor. In the *High-Agency* version, students freely selected preparation tasks (WE, ErrEx, PS or none) before solving problems. In the *Low-Agency* version, the adaptive strategy selected preparation tasks for students based on their performance. Previous research shows that worked examples are more beneficial for novices (Atkinson et al. 2000; McLaren et al. 2008; Sweller et al. 1998). For advanced students, worked examples may become less effective or even lose their effectiveness for learning (Kalyuga et al. 2001; Kalyuga et al. 1998), because the support provided by worked examples is redundant for them. Erroneous examples have so far been shown to be particularly beneficial to students who have amassed a reasonable degree of domain knowledge (Große and Renkl 2007; Tsovaltzi et al. 2012). Therefore, for high prior knowledge students, our adaptive strategy either skips the preparation task altogether (when their performance on previous problems is high), or provides an erroneous example or a problem to solve. Although past research has demonstrated that erroneous examples are more beneficial for students with high prior knowledge, it seems that even students with low prior knowledge can benefit from erroneous examples (e.g., Durkin and Rittle-Johnson (2012), Chen et al. (2016b), Stark et al. (2011)). Therefore, for low prior knowledge students, the adaptive strategy presents either worked examples or erroneous examples, based on their performance on the previous problem. We attempted to answer two research questions:

> **Research Question 1**: *Do the Low- and High-Agency conditions differ on learning outcomes?* Given the results of the Sawyer et al. (2017) study, we expected that the Low-Agency condition would lead to better learning outcomes compared to the High-Agency condition (H1).
> **Research Question 2**: *Are learning outcomes different for students with low or high prior knowledge?* Given the past research showing that the HPK students are good at self-regulating and self-assessing (Zimmerman 2008; Mitrovic 2001), but LPK students commonly benefit from instructional choices being made for them (Zimmerman 2000), we hypothesized that High-Agency would be more beneficial for HPK students (H2a), and the effect of Low-Agency would be more pronounced for LPK students (H2b).

The paper is organized as follows. In the next Section, we overview some studies on the effect of agency on learning in computer-based environments. The following section presents a brief overview of learning from worked examples and erroneous examples, which provides the background for our study. Next, we present the experimental design, followed by the section on results, and finally discussion of our results.

Agency and Learning

Agency, which refers to the level of control that a student has to perform actions in a learning environment, is an important factor and leads to engagement and learning benefits (Bandura 1989; Zimmerman 2008). Scientific research has started to provide evidence that agency can be effective for enhancing motivation, interest, and attitudes that result in positive learning outcomes. Studies of agency have been done with various types of educational environments, such as multimedia learning and educational games, and also with students of different ages. We review some of the approaches in this section, which are summarized in Table 1.

Calvert et al. (2005) report on a study conducted with preschool children in the context of a computer-based storybook. They compared four conditions: 1) full control by an adult who was reading the story and controlling the mouse, while the child was listening; 2) a joint control condition, with the adult reading the story and the child and child controlling the mouse; 3) a child-control condition and 4) a no-exposure condition. The authors reported that children who had control were more attentive and involved in learning than those who were guided by adults; however, they found no difference in the children's memory of the content.

Cordova and Lepper (1996) demonstrate that giving elementary school students control over instructionally irrelevant aspects of an educational game resulted in higher motivation and interest, and led to better learning outcomes on a subsequent math test.

Several studies in the area of multimedia learning showed the *interactivity principle* (Mayer et al. 2003): giving students control over pace and order of instruction reduces cognitive load and increases transfer performance. In one such study conducted with university students, Tabbers and de Koeijer (2010) had a no-control condition, in which students watched a slideshow containing 16 slides, with each slide shown for 13 s with accompanying narration. Students in the other condition were given full control over pace and order of instructional material: they could stop or replay the slides, decide whether to listen to narration, and freely navigate between the slides. The results showed that giving students control over the multiple interactive features led to higher scores on transfer, but at the same time resulted in increased learning time.

Rowe et al. (2011) investigated the relationship between learning and engagement in Crystal Island, a game-based learning environment for microbiology, in which students explore an island where an epidemic has recently spread among a team of scientists. Crystal Island promotes a strong sense of agency, as students decide how to obtain information necessary to solve the problem, by interacting with game characters and other game objects. The authors conducted a study with middle-school students, and found that increased engagement was associated with better learning outcomes and problem solving. Particularly, students who performed better in the game also performed better at post-test, and these students were more successful at gaining information during play.

Table 1 Overview of discussed studies

Paper	Environment	Participants	Conditions	Agency	Effect of agency
(Calvert et al. 2005)	Storybook	53 preschool children	Four	Control over interaction with the book	Increased attention and engagement
(Cordova & Lepper 1996)	Game	72 elementary school children	Four (personalization x choice)	Control over non-pedagogical game elements	Increased motivation and learning
(Tabbers & Koeijer 2010)	Multimedia slideshow	52 University students	Two (control vs no control)	Controlling pacing and order	Increased learning, longer learning time
(Rowe et al. 2011)	Game	153 middle school students	One	Students decide which actions to perform	Engagement highly correlated with learning
(Sawyer et al. 2017)	Game	98 University students	Three (High / Low / None)	Pedagogical (order of actions)	Lower learning gain
(Snow et al. 2015)	Game	University students: 75 in Study 1, 70 in Study 2	One	Students decide which activities to perform	Control choice patterns result in better performance
(Nguyen et al. 2018)	Game	159 middle school students	Two (Low vs High)	Order of games	No difference in learning

Sawyer et al. (2017) explored variations in student agency in Crystal Island, in three conditions. In the High Agency condition, students could freely explore the environment. The Low Agency condition required students to visit a series of game locations in a prescribed order, where they had to complete specific problem-solving actions. Finally, the No Agency condition provided students with a video of an expert playing the game to model an ideal path for solving the problem scenario. The Low Agency students made more incorrect actions, but also achieved the highest learning gains. The authors explain that such results were likely due to the extensive engagement with instructional materials.

Snow et al. (2015) conducted two studies with university students to investigate the effect of agency by analyzing students' choice patterns within the game-based system iSTART-2. Student could choose various learning activities, and could also personalize the environment, which provided students a sense of control. The results showed that the higher quality self-explanations and better performance within the games were closely related to a student's ability to exercise control choice patterns, as opposed to disordered (i.e., random) choice patterns.

Nguyen et al. (2018) also compared Low/High Agency conditions to investigate whether limiting agency could lead to high engagement and improved learning outcomes in a mathematics educational game, Decimal Point. The High Agency condition allowed students to choose how many and in what sequence they will play the game. The Low Agency condition guided students to play games in a prescribed order. Unlike Sawyer et al. (2017) study, they did not find increase in learning with Low Agency compared to when students learned with High Agency. The authors attributed this result to the effects of indirect control or teachers' pressure. Additionally, students in the High Agency condition did not exhibit much agency, due to the design of the game layout which led more than half of the High Agency students to play the games in the same order as the students in the Low Agency condition.

However, research suggests that increasing student agency may not be beneficial for all students (Katz et al. 2006). Agency may lead to non-optimal learning such as increased learning time (Tabbers & de Koeijer 2010), or to difficulties with selecting, organizing and integrating information (Mayer 2004; Kirschner et al. 2006).

Learning from Worked/Erroneous Examples

A worked example (WE) consists of a problem statement, its solution, and additional explanations, and therefore provides a high level of assistance to students. Cognitive Load Theory (CLT) states that problem solving without support produces a high level of cognitive load for novices because of unproductive search procedures (Sweller et al. 1998), as a student needs to do a lot of reasoning while solving a problem with no feedback. *Intrinsic load*, *germane load,* and *extraneous load* are three different loads for the working memory in the CLT. *Intrinsic load* refers to the complexity of the learning materials (the number of interacting information elements a task contains) and the learner's level of prior domain knowledge. The intrinsic load is higher when a novice student is studying a more complicated problem. It is possible to appropriately manage the intrinsic load by dividing the initial learning goal into a series of sub-goals that require fewer processing resources. *Germane load* is considered as the information that is related

to the learning materials, such as self-explanations, which is a metacognitive process in which students explain provided learning material to themselves (Chi et al. 1994; Renkl 1997). *Extraneous load* is caused by information during learning that does not directly contribute to learning. Extraneous load refers to the load imposed on students' working memory that does not contribute to learning. Extraneous load and germane load both depend on the way the task is presented, but only germane load contributes to learning (Clark et al. 2011). In order to solve a problem, a learner must consider both the current problem description and the goal state, find the differences between the problem description and the goal state, and find the problem-solving operators to reduce these differences. Many interacting elements associated with this learning process impose an extraneous load that interferes with learning. WEs may significantly relieve this load on students' working memory thus allowing the students to learn faster and solve more complex problems (Sweller et al. 1998; Sweller 1988).

