

Pure constructivism for teaching a course in database management systems: Experience beats being told

Author

Estivill-Castro, Vladimir

Published

2016

Conference Title

ICERI2016: 9TH INTERNATIONAL CONFERENCE OF EDUCATION, RESEARCH AND INNOVATION

Version

Version of Record (VoR)

DOI

[10.21125/iceri.2016.0396](https://doi.org/10.21125/iceri.2016.0396)

Downloaded from

<http://hdl.handle.net/10072/339916>

Griffith Research Online

<https://research-repository.griffith.edu.au>

PURE CONSTRUCTIVISM FOR TEACHING A COURSE IN DATABASE MANAGEMENT SYSTEMS — EXPERIENCE BEATS BEING TOLD

Vladimir Estivill-Castro

Griffith University (AUSTRALIA)

Abstract

The scientific exploration of how people learn resulted in the theory named Constructivism. Such theory explains how students build their understanding of the world, assembling concepts into knowledge. The primary process consists of experiencing something and then reflecting on those experiences. This paper describes our experience in teaching and undergraduate course (third year) in database management system using laboratories, simulators, problem-solving tutorial session and of course lectures and recorded lectures (on-line delivery). We identified several concepts and tracked down the level of novelty of the idea and also the experience students had with it though the different learning activities in the course. We repeated our approach for four editions of the course. The evidence collected during these four versions demonstrate that concrete experience is invaluable even for the lowest categories of the cognitive domain of Bloom's taxonomy. We introduced other elements, such as authentic assessment, but it seems particularly clear that the concrete experience is the most relevant component to enhance students' progress in the ladder described by Bloom's taxonomy.

Keywords: Constructivism, Concrete experience, Bloom taxonomy, Experiences in Post-graduate education, Courses, Tutorials and Labs.

1 INTRODUCTION

The scientific exploration of how people learn resulted in the theory named Constructivism [7, 3, 11]. Such theory explains how students build their interpretation and understanding of the world, assembling concepts into knowledge. The process consists of experiencing something and then reflecting on those experiences [8]. We describe our experience in teaching an undergraduate course (third year) in databases using laboratories, simulators, problem-solving tutorial session and of course lectures and recorded lectures (on-line delivery). We identified several concepts and tracked down the level of novelty of the idea and also the experience students had with it though the different learning activities in the course. We were fortunate enough that only in this offering students elected several paths for their learning, in particular, some attended lectures while other only used the lecture recordings. Nevertheless, lectures included active techniques (some small experiments).

While Constructivism is typically applied to science delivery as its fundamental principles, consist of analysing the current understanding and then designing an experiment to either revise the views or confirm them, we wanted to adapt this to a more technologically oriented and applied information technology content. In fact, our aim is to review the conceptual understanding that students have of how a database works. Students complete a first-year course using a database in the design of a small Web application. In the second year, students become familiar with query languages (in particular SQL) and the notation for constructing Entity-Relationship diagram. The challenge in the third year course is for students to understand what is under-the-hood. The aim of the course proposed a more advanced understanding. Such understanding is similar to the contrast made in mathematics education [10] between the instrumental teaching of mathematics and the relational teaching of mathematics. We hope students can not only apply and use the technology associated with databases, but they should be able to understand why it is being developed or behaves in certain ways. The aim is for higher levels of Bloom's taxonomy [1] as we hope the student can analyse issues and produce physical designs of databases. Such designs will invariably represent trade-offs in the performance of some aspects at the expense of others. Thus, justifications and explanations will back up the ideas that students combine to synthesise prototypes and products.

A crucial learning objective if that software engineering graduates shall use their understanding of the working of a database server to produce a highly efficient design of database system that performs well for a diversity of needs and workloads. The challenge today is that, in an undergraduate course, it is hard to recreate the environment in which complex database servers operate. For example, to setup

a scenario where thousands of users are concurrently executing transactions is essentially impossible. And if we did, it would be delicate to create the experiment for students to observe different effects caused by the setting of parameters. Thus, the challenge is to create experiences and concrete activities to apply to the learning objectives of the course, and that could be conducted in the simple setting of laboratories, workshops or lectures and within the 100-minutes duration of such sessions.

Parallel with this, we created the assessment instruments that enabled us to rank the level of attainment on a concept as per Bloom's taxonomy [1]. We anticipated that the more experimentation and reflection was provided on a concept, the higher the level students would demonstrate. For example, at the lowest level, we expect students to remember the definition of physical database design, while in the next level up students would be able to distinguish an ER-diagram at the conceptual design stage from the physical design stage. A further level would constitute the manipulation of an ER-diagram to achieve a suitable design and anticipate the SQL-tables and produce corresponding CREATE statements to define such structures. Would students be capable of analysing their ER-diagrams and anticipate properties of the physical design and the implementation, for example, trade-offs of normalisation and de-normalization? Finally, would they be able to create their examples that illustrate properties of their designs?

