

Insights from Student Solutions to MongoDB Homework Problems

Anonymous Author(s)

Abstract

We analyze submissions for homework assignments of 527 students in an upper-level database course offered at a large U.S. public university. The ability to query databases is becoming a crucial skill for technology professionals and academics. Although we observe a large demand for teaching database skills, there is little research on database education. Also, despite the industry's continued demand for NoSQL databases, we have virtually no research on the matter of how students learn NoSQL query languages, such as MongoDB. In this paper, we offer an in-depth analysis of errors committed by student working on MongoDB homework assignments over the course of two semesters. We show that as students use more MongoDB operators, they make more JavaScript-related errors. Additionally, when students face a new functionality of MongoDB operators, such as \$group operator, they usually take time to understand it but do not make the same errors again in later problems. Finally, our analysis suggests that students struggle with advanced concepts in comparable duration and effort. Our results suggest to instructors that in order to improve their curriculum may set more time and effort for the discussed topics in our paper.

CCS Concepts

• **Applied computing** → **Education**; • **Social and professional topics** → **Computer science education**; • **Information systems** → **Information retrieval**; **Query representation**.

Keywords

mongodb, database education, online assessment

ACM Reference Format:

Anonymous Author(s). 2020. Insights from Student Solutions to MongoDB Homework Problems. In *Woodstock '18: ACM Symposium on Neural Gaze Detection*, June 03–05, 2018, Woodstock, NY. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 Introduction

The evolution of big data exposed the limitations of relational (SQL) databases. Internet giants, such as Google and Amazon, developed and used many custom-built databases to work around the shortcomings of SQL databases. MongoDB, a cross-platform document-oriented database that uses JSON-like documents, was introduced in 2007 as part of this NoSQL (Not Only SQL) revolution. Since its inception, MongoDB has gained a great reputation and has

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Woodstock '18, June 03–05, 2018, Woodstock, NY

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

become one of the most used and wanted database query languages [15].

The flexibility of MongoDB's semi-structured documents (JavaScript Object Notation, JSON) and its ability to effectively scale to distributed collections of documents made several large companies shift from SQL to MongoDB. MongoDB provides higher scalability and availability than its SQL-based counterparts [2]. It also allows users to develop user-defined functions written in JavaScript as well as write Map-Reduce programs. Map-reduce infrastructure enables users to write distributed aggregate computations over large, highly-distributed, volumes of data [1].

In response to the new interest in NoSQL, many universities and educational organization have adopted curriculum that includes NoSQL database management systems. For instance, as discussed in Section 2, Mohan showcase their experience with teaching NoSQL paradigms to undergraduates at Rose-Hulman Institute of Technology [12].

Addressing the growing interest of students and instructors in MongoDB, in this paper, we give insight about students' experience when learning MongoDB using a quantitative approach. We classified 76,168 submissions made by 527 students over Fall 2019 and Spring 2020 semesters while working on homework problems at a big Midwestern university in the United States. This contributes to the growing literature about analyzing and understanding students' mistakes when learning a new database language. In this paper, We answer the following research questions:

- What concepts are difficult for students when learning MongoDB?
- What common errors do students make when first learning to query a MongoDB database?

2 Literature Review

Unlike Object-oriented and imperative programming, there is not a lot of research about teaching students to program in database query languages. Though there are no existing studies on the difficulties encountered by students while learning to use MongoDB, there are multiple reports of instructors including MongoDB into their curriculum in university courses [9–12].

Guo et al. integrated MongoDB into security labs [10]. These efforts mainly focus on how to order topics related to database security and how to introduce some of the security drawbacks of MongoDB's implementation. Lei et al. [11], use MongoDB as their NoSQL database Management system in their labs to teach the drawbacks of NoSQL databases. Mohan reported experiences of a database education curriculum that incorporated NoSQL [12]. In Mohan's work, students were exposed to several NoSQL paradigms and had a set of projects, lab and research assignments to complete using the knowledge they gained during the course. The course received positive feedback from students and industry big data engineers for its addressing of NoSQL paradigms. Some other NoSQL

databases have also been incorporated in university curricula. For example, Fowler et al. reported their experience in two database courses with teaching CouchDB, a NoSQL data management system that uses JavaScript as its query language [9]. They mainly focused on measuring students' improvement of understanding NoSQL systems. Although, the previous papers include MongoDB in their studies, but they do not study novices' experiences or difficulties when learning MongoDB.