Numerous studies have investigated the effects of learning from WEs compared to learning from tutored problems solving (TPS) when the ITS has control over learning tasks (Schwonke et al. 2007; Schwonke et al. 2009; McLaren et al. 2008; Salden et al. 2010; McLaren and Isotani 2011). These studies showed that WEs result in shorter learning times, but commonly there was no difference in the knowledge gain compared to learning from TPS. Contrary to those results, Najar and Mitrovic (2014) conducted a study with SQL-Tutor (Mitrovic 2003; Mitrovic and Ohlsson 1999; Mitrovic 1998). (Mitrovic 2003; Mitrovic and Ohlsson 1999; Mitrovic 1998). They compared how students learn from a fixed set of problems, presented to students as examples only (EO), tutored problem only (PO) and alternating examples and tutored problems (AEP). They found that students learned more in the PO and AEP conditions than from EO condition; furthermore, presenting alternating isomorphic pairs of WE and TPS (AEP) to students produced the greatest learning. Also, they found that AEP significantly improved novices' conceptual knowledge in comparison with PO condition, but advanced students did not improve significantly from EO condition. Najar et al. (2014) later compared an adaptive strategy to the alternating worked examples and problem-solving strategy (AEP). Similar to Kalyuga and Sweller (2005) study, the adaptive strategy was based on a measure of cognitive efficiency, where the performance (P) was calculated from the assistance the students received, and the students rated their mental effort (R) after solving each problem. The results showed that the adaptive condition led to better learning outcomes. Additionally, the adaptive condition resulted in shorter learning times for novices compared to their peers in the AEP condition. The advanced students in the adaptive condition learned more than their counterparts in the AEP condition.

In contrast to WEs, erroneous examples (ErrExs) present incorrect solutions and require students to find and fix errors. Presenting students with erroneous examples may help them become better at evaluating problem solutions and improve knowledge of correct concepts (van den Broek and Kendeou 2008; Stark et al. 2011), and procedures (Große and Renkl 2007), which, in turn, may help students to learn material at a deeper level. The presentation of ErrExs can vary, depending on the kind and amount of feedback provided, and the choice and sequencing of the learning activities (e.g. ErrExs provided in addition to problem solving, or WEs). Researchers have started to investigate empirically the use of erroneous examples in order to better understand whether, when, and how the erroneous examples make a difference to learning. Siegler (2002)

demonstrated that learners were more likely to learn and think deeply about correct concepts that applied to a range of problem types while they explained both correct and incorrect solutions during a brief tutoring session in comparison to their peers who only explained correct solutions. Siegler and Chen (2008) compared WEs to ErrExs for mathematical equality problems. Children who studied and self-explained both the correct and erroneous examples had better learning outcomes than those who received and self-explained only correct examples. Große and Renkl (2007) found the learning benefits of ErrExs for students with a high level of prior knowledge, but not for novices.

Tsovaltzi et al. (2012) indicated that 6th-grade students improved their metacognitive abilities after learning from erroneous examples of fractions with inter-active help using an ITS. Erroneous examples with interactive help also improved 9th and 10th grade students' problem-solving skills and conceptual knowledge. The combination of WEs and ErrExs was shown to lead to improvements in both conceptual knowledge and procedural skills in Algebra (Booth et al. 2013). In our previous study (Chen et al. 2016a), we investigated whether ErrEx in addition to WEs and tutored problem solving would lead to better learning. We compared students' performance in two conditions: alternating worked examples and problem solving (AEP) condition and a fixed sequence of worked examples/problem solving pairs followed by erroneous examples and problem-solving pairs (WPEP). The results showed that the addition of ErrExs improved learning on top of WEs and PS. When students were asked to explain why incorrect solutions were wrong, they engaged in deeper cognitive processing. Therefore, they were better prepared for the concepts required in the next isomorphic problem.

The discussed studies focused on investigating the effect of presenting varying levels of assistance (worked examples, tutored problems solving or erroneous examples) to students. However, there is a lack of studies that focus on the effect of agency when learning with ITSs. Students who are attempting to self-regulate often face limitations in their own knowledge and skills, which can cause cognitive overload and decreased interest and persistence (Duffy and Azevedo 2015; Harley et al. 2015). Mitrovic and Martin (2003) investigated the effect of scaffolding and fading problem selection in SQL-Tutor. They found that the fading problem selection strategy was effective, in which the system initially selected the problem for the students and explained why particular problems are good, and over time released the control over problem selection to students. Azevedo et al. (2016) demonstrated that deploying adaptive scaffolding and feedback in self-regulated learning produced better learning outcome compared to no scaffolding and feedback.

The goal of our study is to compare the learning benefits of variations of agency at different levels of prior knowledge (lower, higher). To the best of our knowledge, there are no studies that compare the effectiveness of low- and high-agency in ITSs.

Experimental Design

Participants

The study was performed with the volunteers from COSC265, a second-year database course at the University of Canterbury. Before the study, the students had learned about

SQL in lectures and also had one lab session. There were 67 volunteers who signed the consent form, but 27 participants were excluded because they did not complete all phases of the study. Of the remaining 40 students, 11 were females, 25 were in the age range 18–20, 8 in the age range 21–23 and the remaining 7 were aged between 24 and 29. The majority (75%) were NZ Europeans, two participants were 2 British and the others were Asian.

Pre/Post Tests

At the beginning of the session, the students took an online pre-test. The pre-test had eleven questions. Questions 1 to 6 measured conceptual knowledge and were multi-choice or true-false questions (with the maximum of 6 marks). Questions 7–9 focused on procedural knowledge; question 7 was a multi-choice question (1 mark), question 8 was a true-false question (1 mark), while question 9 required the student to write an SQL query for a given problem (4 marks). The last two questions presented incorrect solutions to two problems and required students to correct them, thus measuring debugging knowledge (6 marks). The maximum mark was 18. Students received the post-test of similar complexity and length to the pre-test after completing all learning activities. The pre/post-tests are given in the Appendix.

Cronbach's alpha value for the pre-test is .386, and for the post-test it is .214. The low values of Cronbach's alpha for knowledge tests that cover a range of various aspects are not unusual (Taber, 2018). There are additional reasons for such low values of Cronbach's alpha. Our pre/post-tests needed to be short, as the duration of the whole session was 100 min. Therefore, we generated a set of questions to get an understanding of the student's domain knowledge. Cronbach's alpha tends to increase with the number of questions, and our tests were short (11 questions each). There are no redundancies in the tests, as each question covers different SQL concepts.

The participants were labeled as LPK students if their pre-test score was less than the Split score (S), defined in Eq. 1. In the equation, M represents the median pre-test score (67%) from our previous study (Chen et al. 2017), while X_n represents the pre-test score of student n. S_n represents the Split score after student n completed the pre-test. Please note that the value of S changes dynamically as students complete the pre-test.

$$S_n = \frac{S_{n-1} + X_n}{2} \quad (S_0 = M, n \geq 1) \tag{1}$$

Materials

The study was conducted in the context of SQL-Tutor, which is a mature, constraint-based ITS for teaching SQL (Structured Query Language) (Mitrovic 2003; Mitrovic and Ohlsson 1999; Mitrovic 1998). We developed three modes of SQL-Tutor to correspond to WE, ErrEx, and PS. Figure 1 shows a screenshot of the problem-solving mode we used in the study. The left pane shows the structure of the database schema, which the student can explore to gain additional information about tables and their attributes, as well as to see the data stored in the database. The middle pane is the problem-solving space. The right pane displays the feedback on the student's solution

Fig. 1 Problem-solving mode of SQL-tutor

once s/he submits his/her solution. SQL-Tutor supports six levels of feedback (Mitrovic and Martin 2000). *Simple* (positive/negative) feedback, which is the lowest level of assistance, specifies whether the solution is correct or not. *Error Flag* feedback indicates the part of the solution that is incorrect (as shown in Fig. 1). The *Hint* level addresses a specific error and states the domain principle violated by the student's solution. *Partial Solution* provides the correct version of a clause in which the student made a mistake. Other two feedback levels are *List All Errors*, which provides Hint-level feedback messages for all mistakes, and *Complete Solution*, which provides the full solution. The default feedback level is *Simple* feedback for the first submission unless overridden by the student. The feedback level is automatically moved up to the Hint level, but the student can ask for any feedback level at the time of submitting the solution.