We repeated our approach for three editions of the course. The evidence collected during these four versions demonstrate that concrete experience is invaluable even for the lowest categories of the cognitive domain of Bloom's taxonomy [1]. We introduced other elements, like authentic assessment, but it seems particularly clear that the concrete experience is the most relevant component. Students who engaged in concrete activities were able to progress.

2 THE COURSE

The course 3612ICT Database Management Systems was developed in 2012 as a 20CP course (double the credit point value of a regular course for the programs it belongs to). In Griffith University one credit point is considered to equate to approximately 15 hours of work. Thus this is a course of about 300 hours of work. It is delivered face-to-face with 24 lectures of 2 hours per week, ten laboratory sessions of 2 hours and ten seminar sessions of 2 hours. The remaining 212 hours are self-study, but lecturers and course conveners hold office hours for consultation. The course belongs to the Bachelor of Software Engineering as a core course, but for the Bachelor of Information Technology and the Bachelor of Multimedia, it is an elective course. Nevertheless, students in the Bachelor of Information Technology have all pre-requisites as core courses in their program, so it is common they chose this course. Since 2012 the course also shares the lectures and seminar with a graduate version of the course for a Masters by course-work program. Master students have slightly more elaborate tasks for the laboratories. The course is delivered at two campuses (Nathan and Gold Coast) by a pool of four academics. Although there have been over ten offerings since 2012, the data and reflection reported in this paper are over four offerings, two concurrent offerings at the two campuses in 2014, one at the Gold Coast in 2015 and one at Nathan campus in 2016. The prerequisites are two units. A 10CP course named 2002ICT Database Design which focuses on using a database with SQL and constructing Entity-Relationship diagrams. A 10CP course on object-oriented programming that provides programming skills in Java.

Over the 4 units, we have that ninety-two percent of students are male and 8% female. Sixty percent are 21-22 years old. Twenty-seven percent 23-25 years old, 7% are 26-30 years-old. Eighty-five percent claim English as their first language, 15% indicate another language. Sixteen percent reached their enrolment in the degree by not regular entry admissions. Less than 2% have a GPA of 7 while 36% have a GPA of "4 to no less than 5" (this is just passing), while 17% do not have a passing GPA.

The laboratories are all conducted using DERBY. The rationale is that DERBY is the implementation of reference for Java Database Connectivity (JDBC), and there are many examples and tutorials on-line of clients programmed in Java. DERBY is very flexible, and can easily install on Windows, MacOS or Linux, and the students can run it with administrator privileges on their computers or the university laboratories. DERBY can be used to teach indexes, index-only queries, privileges, stored procedures transactions and all sorts of client-server applications (2-tier, 3-tier). However, the students are free to choose the software for their project. A very common choice is mySQL, but Postgres is also common.

Officially, students are not to be penalised for missing lectures, or any other session on campus. Officially, students are encouraged to attend as many classes as possible. However, in those four offerings in question we recorded attendance by requesting the completion of activities in class (while

those activities in lectures had no value towards the course grade participation in the seminars awarded 1% per seminar of the final grade). Tutors recorded laboratory attendance, but again, it does not contribute to the grade, although each laboratory contributes 2% of the mark (1% for completion of the laboratory tasks and 1% for a reflective report on the laboratory tasks). All items of assessment are to be performed individually. In all the four offerings under discussion, lectures are recorded and made available soon after. The recording only includes the display of the host computer and the voice of the lecturer (and therefore, discussion on a whiteboard or any activity in the session is not recorded). Participation in seminars is registered by students writing their student identification number and their full name on a piece of paper. The facilitator of the seminar request for a show of hands for those who would like to participate and volunteer alternative answers to an exercise. A participant student is not assessed for the correctness of the contribution, as long it is a genuine attempt, and for all contributing students, the facilitator collects the piece of paper. Facilitators may solicit participation from students present who have not contributed and may give preference to students who have not participated in that session over several interventions by the same student. Seminars are facilitated problem-solving sessions. The exercises or problems to be discussed are released one week before the session and students attending are expected to bring a notebook with solutions attempted. Tutors record the amount of preparation work performed by students just for statistical purposes like the charts presented in this paper. Students should use the discussion by peers under the guidance of the facilitator to complete the exercises or work out solutions. Students are advised that they are to construct their individual submission which is required one week after the session. Submission of a seminar exercises can amount up to 1% per seminar of the course grade.