Previous work in SQL education has looked at thousands of student submission to both homework and tests to understand what types of queries are most difficult for students to write, and what types of errors students make the most [5, 6, 8, 16]. In this paper, we follow their lead, giving a quantitative insight from students' solutions of MongoDB homework questions in order to find common difficulties among novices.

3 Intro to MongoDB

MongoDB [3] is a document-oriented NoSQL database that stores JSON-like data in documents with dynamic schema, so it can store flexible values without fixing the number of fields or type of fields. We will show examples of MongoDB objects and some code snippets that exhibits the basic syntax of MongoDB and the corresponding queries in SQL with similar database schema. The '_id' attribute is an indexed attribute that every object must have, and it acts as the object's primary key. We will show two instances of objects that might appear in collection titled 'Movies', to demonstrate the flexibility in MongoDB, avoiding SQL's rigid relational structure and the need that people normally have to translate objects to relational tables. One object in such a collection might look like:

```
{
  "movie_id": 1,
  "movie_name": "The Imitation Game",
  "release_year": 2014,
  "country": "USA",
  "director": "Morten Tyldum"
}
```

And another object in the 'Movies' collection that doesn't have the information about director but has the attribute 'ratings', might be:

```
{
  "movie_id": 2,
  "movie_name": "The Social Network",
  "release_year": 2010,
  "country": "USA",
  "ratings": 7.7
}
```

From the two instances, we can see that the structures of the items in a collection are similar but not identical, where MongoDB performs greatly as it avoids more complex structure that would be done in SQL and makes sure they can be stored in one single collection, which avoids potential complex joins and greatly improve the performances. This feature also allows the availability and scalability for big data.

MongoDB databases can be queried from many different programming languages, but the most straightforward interface is querying it through a JavaScript shell, so that is what the students

are taught in our class. The following is an example of a simple MongoDB query, written in JavaScript shell, to find the movie ids of all horror movies in our example database.

```
db.Movies.find(
  { movie_genre: "Horror" },
  { _id: 0, movie_id: 1, movie_genre: 1 }
)
```

The following code is the corresponding SQL syntax with an analogous database that has similar attributes in the database.

```
SELECT movie_id, movie_genre
FROM Movies
WHERE movie_genre = "Horror"
```

More complex operations on data in MongoDB are done with MongoDB operators, which start with a '\$' character [4]. For example, \$match is an operator that takes a specified conditions as its inputs in order to filter and produce documents that have met set conditions. Another example, \$unwind is an aggregation operator, which takes a reference to an array and produces multiple objects from the elements of the array. Another aggregation operator \$group, which takes an _id field as its first argument, then, it takes fields combined with accumulator operators to perform basic operations, such as \$sum, on the collection. Finally, \$project is an operator that takes fields as its arguments, then adds, renames, excludes or includes the specified fields in the resulting collection.

Having the structure of the previous collection in mind, we will now show an example which use the previous aggregation operators to produce a collection. This query finds all movies in the collection 2001, calculates their average ratings, and renames the country field to produce_country:

```
db.Movies.aggregate([
  {$match: {release_year: {$gte:2001}}},
  {$unwind: "$country"},
  {$group: {
    _id: "$country",
    ave_ratings: {$avg:"$ratings"}} },
  {$project: {
    ave_ratings:1,
    "produce_country": "$_id",_id:0 } }
])
```

4 Data & Methods

4.1 Data Collection

The data used in this paper are from students at a large public university. This University is research-intensive institution with about 1,800 undergraduate Computer Science majors. The data was collected from the Database Systems course at this university. The Database Systems course is an elective course taken primarily by upper-classmen full-time undergraduates and graduate students with prerequisites covering introduction to programming and data structures. The course is structured to cover three main chapters: *data models, database management systems and query languages* (relational model: relational algebra, SQL and MySQL, graph model: Neo4j and cypher and document-oriented model: JavaScript shells and MongoDB), *database design* (conceptual design and normal forms) and *database implementation* (storage and indexing, query

optimization, concurrency control). The course spends three settings to introduce and study MongoDB. Our data contains submissions from Fall 2019 and Spring 2020 from students in the Database Systems course. Our sample of students was 527 students, including 64 females and 463 males students.