The interface of the WE mode is illustrated in Fig. 2. An example problem with its solution and explanation is presented in the center pane. A student can click the "*Continue*" button to confirm that they s/he has studied the example. The ErrEx mode is illustrated in Fig. 3. An incorrect solution is provided, and the student's task is to analyse the solution, find and correct the error(s). The student can submit the solution to

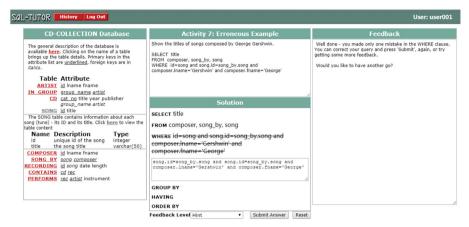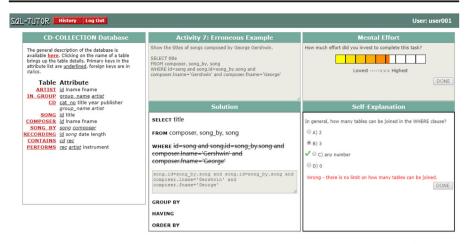**Fig. 2** Worked example mode of SQL-tutor

Fig. 3 Erroneous example mode of SQL-tutor

be checked by SQL-Tutor multiple times, similar to the problem-solving mode. In the example illustrated in Fig. 3, the student has marked the WHERE clauses as being incorrect and has entered answers that s/he believes is correct.

Erroneous solutions presented as ErrEx were selected from the set of incorrect solutions submitted by the participants of the Najar and Mitrovic (2013) study, which used the same set of problems as in our study. We analyzed the 465 submissions to the ten problems corresponding to the erroneous examples in our study from the 2013 problem-only condition. There were on average 5.59 submissions per problem (sd = 2.19). We identified the most frequent misconceptions (or the top two misconceptions) that students had about the relevant domain concepts. Erroneous examples include errors that address these misconceptions.

The students were additionally asked to rate the mental effort (R) and answer the self-explanation prompts required to complete a learning activity (i.e. WE, PS, or ErrEx). Figure 4 illustrates the interface of the mental effort rating bar (Lowest: yellow color, Highest: red color). Research has shown that WEs improve conceptual knowledge more than procedural knowledge, whereas problem solving results in higher levels of procedural knowledge (Schwonke et al. 2009; Kim et al. 2009). For that reason, different types of self-explanation should be provided. Consequently, Najar and Mitrovic (2013) designed the conceptual-focused SE (C-SE) prompts and the procedural-focused SE (P-SE) prompts, to complement learning with WEs and PS. C-SE prompts require the student to answer questions about relevant domain concepts after PS, while P-SE prompts required explanations of solution steps after WEs. A C-SE prompt is presented after a problem is solved, in order to aid the student in reflecting on the concepts covered in the problem they just completed (e.g. *What does DISTINCT in general do?*). On the other hand, P-SE prompts are provided after WEs to assist learners in focusing on problem-solving approaches (e.g. *How can you specify a string constant?*). Therefore, C-SE and P-SE prompts were used in the previous study (Najar and Mitrovic 2013) to increase learning. In our study, in order to keep our experimental design consistent with that of (Najar and Mitrovic 2013), participants received C-SE prompts after problems, and P-SE prompts after WEs, to complement learning activities so that both conceptual and procedural knowledge is supported. Since erroneous

Fig. 4 Mental effort rating and conceptual self-explanation

examples provide both correct and incorrect steps and require students to solve the incorrect steps, which refer to the properties of problems and WEs, we provided P-SE and C-SE prompts alternatively after ErrExs. Figure 4 also illustrates a C-SE prompt, located at the bottom right. The student answered the prompt incorrectly; in return, the system indicated the correct option and provided feedback on the option the student selected. Figure 5 shows a similar example, but with positive feedback in response to the student's answer to the P-SE prompts. Students can attempt each SE prompt only once.

Procedure

The study was conducted in a single, 100-min-long session. Figure 6 illustrates the design of the study. Once participants completed the online pre-test, they were classified as Low Prior Knowledge (LPK) or High Prior Knowledge (HPR) students, based

Fig. 5 Procedural self-explanation

	Low Agency	High Agency
	Online Pre-Test	
	10 Problems and 10 preparation tasks in isomorphic pairs	
Pair 1	**1st task**: Lowest conceptual score: WE; Lowest procedural score: PS; Lowest debugging score: ErrEx **2nd task**: problem	
Pair 2 to 10	**1st task:** LPK: WE, 1- or 2-error ErrEx; HPK: 1- or 2-errors ErrEx, PS or skip **2nd task:** problem	**1st task:** WE, 2- or 1-error ErrEx, PS or skip **2nd task:** problem
	Each problem followed by a C-SE prompt and each example followed by a P-SE prompt	
	Online Post-Test	

Fig. 6 Study design

on their pre-test scores. Then they were randomly assigned to one of the two instructional conditions: (1) Low-Agency condition, which adaptively selected preparation tasks (WE or ErrEx for LPK, and ErrEx or PS for HPK students), or (2) High-Agency condition, in which students could select preparation tasks (WE, ErrEx, PS or skip) by themselves. The participants worked on 20 tasks, organized into ten isomorphic pairs and sorted by increasing complexity. Even-numbered tasks were problems to solve. Odd-numbered tasks are preparatory tasks, and could be presented either as WEs, ErrExs (with one or two errors), or problems to solve. The first preparatory task was different from the others because the student models were empty. For that reason, we used the pre-test score to determine the type of the first preparatory task. If the conceptual score on the pre-test was lower than the procedural and debugging scores, the first preparation task was presented as a worked example. If the student's procedural score was lower than the other two scores, s/he received a problem as the first task. If the lowest score was on debugging questions, the first task was presented as an ErrEx.

Adaptive Strategy

The adaptive strategy uses Cognitive Efficiency (CE) to decide what the preparation task should be, based on the student's problem-solving performance on the previous problem. CE is computed as the quotient between the problem-solving score on the most recent problem (P) and the self-reported mental effort score (R), $CE = P \div R$, as originally proposed in (Kalyuga and Sweller 2005). Both scores had the same range, 0 (lowest) to 9 (highest). The participants were asked to report the effort and answer the SE prompt after each task they completed (Fig. 4).