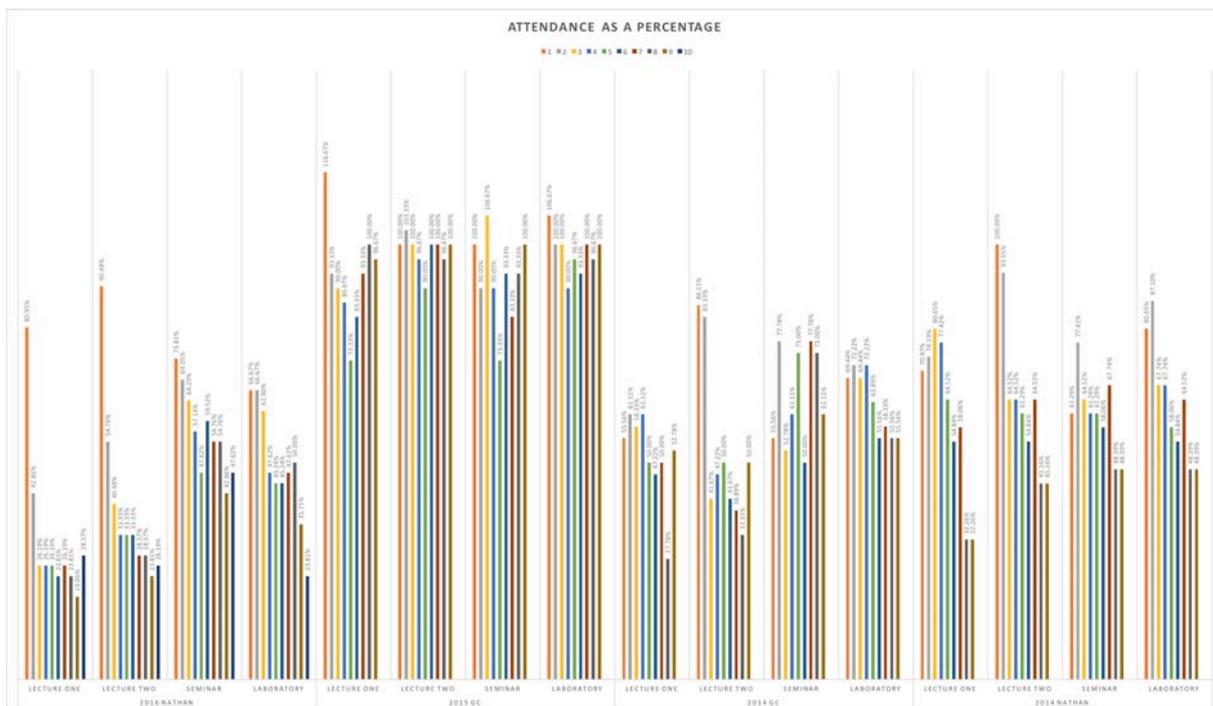


Figure 1 Students attendance to the two lectures, seminar and laboratory sessions for four offerings of the same course

Figure 1 shows the percentage of attendance over ten weeks of delivery of the above-mentioned course offerings. In all these courses offering, the same lecturer was in charge of lectures, and in charge of managing tutors. For both of the 2014 offerings and the 2015 offerings, the lecturer was also in charge of the seminars, while in 2016, the lecturer was in charge for 5 of the seminars (tutors took over for the last five seminars). In all offerings, the laboratories were managed by casual tutors.

In Figure 1, the ratio is relative to the number of students that completed the course and were awarded a mark. Typically, about 30% more students enrol, but they drop the course by week 4 of the semester. Thus, on occasion, there is a figure above 100%. But the reader may notice that most of the time this is not the case. Both offerings at Nathan campus show a decline in engagement as weeks go by. At the Gold Coast campus, this is less marked, and for the 2015 offering, the Gold Coast student show high engagement across the four activities (the two lectures, the seminar and the laboratory). All

of them receive above 75% attendance above all weeks and averages of 90% attendance across all activities. By contrast, the 2016 Nathan offering displays lectures with an average of 25% participation, seminars and laboratories are higher (perhaps because of the incentive of direct impact on the course grade), but not above 50%. In 2014 attendance at the Gold Coast is slightly higher than at Nathan, but the distinction is not significant. In the earlier offering of the course, the two lectures, the seminar and the laboratory session were scattered over the week. Attendance was even lower. These four offerings have a lecture at 8 AM and the second lecture 11 AM after or noon, while the seminar was scheduled for the afternoon. Laboratories are scheduled concurrently because the class is divided into two groups, and two tutors are staffed for this activity. One can see a slightly higher attendance for the second later lecture over the early lecture.

Students at the Gold Coast are on average better performers than at Nathan for the corresponding degrees (their entry scores from high school are higher for the Gold Coast than for Nathan). But the high variability of attendance motivated to research further how the activities and learning opportunities are being used by the students. We want to investigate how this diverse levels of participation reflect on the learning, and what instruments are providing the learning.