The students in the databases course completed their homework using an online homework and exam platform, which we will call OHEP for anonymization purposes [7]. Upon submitting a MongoDB query in OHEP, the students were given instant feedback from auto-graders to help them modify their MongoDB queries to achieve the desired output. Students had unlimited number of attempts and no constraint in what order they may approach the questions. After close inspection of the statistics we found that the trends in student errors and number of submissions were similar for both semesters included in the study, so we will report them in aggregate.

For each MongoDB problem, the students were given a general description of the collections related to the problem. Although, Students did not have access to the objects of the collections themselves. Above the description, students see the problem that their MongoDB query should address. Students wrote their MongoDB queries in a JavaScript MongoDB enabled shell, where they used JavaScript as a query language.

Student had small text editor where they wrote their queries and had the option of saving and grading them or just saving them. When the student chooses to grade their submission, the query runs against the MongoDB objects. Should the query face a JavaScript error or a MongoDB error, it would be reported back to the student with a descriptive message from the JavaScript MongoDB enabled shell. If the query successfully runs, the student can see whether they have an incorrect result set (did not receive full points) or whether they achieved the desired output (received full points).

Also, it is worth mentioning that 2715 students' submissions were omitted from this study, because they had a common error where students copied JavaScript code from online sources, causing the interpreter to fail due to an unexpected unicode character. In addition, we do not include question ten of the homework in our analysis, because it was optional and not many students attempted to solve it.

4.2 Data Handling

This study was approved by the local Institutional Review Board. All graduate and undergraduate researchers completed IRB-mandated training in responsible conduct of research. As required by the IRB, to protect students' privacy, we assigned a unique numeric identifier to each student.

4.3 Categorization

Prior work in understanding student struggles while writing database queries with SQL have partitioned errors into different categories based on the type of error returned by the SQL engine [6, 8, 17]. However unlike most SQL engines, which will only return a single error message to the user at a time, JavaScript shells, which are used in executing MongoDB queries, may report any number of errors when you run your code. This significantly complicates the process of classifying student errors. In order to simplify

the process of arriving at a meaningful categorization, we chose to classify each submission using only the first error which was returned.

Another aspect which complicates the categorization of MongoDB errors is the fact that MongoDB queries are written through JavaScript library, rather than having its own dedicated language like SQL or Neo4j's Cypher query language [13, 14]. In practice, this means that code written to query a MongoDB database could fail on a JavaScript error before the MongoDB library code is even invoked, or the JavaScript could run successfully, only to have the MongoDB library return an error saying that though it is valid JavaScript, it does not comply with the particular usage rules of the library.

Thus we arrive at a way to partition student submissions into 4 categories: *JavaScript Errors*, where the first error returned by the JavaScript library comes from a JavaScript failure, rather than from within the MongoDB library, *MongoDB errors*, where first error is an error message given by MongoDB library denoting that the query is incorrect, *Incorrect result set*, where the query is valid and is executed by MongoDB, but returns an incorrect result, and *correct solutions*, where the query is executed by MongoDB, and the returned result set matches the expected result set.

4.4 Overview of Homework Assignments

The ten homework questions are designed to address the following concepts, and students were free to move between the problems however they wanted:

- (1) Selection and projection using `find()` (Basic) (1 question)
- (2) Selection and projection using `find()` (Advanced) (1 question)
- (3) Aggregation pipeline with group and project (1 question)
- (4) Aggregation pipeline with match, sort, group and project (1 question)
- (5) Aggregation pipeline with unwind operator (1 question)
- (6) Querying arrays (1 question)
- (7) Querying embedded documents (1 question)
- (8) Querying linked documents using cursor (1 question)
- (9) Map Reduce (1 question)
- (10) Extra Credit: Map Reduce (Advanced) (1 question)

4.5 Data Overview

We will now show an example of one of the homework problems and student solutions to that problem to demonstrate the information that is available to us in our data set. We will be looking at a student's work while trying to solve the following homework problem (the questions from Spring 2020 addressing topic (3)):

*Given a MongoDB database with two collections, **Movies** and **Actors**. For each movie genre, display the genre and the average ratings of movies under that genre. Your output field should only the name of the genre (rename it as "movie_genre") and average ratings of movies (output it as "ave_ratings").*

Database Description:

Movies collection: each Movie has its unique id (`$movie_id`), name (`$movie_name`), country (`$country`),

director (\$director), releasing date (\$release_year), ratings (\$ratings), genre (\$genre), and \$actors [an array of actor id(\$actor_id)]. Every movie is associated with only one genre and one director.