In constraint-based tutors, domain knowledge is represented as a set of constraints (Mitrovic 2003; Ohlsson 1994). Each constraint has two conditions, the relevance and satisfaction condition. When the student's solution is matched to a constraint, if the relevance condition of a constraint is met, the satisfaction condition is checked next. Therefore, a relevant constraint can either be violated (when the satisfaction condition is not met) or satisfied. A solution is incorrect if it violates one or more constraints; therefore, the solution can be scored based on the violated or satisfied constraints. SQL-Tutor contains six key concepts, represented by the SELECT, FROM, WHERE,

GROUP BY, HAVING and ORDER BY clauses. Each concept can be scored according to how many constraints are violated for that concept. The student's score for a clause is calculated using Eq. 2, in which C_v represents the number of violated constraints, while C_r represents the number of relevant constraints. When a solution does not violate any constraints for a clause, its score C is 1.

$$C = 1 - {}^{C_v}\!/_{C_r} \tag{2}$$

However, Eq. 2 does not produce accurate scores when there are several violated constraints that come from the same mistake. For instance, if a solution missed one attribute in the FROM clause, several constraints will be violated. Equation 2 results in a big penalty in that case. To deal with this situation, we investigated Eq. 3 instead.

$$C = \begin{cases} \log_{(1/C_r)}(C_v/C_r), & 0 < C_v < C_r \\ 1, & C_v = 0 \end{cases} \tag{3}$$

We compared the scores produced by a human marker for the problem-solving question from the pre-test (Question 9). The mean score for 58 solutions was .77 (sd = .303). Equation 3 produces scores with the mean of .84 (sd = .26). The correlation between manual scores and the scores produced by Eq. 3 is significant and high ($r = .864$, $p = 0$). However, a student's incorrect solution may not violate all relevant constraints. For example, one solution for Question 9 violated 5 out of 10 relevant constraints, and the human marker allocated 0 marks to it, while Eq. 3 resulted in the score of .301. For solutions with a higher number of relevant constraints, the difference between manual and automatically-calculated scores was larger. To handle this situation, we used Eq. 4. C is 0 if the number of violated constraints is equal to the number of relevant constraints, as in Eq. 2. The scores produced by Eq. 4 had the mean of .808 (sd = .282), and the correlation was stronger ($r = .921$, $p = .000$).

$$C = \begin{cases} \log_{(1/C_r)}(C_v/.5C_r), & 0 < C_v < C_r \\ 1, & C_v = 0 \\ 0, & C_v = C_r \end{cases} \tag{4}$$

Equation 5 calculates the solution score P as the sum of scores for all clauses the student specified (with a maximum of 6 clauses). Note that the clause score is zero and Eq. 4 is not applied if the clause is empty. The weight of a clause (W_i) is calculated on the basis of the ideal solution for a problem. C_t is the number of constraints relevant for the ideal solution. The weight of a clause (W_i) is calculated as a quotient of the number of relevant constraints for that clause (C_{ci}) and C_t, as shown in Eq. 6.

$$P = \sum_{i=1}^{n} W_i\, C_i \tag{5}$$

$$W_i = {}^{C_{ci}}\!/_{C_t} \tag{6}$$

The maximum value for P when using Eq. 5 is 1 (when the student's solution is correct). Since the maximum value of R is 9, we need to have the same maximum value for performance, which gives us the final Eq. 7:

$$P = 9 \sum_{i=1}^{6} W_i C_i \qquad (7)$$

The CE score is computed after the student provides the mental effort rating. Figure 7 shows the relationship between CE and preparation tasks, while Figs. 8 and 9 illustrate how the preparation task (i.e. the first element of a pair of learning activities) is selected, based on CE and students' prior level of knowledge. There were two types of erroneous examples: ErrExs with one error (1-error ErrEx) or two (2-error ErrEx). For HPK students, if CE was higher than 1, that illustrated very high problem-solving performance, and the preparation task was skipped. CE below 1 and greater than 0.75 showed a relatively good performance on the previous problem, and the preparation task chosen was a problem to be solved. An HPK student received a 2-error ErrEx before the next problem if CE was between 0.75 and 0.5, otherwise, they received a 1-error ErrEx if CE was lower than 0.5. For LPK students, if CE was higher than 0.5, the preparation task was a 2-error ErrEx. If CE was below 0.5 and greater than 0.25, they received a 1-
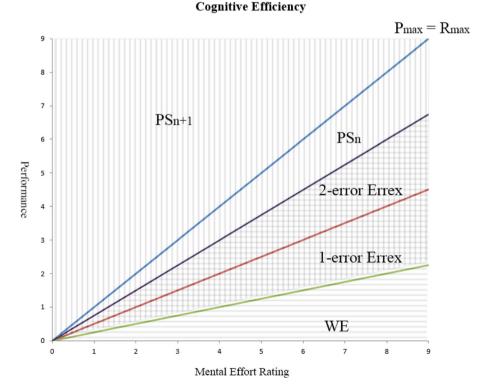

Fig. 7 Relationship between CE and preparation tasks

Fig. 8 Adaptive selection of learning activities for LPK students

error ErrEx as the preparation task. A worked example was provided to an LPK student if his/her CE was below 0.25.

Learning Activity Selection in the High-Agency Condition

The High-Agency condition allowed students to select preparation tasks on their own, as illustrated in Fig 10.

Results

Our study was conducted at a time when the participants had assessments due in other courses they were taking. Since participation was voluntary, only 40 students completed all phases of the study. Such a big attrition rate necessitated further investigation. We compared the incoming knowledge (i.e. the pre-test scores) of the participants who completed the study with those who abandoned it, in order to identify whether they were comparable or whether it was the weaker students who did not complete the study.When comparing the pre-test scores (Table 2), we found no significant differences between the scores of those students who completed or abandoned the study. As we mentioned above, the pre−/post-test consisted of conceptual, procedural, and debugging questions. There were also no significant differences in the scores for

Fig. 9 Adaptive selection of learning activities for HPK students

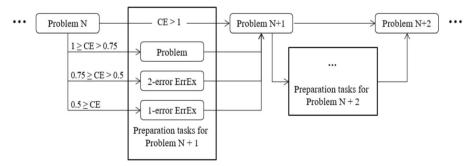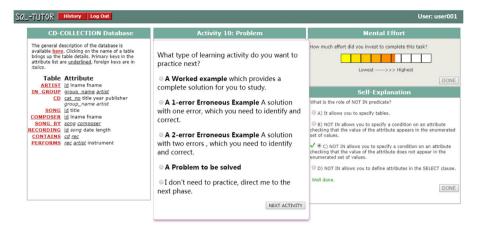Fig. 10 Self-selection prompt

conceptual, procedural, and debugging questions. Therefore, the 40 remaining partic-
ipants had the same level of background knowledge as the other participants.

Research Question 1: Do the Low- and High-Agency Conditions Differ on Learning Outcomes?

There were 20 participants in the Low-Agency condition. We removed an outlier from
the High-Agency condition, leaving 19 participants. Table 3 presents the test scores for
the participants in the two conditions. We developed a repeated measures mixed effects
model, with the knowledge score as the within-subject factor with two levels (pre- and
post-test score), and the group as the between-subject factor. There was no significant
interaction between the test scores and the group, $F(1,37) = .002$, $p = .963$. There was a
significant difference in test scores between the pre- and post-test, $F(1,37) = 55.56$,
$p < .001$, partial $\eta^2 = .661$, but not a significant difference between groups, $F(1,37) = 1968.506$, $p = .094$, partial $\eta^2 = .074$. Therefore, the participants in both group
significantly improved their knowledge from the pre- to the post-test.

We calculated the effect size (Cohen's d), with the following assumption: $d \geq 0.8$
(large effect), $d \geq 0.5$ (medium effect) and $d \geq 0.2$ (small effect) (Cohen 1988). The
effect size for the post-test ($d = 0.47$) was average. In the Low-Agency condition, the
pre-test and post-test scores were positively correlated, and the correlation was

Table 2 Pre-test scores (%) for participants who completed/abandoned the study

	Completed (N = 40)	Abandoned (N = 27)
Overall	61.08 (13.5)	57.65 (11.82)
Conceptual	47.92 (16.96)	50.62 (18.19)
Procedural	72.97 (19.2)	66.38 (25.59)
Debugging	62.36 (23.55)	55.95 (16.36)

Note: all tables present means and standard deviations (given in parentheses).

Table 3 Statistics for the two conditions

	Condition	
	Low Agency (N = 20)	High Agency (N = 19)
Pre-Test (%)	63.50 (12.42)	58.17 (14.62)
Post-Test (%)	82.46 (9.07)	76.92 (13.98)
Pre/Post-test Correlation	r = .48, p = .032	r = .075, ns
Learning time	91.82 (40.13)	97.06 (49.54)

significant. On average, the participants spent 94 min interacting with the learning tasks, and there was no significant difference in interaction time between the two conditions.