It should be noted that although there is a prescribed textbook, in all offerings, about 10% of the students purchase a new or used textbook, and this was across all offerings. Regarding the utilisation of the recorded lectures, there is perhaps more alarming or puzzling data.

The other common element between all four offerings refers to the use of recorded lectures or recorded seminars.

- The average number of students who reviewed the early lecture is two students, and they spent an average of 10 minutes.
- The average number of students who reviewed the later lecture is also two, and they spent an average of 10 minutes.
- The average number of students who reviewed the seminar is five, and they spent an average of five minutes.

For the four units, these values include the 12 weeks for lectures and additional two seminars where a course revision and a final exam revision are performed. The data implies that 4% of the student body use the recorded lectures and 96% do not. Similarly, 10% of the student body use the recorded workshops and 90% do not. Those who use the recorded lectures use it for less than 5% of the recorded duration of the lecture. This data means that those who use the recorded workshops use them for less than 10% of the recorded workshop.

Similar data was collected about the utilisation of the consultation hours. Remarkably, the data is even less encouraging. The number of visits over the entire semester is less than 4. Enquiries by e-mail are an average of 10 a semester across the four offerings. Cases of plagiarism that were referred for integrity management are an average of 3 across all offerings.

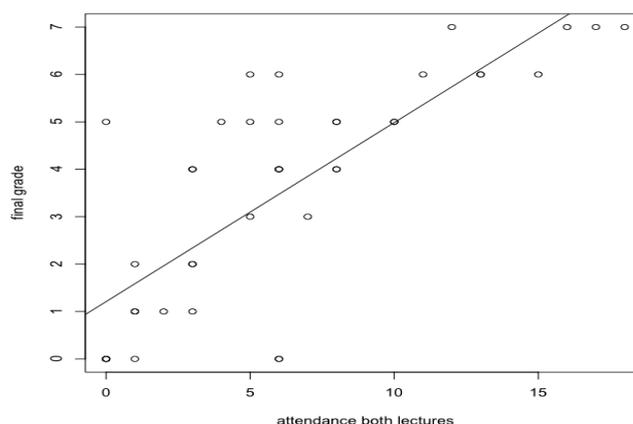


Figure 2 Correlation between attendance to lectures and final grade

Figure 3 shows the marks achieved in the four offerings over analysis. Perhaps it is not surprising that attendance to laboratories or attendance to seminars is highly correlated with the grade achieved.

Since participation in both of these learning activities usually, results in a submission that contributes to the final grade. However, it is perhaps more interesting that attendance to lectures is highly correlated with a p-value that rejects the null hypothesis (typical p-values are of the order of 2.9×10^{-10}).

Figure 2 shows Pearson's product-moment correlation (0.79) between final grade and attendance to both lectures for the Nathan offering 2016. It is not possible to say exactly what is cause and effect: all forms of attendance (even to the early lecture or to the later lecture) are also correlated. In general, students who engage attend all activities, but we do have the student who cannot attend one activity for work commitments and attends all others. Simply put, it is clear that engagement demonstrated by attendance is highly correlated with succeeding in the course. Interestingly, despite showing these statistics to students, as illustrated in Figure 1, participation can be very low.