Actors collection: each actor has his/her unique id (\$actor_id), name (\$actor_name), and his/her birth country (\$birth_country).

One student first attempted to answer the prompt with the following query:

```
db.Movies.aggregate({
  $group: {
    _id: "$genre",
    ave_ratings: { $avg : "$ratings"}
  }
})
```

And they received the following feedback message:

Expected results

=====

```
{
  "movie_genre": "Horror",
  "ave_ratings": 5.869565217391305
}
```

Actual results

=====

```
{
  "_id": "Horror",
  "ave_ratings": 5.869565217391305
}
```

The problem here was the student did not use the \$project operator in order to rename the _id field as movie_genre. Then, the student came up with this MongoDB query in order to solve the renaming issue:

```
db.Movies.aggregate({
  $group: {
    movie_genre : "$genre",
    ave_ratings : { $avg : "$ratings"}
  }
})
```

but then they received the following error message:

```
Error: command failed: {
  errmsg : "The field 'movie_genre' must be an accumulator object",
  code : 40234,
  codeName : "Location40234"
}
```

The problem here was the misuse of \$group operator. the \$group operator takes the first argument as a field name to group the resulting set by, then takes the following arguments to calculate the aggregate operations on the specified fields. This is a very common problem when students try to learn the functionality of \$group. This is shown in Table 4, where you can see that students experienced this error many times on this particular question.

Category	Percentage
(1) Incorrect result set	47 %
(2) MongoDB error	32 %
(3) JavaScript error	6 %
(4) Correct	14 %

Table 1: This is the General breakdown of our Categorization

JavaScript Error Codes	Percentages
(1) Syntax error	68 %
(2) Type error	27 %
(3) Failed to parse	3 %
(4) Unknown error	2 %

Table 2: Breakdown of JavaScript error messages percentages

After a few more tries, the student came up with this MongoDB query incorporating \$project operator in order to rename _id as movie_genre. The \$project operator takes the fields' names, and either excludes the field, or includes it in the resulting set, or renames the field in the resulting set.

Finally, the student was able to get the correct answer:

```
db.Movies.aggregate({
  $group: {
    _id : "$genre",
    ave_ratings : { $avg : "$ratings" }
  }
}, {
  $project: {
    movie_genre: "$_id",
    ave_ratings : 1,
    _id: 0
  }
})
```

5 Results

Table 1 shows the breakdown of our categorization in the left column, and the percentages of each main category in respect to all submissions in the right column. In this table, we see that the majority of submissions are classified under Incorrect result set, which accounts for 47% of all submissions. Incorrect result set category have a variety of causes, such as students' lack of knowledge about some of MongoDB operators. Section 4.5 includes one of the students' MongoDB queries that produced an undesirable result set by not incorporating \$project in their MongoDB query.

Tables 2 and 3 show the most common JavaScript and MongoDB errors in their right column and their percentages of appearing among MongoDB or JavaScript errors. As can be seen in the first row of Table 3, the most prominent MongoDB error is Reference error, which happens when students do not address the specific fields correctly. Hence, Reference error have different cause, such as misspelling fields' names or updating fields properties through

MongoDB Error Codes	Percentages
(1) Reference error	60 %
(2) Field must be an accumulator object	8 %
(3) Unknown operator \$and	6 %
(4) Undefined field	4 %
(5) A pipeline stage specification object must contain exactly one field	4 %
(6) Unrecognized pipeline stage name: <code>_id \$</code>	3 %
(7) Unknown group operator	2 %
(8) Cannot return an array from Map Reduce	1 %
(9) Assert failed	1 %
(10) Namespace not found	1 %
(11) Unrecognized expression	1 %
(12) An object representing an expression must have exactly one field	1 %
(13) Illegal number of arguments used in <code>\$gte</code> .	1 %
(14) Fields' name shouldn't begin with '\$'	1 %
(15) Path option to <code>\$unwind</code> stage should be prefixed with a '\$'	1 %
(16) Bad projection specification, excluding more fields than required	1 %

Table 3: Breakdown of MongoDB error messages percentage-wise

other operators and not using updated proprieties in later operators. Furthermore, Table 4 fifth column, specifically in questions 5 to 8, shows that on more complex of a MongoDB concepts, students have more Reference errors. Likewise, the authors previous work shows similar students' performance trends between advanced SQL queries and Undefined Column errors, [8].