As explained earlier, preparation tasks for the Low-Agency condition were selected depending on Cognitive Efficiency (CE) scores on the previous problem and the students' prior knowledge. Therefore, HPK students in the Low-Agency condition could receive PS, 2-error/1-error ErrEx as the preparation task, or skip to the next PS, while LPK students in the Low-Agency condition could receive a 2-error/1-error ErrEx or a WE. The students in the High-Agency condition could select any type of learning activity as the preparation task or choose to skip the preparation task entirely to move on to the next PS. The CE scores were calculated in both conditions after each problem was solved. Table 4 reports the CE scores, and the number of activities of different types the participants completed in the two conditions. There was no significant difference between the two conditions on the CE scores. On average, the students completed 18 learning activities, ten of which were problems to be solved.

Table 4 also reports the number of preparatory tasks the students in the two conditions completed. The Low-Agency group received significantly fewer problems (p = .001) and WEs (p < .001), but more ErrExs (p < .001) than the High-Agency group. The participants in the High-Agency group selected approximately the same number of preparation tasks of different types (i.e. problems, ErrExs, WEs and skips),

Table 4 Student performance in the two conditions

	Condition		U, p
	Low Agency (N = 20)	High Agency (N = 19)	
Cognitive Efficiency (CE)	1.9 (0.78)	1.76 (0.62)	ns
Post-Test/CE Correlation	r = .53, p = .009	r = .59, p = .008	
Number of learning activities	18.1 (3.21)	17.79 (2.92)	ns
Number of problems	0.50 (0.95)	2.37 (2.45)	305, .001***
Number of ErrExs (1/2-error)	6.80 (3.58)	2.47 (2.46)	64.5, .000***
Number of WEs	0.80 (0.77)	2.89 (2.42)	66.5, .000***
Skips	1.90 (3.21)	2.21 (2.92)	ns

***, ** and * denote significance at the p = .001, .01 and .05 levels; "ns" stands for "not significant")

while for the Low-Agency group the adaptive strategy selected the activities according to the CE scores.

Research Question 2: Are Learning Outcomes Different for Students with Low/High Prior Knowledge?

Once a student submitted the pre-test, SQL-Tutor classified him/her immediately as HPK or LPK, as described previously. To confirm whether Eq. 1 identified HPK/LPK students correctly, we additionally used the median split on the pre-test to classify students. The median split on the pre-test results in 22 LPK and 18 HPK students, which is the same as using Eq. 1.

Table 5 presents the test scores and the normalized learning gains for the LPK/HPK students in the two conditions. We constructed a generalized mixed model with the normalized learning gain as the response variable, and the group (i.e. agency) and the level of student's prior knowledge (i.e. LPK or HPK) as the between-subject factors. The interaction between group and level was not significant, $F(1,35) = 4.062$, $p = .052$, partial $\eta^2 = .104$. There was a significant main effect of the level, $F(1,35) = 5.87$, $p = .021$, $\eta^2 = .144$. Therefore, the normalized gain of LPK students was significantly higher than the normalized gain of HPK students.

Table 6 presents the CE scores and information about the activities the students performed in the two groups. We developed a general linear model, with the student level (LPK or HPK) and group (i.e. agency) as fixed factors. There was no significant interaction between group and level for CE, but there was a significant main effect of level, $F(1.35) = 5.386$, $p = .026$, partial $\eta^2 = .133$. For the total number of learning activities the participants completed, there was no significant interaction between group and level, but there was a significant main effect of level, $F(1.35) = 6.337$, $p = .017$, partial $\eta^2 = .153$. For the number of problems received as preparatory activities, there was no significant interaction between group and level, but there was a significant main effect of group, $F(1.35) = 9.997$, $p = .003$, partial $\eta^2 = .222$. For ErrEx, there was a significant interaction between group and level, $F(1,35) = 6.68$, $p = .014$, partial $\eta^2 = .16$. For the number of WEs, there was no significant interaction between group and level, but there was a significant main effect of group, $F(1.35) = 11.225$, $p = .002$, partial $\eta^2 = .013$. For the number of skips, there was no significant interaction between group and level, but there was a significant main effect of level, $F(1.35) = 6.337$, $p = .017$, partial $\eta^2 = .072$.

Table 5 Detailed test scores for LPK/HPK students

		Pre-test %	Post-test %	Normalized learning gain
Low Agency (N = 20)	LPK (10)	53.62 (8.84)	78.40 (9.82)	.53 (.24)
	HPK (10)	73.38 (5.53)	86.59 (6.33)	.49 (.27)
High Agency (N = 19)	LPK (12)	49.05 (8.51)	77.20 (14.69)	.55 (.28)
	HPK (7)	73.80 (7.53)	76.45 (13.79)	.10 (.49)

Table 6 CE and the number of activities for LPK/HPK students

	Condition	Level	
		LPK	HPK
Cognitive Efficiency (CE)	Low Agency	1.59 (.49)	2.21 (.90)
	High Agency	1.62 (.51)	2.01 (.74)
Number of learning activities	Low Agency	20 (0)	16.20 (3.71)
	High Agency	18.08 (2.71)	17.29 (3.40)
Number of problems	Low Agency	0 (0)	1 (1.16)
	High Agency	2.25 (2.45)	2.57 (2.64)
Number of ErrExs (1/2-error)	Low Agency	9 (0.82)	4.60 (3.95)
	High Agency	2.42 (2.15)	2.57 (3.10)
Number of WEs	Low Agency	1 (0.82)	0.60 (0.70)
	High Agency	3.33 (2.67)	2.14 (1.86)
Skips	Low Agency	0 (0)	3.80 (3.71)
	High Agency	1.92 (2.71)	2.71 (3.40)

In the High-Agency condition, students selected the preparation task on their own. There was no significant difference between LPK and HPK students on the post-test scores (Table 5). Surprisingly, the CE scores of HPK and LPK students in the High-Agency condition were approximately the same, and they completed the same number of activities (Table 6). To further investigate this interesting finding, we analyzed the student's task selection 'step size' and self-assessment accuracy between LPK/HPK students in the High-Agency condition, based on the Cognitive Efficiency and students' task selection. Figure 11 presents the relationship between the student's selection (High-Agency) and the system's selection (Low-Agency), which could be used to infer a recommended 'step size' for task selection (e.g., a student selected WE as the preparation task and the system selected PS as the preparation task means a step size of +3). A positive step size means a recommendation to select a more challenging preparation task, a step size of 0 means a student selected the same preparation task as the system's selection, and a negative step size means a recommendation to select a simpler preparation task.

	WE	+4	+3	+2	+1	0
Student Selection	2-error ErrEx	+3	+2	+1	0	-1
	1-error ErrEx	+2	+1	0	-1	-2
	PS	+1	0	-1	-2	-3
	Skip to next PS	0	-1	-2	-3	-4
		Skip to next PS	PS	2-error ErrEx	1-error ErrEx	WE
		System Selection				

Fig. 11 Step Size of Preparation Task Selection

The mean step size for the High-Agency students is 0.49 (sd = 1.28). Specifically, the mean step size for the HPK students is 1.32 (sd = 1.14), while for the LPK students the mean step size is −0.07 (sd = 1.08). The difference is significant (U = 21, p = .038). The selections of LPK students were close to the selections the adaptive strategy would make, which explains the significant improvement for LPK students from the pre- to post-test.

Discussion and Conclusions

Some previous studies found that the increased student agency resulted in better learning outcomes (Snow et al. 2015; Rowe et al. 2011), while Sawyer et al. (2017) found that in their study the Low-Agency condition led to higher learning gains. In our study, the participants in both Low- and High-Agency conditions needed to solve a fixed sequence of 10 problems. Before the problems, the students received preparatory tasks. The adaptive strategy used in the Low-Agency condition provided WE or ErrEx as preparatory tasks to students with lower prior knowledge; for the students with higher prior knowledge, the preparatory tasks could have been skipped, if the students demonstrated high performance on previous problems, or they received ErrEx or problems to solve. In the study, we compared this Low-Agency condition to the High-Agency condition, which enabled students to select preparatory learning activities on their own.

The students improved significantly from the pre-test to post-test in both groups. Even though Low-Agency students had higher means of post-test scores and CE scores than High-Agency students, we found no significant differences in post-test scores and CE scores between the Low- and High-Agency students; therefore Hypothesis 1 was not confirmed.