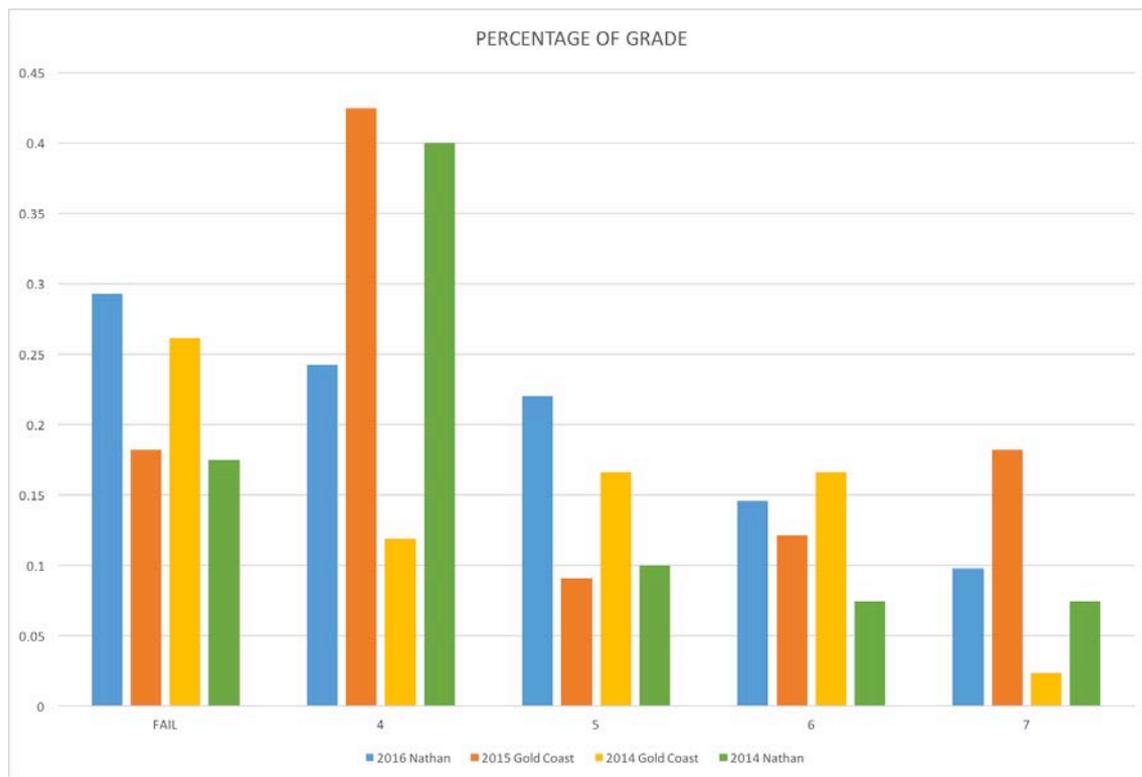


Figure 3 Distribution of grades for the four offerings of the course.

We also conduct a review of pre-requisite material in the first two weeks of the course. The first workshop consists of solving a final exam of the Database Design course. The conditions are not of an exam setting (students solve this outside class and work on it for one week before the seminar and then they have an additional week to complete it). Nevertheless, performance at this first workshop is also highly correlated with the final grade achieved in this course. Again, we cannot speculate as perhaps the causing variable is the maturity and self-regulation attitudes of the student who determine performance in the pre-requisite course, performance in prerequisite exercises, and grade in the current course. Nevertheless, in the opinion of this author, it is surprising that pre-requisites are no longer enforced as University policy and secondly, as we mentioned before, attendance in courses where face-to-face is the delivery mode is also not compulsory.

3 WHAT DO THE STUDENTS USE?

The context described above motivates us to explore and investigate further what is useful and what enables students to progress to higher levels of the Bloom taxonomy with specific concepts. Figure 4 is a high-level description of the concepts involved. The type of inquiry we proposed here consisted on having options in the final exam. For example, there would be a question on transactions, but the question would have two variants, one with client-side transactions (via JDBC) and the other as server side transactions via stored procedures. Lectures discuss both types of transactions and additional

theory on transaction management, but only half the class would have a practical, concrete exercise on client-side transactions while the other half would have a practical exercise on server-side transactions. We kept track then of which students selected which question, and to what extent they could construct a more advanced solution in that they could perform an analysis or a critique of a specific ER design.

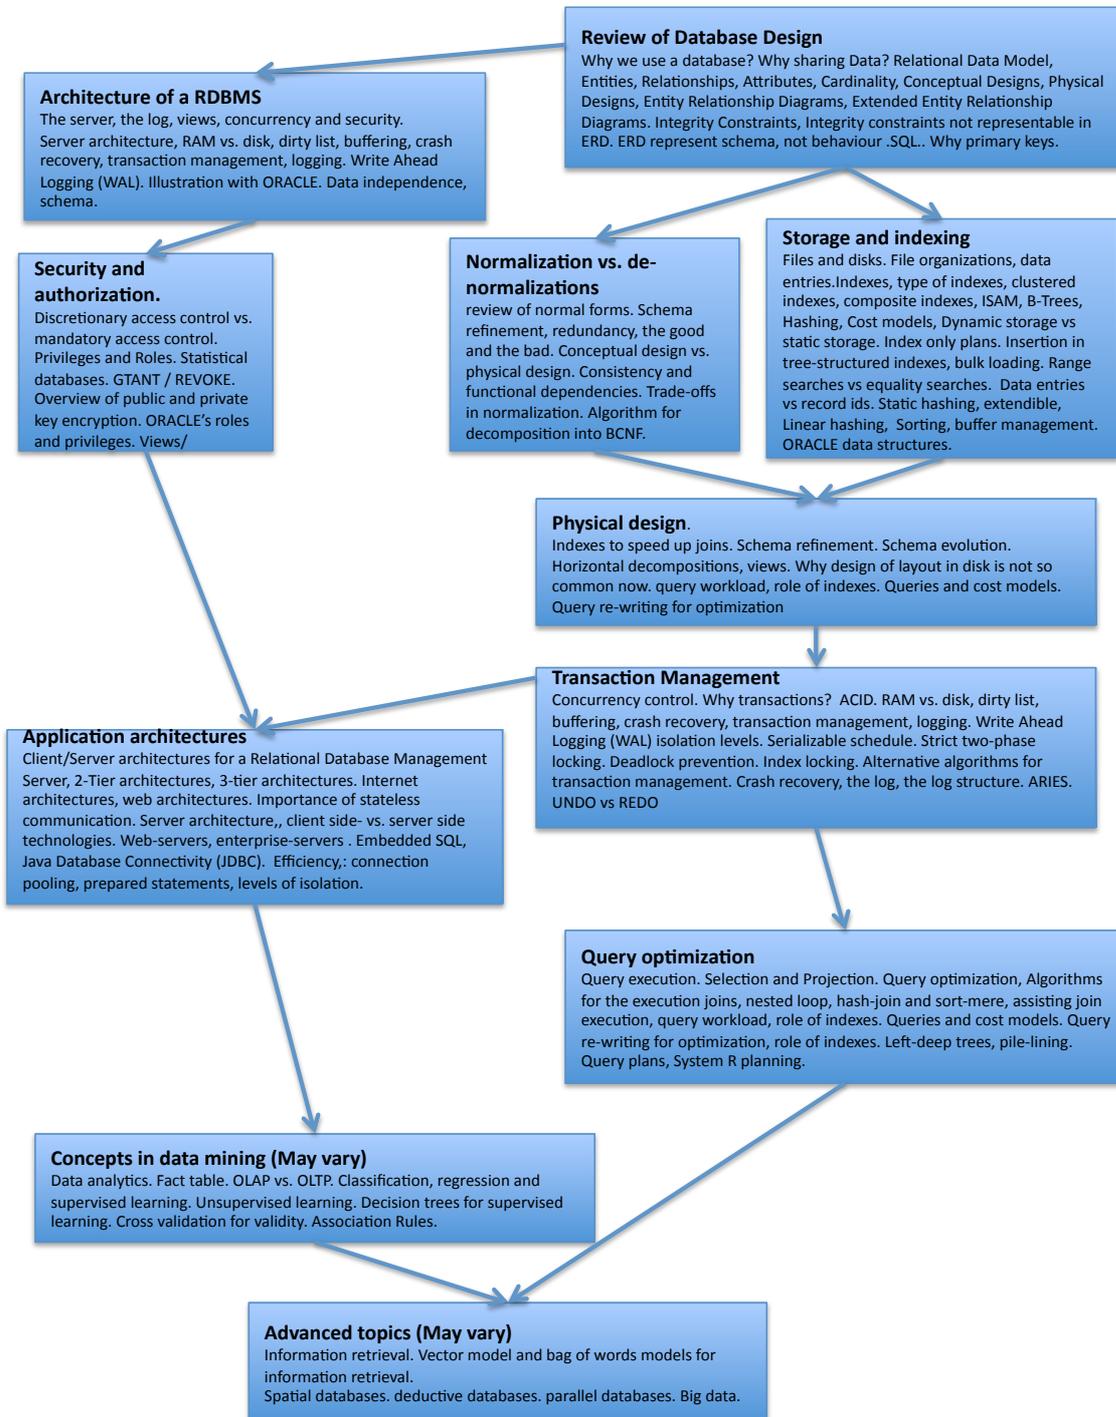


Figure 4 Concept map for the unit under analysis.

For each concept, we pose two questions; however, the questions have a Bloom taxonomy level of sophistication. The first level was the recall level, and the question just required to duplicate a definition or match and order some concepts and potentially recognise and order some steps in a process. The next part of the question demanded to give examples or to classify instances into

categories, to correspond with the level named comprehension. The level of application was evaluated by a third section of the question that required to prepare a solution, apply some method or compute an answer, it may include the demonstration of the effects of an algorithm, or to interpret a log of transactions. For the level of analysis, students were required to discover potential areas improvement in an answer, or to contrast advantages and disadvantages, suggesting a situation where the conditions may change and demand new interpretations, perhaps producing an illustration of scenario which would require a reverse effect. The fifth level of synthesis requested to create new designs, to explain and justify the solutions created and to arrange requirements that may not be possible to satisfy. For the level of evaluation, we measured the detail of justifications for designs or the ability to appraise the context of a scenario and to relate concepts together. The clarity of explanations and the particulars of the evaluation was important here.