We were also interested in whether Low- and High-Agency had differential effects for students with different prior knowledge. The LPK students had significantly higher learning gain compared to HPK students. The HPK students improved significantly from pre- to post-test only in the Low-Agency condition. Unlike other studies, such as (Zimmerman 2008; Mitrovic 2001), in which advanced students performed better when given freedom and control to perform actions, we did not find any significant improvements for HPK students in the High-Agency condition. On the contrary, HPK students in the Low-Agency condition had higher post-test scores than the counterparts in the High-Agency condition with a large effect size (d = .95). Therefore Hypothesis 2a was rejected. There was no significant difference in learning gains between LPK students in the two conditions; therefore, Hypothesis 2b was also not confirmed. The Low-Agency condition was beneficial for both LPK/HPK students.

To determine why LPK and HPK students performed similarly on the post-test in the High-Agency condition, we proposed the 'step size' to infer whether students selected harder or simpler preparation tasks compared to the adaptive strategy used in the Low-Agency group. The results revealed that HPK students selected significantly more challenging learning activities in comparison to LPK students, but the selections made by LPK students were similar to the system selections. The findings suggest that the

adaptive strategy in the Low-Agency condition was efficient in selecting learning activities for both LPK and HPK students.

One limitation of the presented study is the small sample size. In order to achieve a power of 0.8, with the effect size d = 0.47 (i.e. the effect size on the post-test scores for the two conditions), it is necessary to have 138 participants (69 in each condition). Our study was conducted in an introductory database course taught at the University of Canterbury, which normally has about 200 enrolled students. The timing of the study coincided with assignments and lab tests in other courses the participants were taking, therefore, many students did not attend the scheduled labs, and some participants did not complete the study.

Another limitation of the study is in the difference in the types of preparatory activities the participants worked on in the two conditions. We found significant differences in the types of activities students from the High Agency group selected in comparison to the Low Agency group. In the Low Agency group, the adaptive strategy selected a high number of erroneous examples, while in the High Agency group the participant worked relatively equally on all types of learning activities. The preparatory activities affect learning differently; therefore our conclusions are not solely due to the difference in agency.

Several exciting research questions remain to be answered. We need to understand better the role of prior knowledge in learning from examples. All participants in our studies were familiar with SQL because they learned SQL in the lectures before participating in the studies. Even though our adaptive strategy is beneficial for students with different levels of prior knowledge, the results of our studies may be different with students who are fresh to the domain of SQL queries; it would be interesting to investigate the learning effect of using examples with this group of students.

Our adaptive strategy selects the learning activities for students based on their cognitive efficiency score on previous problems. The performance is computed from the student's score on the first submission of a problem. However, students may simply ask for feedback by submitting an empty solution initially. Therefore, in future work, the performance scores could be calculated more precisely by adding the time control as well as the feedback element that may affect students' learning during problem solving. Additionally, as we mentioned above, constraint-based SQL-Tutor models students by comparing students' solutions to ideal solutions provided by the teacher. A violated constraint represents an error, which translates to incomplete or incorrect knowledge. Our adaptive strategy is based on the number of violated and relevant constraints, but it does not consider how well the student knows each constraint. One of the future directions is to further enhance the adaptive strategy, in which the calculation of performance will take into account the complete student model rather than only violated/satisfied constraints from the most recent problem.

We proposed a High-Agency strategy that allowed students to select learning activities on their own. We found, like the Mitrovic and Martin (2003) study, that LPK students who selected learning activities themselves performed as well as LPK students who received learning activities adaptively. HPK students chose more challenging learning activities when they did not receive any instruction on the activity selection. Thus, they may not have been able to identify gaps or misconceptions in their knowledge, which could have helped them to select appropriate learning support on

their own. Furthermore, students who are attempting to self-regulate often face limitations in their own knowledge and skills, which can cause cognitive overload and decreased interest and persistence (Duffy and Azevedo 2015; Harley et al. 2015). Azevedo et al. (2016) demonstrated that deploying adaptive scaffolding and feedback in self-regulated learning produced better learning outcome compared to no scaffolding and feedback. Therefore, using adaptive scaffolding or feedback to guide students in High Agency would be an interesting topic for future research, particularly for HPK students.

Appendix

Pre-test

1. What clause of the SELECT statement allows tuples to be retrieved?

    a. SELECT
    b. FROM
    c. WHERE
    d. GROUP BY
    e. HAVING
    f. ORDER BY

2. What clause of the SELECT statement allows conditions to be specified on groups of tuples?

    a. SELECT
    b. FROM
    c. WHERE
    d. GROUP BY
    e. HAVING
    f. ORDER BY

3. What is the effect of the ORDER BY clause?

    a. Sorts tuples in a specified order
    b. Eliminates duplicate tuples
    c. Groups tuples
    d. Eliminates tuples that do not meet a specified condition

4. Which of the following clauses is mandatory in a nested query?

    a. ORDER BY

    b.   WHERE
    c.   SELECT
    d.   GROUP BY

5. Which predicate allows to check whether the value of an attribute is a member of the list of pre-specified values?

    a.   NOT EXISTS
    b.   MEMBER
    c.   EXISTS
    d.   IN

6. The attributes of tables specified in the outer query are accessible in the nested query. (True or False)

7. Which of the following should be used to fill the blank below to find the mean price?
      SELECT _____ FROM BOOK

    a.   MAX(price)
    b.   COUNT(price)
    c.   AVG(price)
    d.   SUM(price)

8. Two tables are given:

    STUDENT(StudNo, Name, Department)
    GRADES(StudNo, Course, Grade)

What is the effect of the following statement:

  SELECT name FROM student
  WHERE EXISTS (select * from grades
              where student.studno = grade.studno AND
              Course LIKE 'MATH___');

a. Find students who have passed no mathematics courses.
b. Find students who have passed no courses.
c. Find students who have passed some mathematics courses.
d. Find students who have passed at least one course.

Questions 9 and 11 are based on the following schema:

    DEPARTMENT(dname, dnumber, mgr, mgrstartdate)
    EMPLOYEE(ird, lname, minit, fname, bdate, address, sex, salary, supervisor, *dno)*
    DEPT_LOCATIONS(*dnumber,* dlocation)
    PROJECT(pname, pnumber, plocation, *dnum)*
    WORKS_ON(*eird, pno,* hours)
    DEPENDENT(*eird,* dependent_name, sex, bdate, relationship)

9. Find the first and last names of all employees who work in the Research department. (4 marks)
10. We need to find IRDs of employees who have no dependents. Is the following query correct? If not, specify the correct query. (3 marks)
     Select ird
     From employee
     Where 0 = (select count(*) from dependent where ird = eird);
11. We Need to Show for each Employee his/her IRD and how Many Projects he/she Works on. Is the Following Query Correct? If Not, Specify the Correct Query. (3 Marks)

     Select eird, count(*)
     from works_on;

Post-test

1. What clause of the SELECT statement allows the resulting table to be sorted?

    a. SELECT
    b. FROM
    c. WHERE
    d. GROUP BY
    e. HAVING
    f. ORDER BY

2. What clause of the SELECT statement allows conditions to be specified on tuples?

    a. SELECT
    b. FROM
    c. WHERE
    d. GROUP BY
    e. HAVING
    f. ORDER BY

3. What does Distinct do in an SQL query?

    a. Sorts the records in ascending order
    b. Returns only different values
    c. Sorts the result using a specified attribute
    d. Allows to have duplicated records in a database

4. Which aggregate function can be used to return the number of tuples?

    a. SUM
    b. COUNT
    c. MAX
    d. AVG

5. NOT IN allows you to specify a condition on an attribute checking that the value of the attribute does not appear in the enumerated set of values. (True or False)
6. The HAVING clause is applied to each group of tuples. (True or False)
7. We need to find the mean price of books for each genre. The query below is incorrect because:

SELECT GENRE, TITLE, AVG(PRICE) FROM BOOK

GROUP BY GENRE;

  a. TITLE should be added to the GROUP BY clause
  b. The GROUP BY clause is not needed
  c. TITLE should be removed from the SELECT clause
  d. PRICE should be added to the GROUP BY clause

8. Two tables are given:
     STUDENT(StudNo, Name, Department)
     GRADES(StudNo, Course, Grade)

What is the effect of the following statement:
   SELECT StudNo, Name FROM student
   WHERE StudNo IN (select StudNo from grades
   where Course = 'COSC265');

  e. Find students who have failed COSC265.
  f. Find students who have passed some courses.
  g. Find students who have taken COSC265.
  h. Find students who have passed COSC265.