The concept map of Figure 4 helps to construct questions for the recall level. Most questions re of the form, provide the definition of (a concept). For instance, "in ER-modelling what is an entity?" However, for comprehension, we need to be more creative although questions that create contrast are typical "explain the difference between conceptual design and physical design" or "explain the difference between client-side transaction and server-side transaction". Students are warned that they cannot answer with tautologies. That is, no marks are awarded if the answer does not provide additional information that explains the difference and creates contrast. For example, a question that asks for the difference between the "UNDO" and the "RE-DO" on crash-recovery in transactions, cannot be answered by "Undo is when we UNDO the work and re-do is when we do it again" (although it is somewhat disappointing that such answers do appear in exams).

The types of exercises we constructed were on several topics.

- 1 Database Design (Conceptual vs. Physical).
- 2 Physical design trade-offs.
- 3 Transactions design and construction (server-side vs client-side)
- 4 Data Modelling. Object-Oriented Modelling vs. Entity Relation Modelling
- 5 Machine learning (supervised vs. unsupervised learning).
- 6 Security (database privilege administration vs. application architecture administration).
- 7 Design constraints (Data definition constraints vs. other integrity constraints)
- 8 Query evaluation plans and join evaluations trade-offs (algorithms for joins, normalisation vs. de-normalization).

In this list, each item has two concepts, and although both were covered in the lectures, one of the aspects was covered by a concrete exercise in seminars and laboratories for approximately half of the student cohort.

We evaluated first if the students select the question that elaborates more on the aspect they had a concrete, practical exercise. Naturally, as expected, students chose the variant of the problem they had more exposure through the practical exercise, although not every student does. Second, we analysed the correlation between the depth of self-reflection and grade achieved in the laboratory exercise and the depth in the Bloom taxonomy that students achieve. Again, there is also significant correlation. For example, those students that observed the use of hash-indexes and experimented with them deeply in the laboratories could also explain better in the exam question regarding linear hashing and extendible hashing. Students who practiced the distinction of join algorithms and the differences between which table to sue as the inner table could argue better the exam question that demanded a join between two tables.

The theory of CONSTRUCTIVISM would predict that the more a student confronted and practised with a concept, the higher the level of understanding and command of the concept. We analysed as binary variables whether students belonged to the cohort that had a practical exercise with a concept and whether students selected the questions for the concepts. We used the Pearson correlation coefficient again because, for two binary variables, this returns the phi coefficient. In all cases, we had positive correlation coefficient. Many times we had a p-value that showed statistical significance; however, in some few cases, the p-value was not small enough to imply significance. One of such situations was the practice with extended ER-diagrams and Object-oriented and the selection of questions regarding the conversion of Extended ER-models to an implementation with relational databases, followed by a discussion of advantages and disadvantages of the resulting physical design. However, practice with

the server-side transaction or client-server-side transactions is a significant predictor of the type of question students will choose. Similarly, practice with particular types of query plan formulations is a predictor for the questions students will select in the final exam for the concept of query plans.

We observed that there is a positive correlation between the mark the student received in the reflective exercise regarding a concept and the level achieved in the Bloom taxonomy. However, the feedback received did not seem to make much difference. That is, students who did poorly in a concrete exercise in a concept do choose questions related to what they have practised more before, but their performance is not significantly better. This relation probably confirms the main point of this paper, if students do find the answer through working the problems themselves to a standard that improves their competence, they retain such expertise, and thus selected the appropriate question for their expertise in the final exam. If they did well already, they do well again. If they did not do so well, then being told the answer through the feedback does not change things too much. It is clear that students can reach the knowledge/recall and the comprehension levels of the Bloom taxonomy if they complete the concrete exercises for the corresponding concepts successfully. Students who do not complete such activities can hardly produce correct definitions or are easily confused with multiple-choice questions and matching questions as well as being unable to discuss differences and commonalities with similar concepts.

We tried to identify why is it that some students chose another problem rather than the one they had more concrete practice. We examined if a student's selection of the different alternative to the one they had more practice had any correlation with the final mark in the exam. While this author's intuition is that weaker students cannot even tell if the question is something they should answer, the numbers are small, and the statistical analysis reveals just a slight correlation but the p-values are not extreme enough to make a claim of this. We believe that further data and larger numbers eventually will confirm this. What we are still puzzled is why there are a few students that are in the group of a passing mark or a credit grade, and they still seem to choose the questions for which they have had less concrete practice. We can only speculate that this is due to many other factors, like the pressure of the exam setting and students wrong assessment of how long would an exercise require. Their experience may guide them to an exercise they have no previous experience over the one they should work better given they are reworking a previous problem.