Questions 9 and 11 are based on the following schema:

     DEPARTMENT(dname, dnumber, mgr, mgrstartdate)
     EMPLOYEE(ird, lname, minit, fname, bdate, address, sex, salary, supervisor, dno)
     DEPT_LOCATIONS(dnumber, dlocation)
     PROJECT(pname, pnumber, plocation, dnum)
     WORKS_ON(eird, pno, hours)
     DEPENDENT(eird, dependent_name, sex, bdate, relationship)

9. Select names of all departments located in Auckland. (4 marks)
10. We need to retrieve the IRDs of employees who work on any project controlled by the Planning department. Is the following query correct? If not, specify the correct query. (3 marks)
      Select distinct eird From works_on, project
      Where pno = pnumber and dnum = 'Planning';
11. We need to retrieve the IRD of each employee who works on more than two projects. Is the following query correct? If not, specify the correct query. (3 marks)

Select eird
From works_on
Group by eird
Having count(pno) > 2;

References

Atkinson, R. K., Derry, S. J., Renkl, A., & Wortham, D. (2000). Learning from examples: Instructional principles from the worked examples research. *Review of Educational Research, 70*(2), 181–214. https://doi.org/10.3102/00346543070002181.

Azevedo, R., Martin, S. A., Taub, M., Mudrick, N. V., Millar, G. C., & Grafsgaard, J. F. (2016). Are pedagogical agents' external regulation effective in fostering learning with intelligent tutoring systems? In A. Micarelli, J. Stamper, & K. Panourgia (Eds.), 13th International Conference on Intelligent Tutoring Systems, Zagreb, Croatia: 197–207: Springer.

Bandura, A. (1989). Human agency in social cognitive theory. *American Psychologist, 44*(9), 1175.

Bandura, A. (1991). Social cognitive theory of self-regulation. *Organizational Behavior and Human Decision Processes, 50*(2), 248–287. https://doi.org/10.1016/0749-5978(91)90022-L.

Booth, J. L., Lange, K. E., Koedinger, K. R., & Newton, K. J. (2013). Using example problems to improve student learning in algebra: Differentiating between correct and incorrect examples. *Learning and Instruction, 25*, 24–34. https://doi.org/10.1016/j.learninstruc.2012.11.002.

Calvert, S. L., Strong, B. L., & Gallagher, L. (2005). Control as an engagement feature for young children's attention to and learning of computer content. *American Behavioral Scientist, 48*(5), 578–589.

Chen, X., Mitrovic, A., & Mathews, M. (2016a). Do Erroneous Examples Improve Learning in Addition to Problem Solving and Worked Examples? In A. Micarelli, J. Stamper, & K. Panourgia (Eds.). *13th International Conference on Intelligent Tutoring Systems, Zagreb, Croatia:* 13–22: Springer.

Chen, X., Mitrovic, A., & Mathews, M. Do Novices and Advanced Students benefit from Erroneous Examples differently? In W. Chen, J. C. Yang, S. Murthy, S. L. Wong, & S. Iyer (Eds.), *24th International Conference on Computers in Education, Mumbai, India,* 2016b (pp. 46–53): Asia-Pacific Society for Computers in Education.

Chen, X., Mitrovic, A., & Mathews, M. (2017). Does Adaptive Provision of Learning Activities Improve Learning in SQL-Tutor? In E. André, R. Baker, X. Hu, M. M. T. Rodrigo, & B. du Boulay (Eds.), *18th International Conference on Artificial Intelligence in Education, Wuhan, China:* 476–479: Springer.

Chi, M. T., Leeuw, N., Chiu, M. H., & LaVancher, C. (1994). Eliciting self-explanations improves understanding. *Cognitive Science, 18*(3), 439–477. https://doi.org/10.1016/0364-0213(94)90016-7.

Clark, R. C., Nguyen, F., & Sweller, J. (2011). *Efficiency in learning: Evidence-based guidelines to manage cognitive load (illustrated ed.).* San Francisco, CA: John Wiley & Sons.

Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Hillsdale, N.J: Erlbaum.

Cordova, D. I., & Lepper, M. R. (1996). Intrinsic motivation and the process of learning: Beneficial effects of contextualization, personalization, and choice. *Journal of Educational Psychology, 88*(4), 715–730.

Duffy, M. C., & Azevedo, R. (2015). Motivation matters: Interactions between achievement goals and agent scaffolding for self-regulated learning within an intelligent tutoring system. *Computers in Human Behavior, 52*, 338–348.

Durkin, K., & Rittle-Johnson, B. (2012). The effectiveness of using incorrect examples to support learning about decimal magnitude. *Learning and Instruction, 22*(3), 206–214. https://doi.org/10.1016/j.learninstruc.2011.11.001.

Große, C. S., & Renkl, A. (2007). Finding and fixing errors in worked examples: Can this foster learning outcomes? *Learning and Instruction, 17*(6), 612–634. https://doi.org/10.1016/j.learninstruc.2007.09.008.

Harley, J. M., Bouchet, F., Hussain, M. S., Azevedo, R., & Calvo, R. (2015). A multi-componential analysis of emotions during complex learning with an intelligent multi-agent system. *Computers in Human Behavior, 48*, 615–625.

Hübscher, R., & Puntambekar, S. (2001). *Navigation support for learners in hypertext systems: Is more indeed better?* Paper presented at the 10th international conference on artificial intelligence in education, San Antonio, Texas, United States,

Kalyuga, S., Chandler, P., & Sweller, J. (1998). Levels of expertise and instructional design. *Human Factors, 40*(1), 1–17.

Kalyuga, S., Chandler, P., Tuovinen, J., & Sweller, J. (2001). When problem solving is superior to studying worked examples. *Journal of Educational Psychology, 93*(3), 579–588 doi:When problem solving is superior to studying worked examples.

Kalyuga, S., & Sweller, J. (2005). Rapid dynamic assessment of expertise to improve the efficiency of adaptive e-learning. *Educational Technology Research and Development, 53*(3), 83–93 doi:Rapid dynamic assessment of expertise to improve the efficiency of adaptive e-learning.

Katz, I., Assor, A., Kanat-Maymon, Y., & Bereby-Meyer, Y. (2006). Interest as a motivational resource: Feedback and gender matter, but interest makes the difference. *Social Psychology of Education, 9*(1), 27–42.

Kim, R. S., Weitz, R., Heffernan, N. T., & Krach, N. (2009). Tutored problem solving vs. "pure" worked examples. In N. Taatgen & H. van Rijn (Eds.), *31st annual conference of the cognitive science society* (pp. 3121–3126). Austin, TX: Cognitive Science Society.

Kirschner, P. A., Sweller, J., & Clark, R. E. (2006). Why minimal guidance during instruction does not work: An analysis of the failure of constructivist, discovery, problem-based, experiential, and inquiry-based teaching. *Educational Psychologist, 41*(2), 75–86. https://doi.org/10.1207/s15326985ep4102_1.

Mayer, R. E. (2004). Should there be a three-strikes rule against pure discovery learning? *American Psychologist, 59*(1), 14–19.

McLaren, B. M., & Isotani, S. (2011). When is it best to learn with all worked examples? In G. Biswas, S. Bull, J. Kay, & A. Mitrovic (Eds.), *15th International Conference on Artificial Intelligence in Education, Auckland, New Zealand:* 222–229): Springer. doi:https://doi.org/10.1007/978-3-642-21869-9_30.

McLaren, B. M., Lim, S. J., & Koedinger, K. R. (2008). When and how often should worked examples be given to students? New results and a summary of the current state of research. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *30th annual conference of the cognitive science society* (pp. 2176–2181). Austin, TX: Cognitive Science Society.

Mitrovic, A. (1998). Experiences in implementing constraint-based modeling in SQL-tutor. In B. P. Goettl, H. M. Halff, C. L. Redfield, & V. J. Shute (Eds.), 4th international conference on intelligent tutoring systems, San Antonio, Texas, USA, (pp. 414-423): Springer. doi:https://doi.org/10.1007/3-540-68716-5_47.

Mitrovic, A. (2001). Investigating students' self-assessment skills. In M. Bauer, P. J. Gmytrasiewicz, & J. Vassileva (Eds.), *8th International Conference on User Modeling, Sonthofen, Germany:* 247–250: Springer.

Mitrovic, A. (2003). An intelligent SQL tutor on the web. *International Journal of Artificial Intelligence in Education, 13*, 173–197.

Mitrovic, A., & Martin, B. (2000). Evaluating the effectiveness of feedback in SQL-Tutor. In *International Workshop on Advanced Learning Technologies, Palmerston North, New Zealand,* (pp. 143–144): IEEE. doi:https://doi.org/10.1109/IWALT.2000.890591.

Mitrovic, A., & Martin, B. (2002). Evaluating the effects of open student models on learning. In P. De Bra, P. Brusilovski, & R. Conejo (Eds.), *2nd International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems, Malaga, Spain,* (pp. 296–305): Springer. doi:https://doi.org/10.1007/3–540-47952-X_31.

Mitrovic, A., & Martin, B. (2003). Scaffolding and fading problem selection in SQL-tutor. In U. Hoppe, M. F. Verdejo, & J. Kay (Eds.), 11th international conference on artificial intelligence in education, Sydney, Australia: 479–481: IOS Press.

Mitrovic, A., & Ohlsson, S. (1999). Evaluation of a constraint-based tutor for a database language. *International Journal of Artificial Intelligence in Education, 10*, 238–256.

Najar, A. S., & Mitrovic, A. (2013). Examples and tutored problems: How can self-explanation make a difference to learning? In H. C. Lane, K. Yacef, J. Mostow, & P. Pavlik (Eds.), *16th International Conference on Artificial Intelligence in Education* (pp. 339–348). Memphis, TN, USA: Springer. https://doi.org/10.1007/978-3-642-39112-5_35.

Najar, A. S., & Mitrovic, A. (2014). Examples and tutored problems: Is alternating examples and problems the best instructional strategy? *Research and Practice in Technology Enhanced Learning, 9*(3), 439–459.

Najar, A. S., Mitrovic, A., & McLaren, B. M. Adaptive support versus alternating worked examples and tutored problems: Which leads to better learning? In V. Dimitrova, T. Kuflik, D. Chin, F. Ricci, P. Dolog, & G.-J. Houben (Eds.), 22nd International Conference on User Modeling, Adaption and Personalization, Aalborg, Denmark, 2014 (Vol. 8538, pp. 171-17a): Springer, Cham. doi:https://doi.org/10.1007/978-3-319-08786-3_15.

Nguyen, H., Harpstead, E., Wang, Y., & McLaren, B. M. (2018). Student agency and game-based learning: A study comparing low and high agency. In C. P. Rosé, R. Martínez-Maldonado, H. U. Hoppe, H. Ulrich, R.

Luckin, M. Mavrikis, et al. (Eds.), *19th international conference on artificial intelligence in education* (pp. 338–351). London, UK: Springer.

Ohlsson, S. (1994). Constraint-based student modeling. In J. E. Greer & G. I. McCalla (Eds.), *Student modelling: The key to individualized knowledge-based instruction* (pp. 167–189). Berlin, Heidelberg: Springer.

Renkl, A. (1997). Learning from worked-out examples: A study on individual differences. *Cognitive Science, 21*(1), 1–29. https://doi.org/10.1016/S0364-0213(99)80017-2.

Rowe, J. P., Shores, L. R., Mott, B. W., & Lester, J. C. (2011). Integrating learning, problem solving, and engagement in narrative-centered learning environments. *International Journal of Artificial Intelligence in Education, 21*(1–2), 115–133.

Salden, R. J., Koedinger, K. R., Renkl, A., Aleven, V., & McLaren, B. M. (2010). Accounting for beneficial effects of worked examples in tutored problem solving. *Educational Psychology Review, 22*(4), 379–392. https://doi.org/10.1007/s10648-010-9143-6.

Sawyer, R., Smith, A., Rowe, J., Azevedo, R., & Lester, J. (2017). Is more agency better? The impact of student agency on game-based learning. In A. Elisabeth, B. Ryan, H. Xiangen, R. M. M. T., & d. B. Benedict (Eds.), *18th International Conference on Artificial Intelligence in Education, Wuhan, China,* (pp. 335–346): Springer.

Schunk, D. H., & Zimmerman, B. J. (2007). Influencing children's self-efficacy and self-regulation of reading and writing through modeling. *Reading & Writing Quarterly, 23*(1), 7–25. https://doi.org/10.1080/10573560600837578.

Schwonke, R., Renkl, A., Krieg, C., Wittwer, J., Aleven, V., & Salden, R. (2009). The worked-example effect: Not an artefact of lousy control conditions. *Computers in Human Behavior, 25*(2), 258–266. https://doi.org/10.1016/j.chb.2008.12.011.

Schwonke, R., Wittwer, J., Aleven, V., Salden, R., Krieg, C., & Renkl, A. (2007). Can tutored problem solving benefit from faded worked-out examples. In S. Vosniadou, D. Kayser , & A. Protopapas (Eds.), *2nd European Cognitive Science Conference, Greece:* 59-64: Erlbaum.

Siegler, R. S. (2002). Microgenetic studies of self-explanation. In N. Granott & J. Parziale (Eds.), *Microdevelopment: Transition processes in development and learning* (pp. 31–58). New York, US: Cambridge University Press.

Siegler, R. S., & Chen, Z. (2008). Differentiation and integration: Guiding principles for analyzing cognitive change. *Developmental Science, 11*(4), 433–448. https://doi.org/10.1111/j.1467-7687.2008.00689.x.

Snow, E. L., Allen, L. K., Jacovina, M. E., & McNamara, D. S. (2015). Does agency matter?: Exploring the impact of controlled behaviors within a game-based environment. *Computers & Education, 82*, 378–392. https://doi.org/10.1016/j.compedu.2014.12.011.

Stark, R., Kopp, V., & Fischer, M. R. (2011). Case-based learning with worked examples in complex domains: Two experimental studies in undergraduate medical education. *Learning and Instruction, 21*(1), 22–33. https://doi.org/10.1016/j.learninstruc.2009.10.001.

Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science, 12*(2), 257–285. https://doi.org/10.1207/s15516709cog1202_4.

Sweller, J., Van Merrienboer, J. J., & Paas, F. G. W. C. (1998). Cognitive architecture and instructional design. *Educational Psychology Review, 10*(3), 251–296. https://doi.org/10.1023/A:1022193728205.

Tabbers, H. K., & de Koeijer, B. (2010). Learner control in animated multimedia instructions. *Instructional Science, 38*(5), 441–453. https://doi.org/10.1007/s11251-009-9119-4.

Tsovaltzi, D., McLaren, B. M., Melis, E., & Meyer, A.-K. (2012). Erroneous examples: Effects on learning fractions in a web-based setting. *International Journal of Technology Enhanced Learning, 4*(3–4), 191–230. https://doi.org/10.1504/IJTEL.2012.051583.

van den Broek, P., & Kendeou, P. (2008). Cognitive processes in comprehension of science texts: The role of co-activation in confronting misconceptions. *Applied Cognitive Psychology, 22*(3), 335–351. https://doi.org/10.1002/acp.1418.

van Gog, T. (2011). Effects of identical example–problem and problem–example pairs on learning. *Computers & Education, 57*(2), 1775–1779. https://doi.org/10.1016/j.compedu.2011.03.019.

Zimmerman, B. J. (2000). Self-efficacy: An essential motive to learn. *Contemporary Educational Psychology, 25*(1), 82–91. https://doi.org/10.1006/ceps.1999.1016.

Zimmerman, B. J. (2008). Investigating self-regulation and motivation: Historical background, methodological developments, and future prospects. *American Educational Research Journal, 45*(1), 166–183. https://doi.org/10.3102/0002831207312909.