4 STUDENT EVALUATIONS

Two standard student surveys of Griffith University provide feedback for all four offerings. The student evaluation of teaching (SET) focuses on the delivery of lectures, laboratories and the seminar, while the Student Evaluation of Course (SEC) aims at collecting feedback on the unit itself. Evaluations of teaching are varied, and for the common element of the same lecturer for all four offerings evaluations are usually very positive, with all offering having at least a nomination for best lecturer and the 2015 offering at the Gold Coast scoring the maximum possible score. But in 2016 scores were low. Similarly, the evaluations of the course. They are typically aligned with the distribution of marks (see Figure 3) and show that student feedback has a bias towards the expectation of grade [6]. The literature on student evaluations is extensive, but it is accepted that SET and SEC results are influenced by external factors. Those factors include the length of the course, the size of the class, and the gender and age of the teacher [2]. Many other factors that have no direct connection to teaching effectiveness but such factors do play a role. The one identified here is students' expectation of what their final grade will be [9].

That is, the SEC score is favourable when there is a low failing rate and unfavourable when there is a high failure rate. The most common element in the student surveys is the perception that workload is extremely high for the subject. This feedback is precisely why despite the evidence that concrete learning activities result in better learning of the concept, such learning activities when they happen in the classroom they are not graded. We face a dilemma. Students seem to engage [4] (attend) a session if the activity in that session counts in some way to the final grade, but they also expect a workload for a 20CP unit equivalent to 10CP units they have completed in the first two years.

As a result of this, the faculty has concluded that student expectations [5] of a course are commensurate with their previous experience and that third-year courses for 20CP are to be eliminated. The unit will, therefore, be offered for the last time in 2017, and all courses for 20CP are eradicated.

5 CONCLUSIONS

We have collected data of four offerings of a demanding course in database technology and administration (it is designed to be double the load of most of the courses the students complete in the corresponding program). Our data suggests that four of the four offerings under analysis, students' concrete experience with some concepts is invaluable, especially for the lowest categories of the cognitive domain of Bloom's taxonomy. That is, students can remember a definition, or produce an example and indicate what technique to apply when they have had concrete exposure to the concept. Despite introducing small concrete experiences in lectures, these remain unassessed, and attendance to lectures is the weakest engagement. Other factors play a role, like command of the pre-requisite knowledge. The course also has an element of authentic assessment, which consists of a realistic project of developing a non-trivial client-server application where a database backend is the main component of the solution. Progress in the project is complemented with formative assessment by setting up three milestones (conceptual design, physical design, operational prototype). The project seems to have no impact. In fact, completion of the laboratory exercises is the most determinant factor in completing the project. This relationship between laboratories and project success is another reason why completion of the laboratory exercises has a strong correlation with success in the course. Therefore, concrete experiences have emerged as the most relevant component so that students' progress in the ladder described by Bloom's taxonomy.

REFERENCES

- [1] B.S. Bloom, editor. *Taxonomy of Educational Objectives: Handbook I; Cognitive Domain*. David McKay Co, New York, 1956.
- [2] A. and Boring. Student evaluations of teaching (mostly) do not measure teaching effectiveness. *ScienceOpen Research*, January 7th 2016.
- [3] R. Driver, H. Asoko, J. Leach, P. Scott, and E. Mortimer. Constructing scientific knowledge in the classroom. *Educational Researcher*, 23(7):5–12, 1994.
- [4] J. A. Fredericks, P. C. Blumfield, and A. H. Paris. School engagement: potential of the concept, state of the evidence. *Review of Educational Research*, 74(1):59–110, 2004.
- [5] B. Glesnaer Fines. The impact of expectations on teaching and learning. *Gonzaga Law Review*, 38(1):89–128, 2002.
- [6] B.W. Griffin. Grading leniency, grade discrepancy, and student ratings of instruction. *Contemporary Educational Psychology*, 29(4):410–425, 2004.
- [7] B. Krueger, J. J. Loughran, and R. Duit. Chapter 13 : Constructivism. In J. Wallace and W. Louden, editors, *Dilemmas of science teaching*, pages 191–204, London, 2005. Routledge.
- [8] J. J. Mintzes and J. H. Wandersee. Chapter 2 : Reform and innovation in science teaching : a human constructivist view. In J. J. Mintzes, J. H. Wandersee, and J. Novak, editors, *Teaching science for understanding: a human constructivist view*, pages 29–58, San Diego, CA, 1998. Academic Press.
- [9] R. Remedios and D. A. Lieberman. I liked your course because you taught me well: The influence of grades, workload, expectations and goals on students' evaluations of teaching. *British Educational Research Journal*, 34:91–115, 2008.
- [10] [R.R Skemp. Relational understanding and instrumental understanding. *Mathematics Teaching*, 77:20–26, 1976.
- [11] R. Tyler. Chapter 2: Constructivist and socio-cultural views of teaching and learning. In G. Venville and V. Dawson, editors, *The art of teaching science — for middle and secondary school (2nd edition)*, pages 23–40, Cross Nest NSW 2065, Australia, 2012. Allen Unwin, Griffith Press.