Table 3 reveals some problems students face when writing MongoDB queries' operators against a certain fields or operators with or without prefixing \$, specifically seen in rows 6, 14 and 15. Table 2 third entry shows students' problems with parsing their MongoDB queries in the JavaScript shell. Table 4 third column indicates that Type errors spike in question 3. Type errors are associated with students not giving the correct corresponding argument for the written MongoDB operators. Therefore, since question 3 is the first question that assesses aggregate operators, such as \$group and \$project, students spend more time and effort getting accustomed to the dynamics of aggregate operators.

Table 5 shows the distribution of students' individual submissions per question by median time (second column), the numbers of submissions (third column), the number of students who attempted the questions (fourth column) and the number of students who completed the question (fifth column). Table 5 fourth and fifth columns show us that most students completed each required question. Hence, we couldn't study the effect of students' preference of techniques on their ability to complete questions. Also, Table 5 confirms that students take more time and effort when introduced to a new MongoDB operator, specifically question 3, where aggregate operators \$group and \$project are first assessed, have the highest median time. The median time is calculated from our median of

students' session duration sample for each question, where we calculate sessions using the time elapsed from their first submission and last submission for each question.

6 Discussion

In section 5, we mentioned that as students get assessed in new MongoDB operators, they make more mistakes and need more time to generate the expected dataset. This argument finds further validity in Table 4. The fifth column of Table 4 indicates that advanced MongoDB queries and operators had a comparable and high number of Reference errors.

Another student behavior we found interesting is that when students get assessed in a new MongoDB operator, students spend more time and effort to learn and tend to make less of that mistake as time progress. Table 4 sixth column shows that students make the highest number of *Field must be an accumulator object* in question 3, which is the question in which the \$group operator is first assessed. Considering that the \$group operator was also assessed in most of the remaining questions (i.e., questions 4-8), the *Field must be an accumulator error* appear significantly less in the later questions since students have already solved question 3 and are now more comfortable with the \$group operator.

In addition, we observe in Table 5 that students spend comparable time and effort when approaching advanced concepts in MongoDB, specifically in questions 5 through 9. Also, Table 4 second and fifth columns show that students made comparable amount of reference errors and syntax errors in questions 5 to 9.

6.1 Limitations

For the sake of simplicity, we only considered the first appearing error message for the student. Considering MongoDB's integration into JavaScript and manually inspecting some of the queries, we noticed some JavaScript errors may cause additional error messages. A thorough examination of all error messages may give more in-depth insight into novices' experience with MongoDB. Also, our study only considered queries as individual entities, so, considering the context of submissions, including what error messages students received before and after the current submission, can reveal more details about students' difficulties when learning MongoDB.

Our paper only had two problem statements for each concept. Hence, we couldn't give insight about the effect of problems' wording on students' understanding. In addition, most of students completed the question prompts and students had a plenty of time to do the homework. This made it impossible to gauge the difficulty of the questions by looking at completion rates, due to strong ceiling effects. In addition, our data did not include students' ethnic or cultural background.

Another limitation is that we examine only the error output of the queries, not the queries themselves. We could gain much greater insight into why students obtained so many "Incorrect result set" errors by doing a qualitative analysis of the code that the students wrote. Future work could also benefit greatly from talk-aloud interviews with students in order to gain better understanding of student's thought processes while writing database queries.

	Javascript Errors				MongoDB Errors		
	Syntax error (3)	Type error (487)	Failed to parse (9)	Unknown error (8)	Reference error (666) (139)	Field must be an accumulator object (40234)	Unknown operator \$and (2)
(1) Selection and projection using find() (Basic)	108	46	0	0	890	9	145
(2) Selection and projection using find() (Advanced)	672	34	3	0	1774	14	12
(3) Aggregation pipeline with group and project	372	116	76	0	854	1250	7
(4) Aggregation pipeline with match, group, sort and project	353	473	19	0	768	137	244
(5) Aggregation pipeline with unwind operator	343	68	9	17	1873	150	213
(6) Querying arrays	402	122	6	0	2395	251	284
(7) Querying embedded documents	496	42	1	0	2413	42	390
(8) Querying linked documents using cursor	428	216	46	0	2150	100	146
(9) Map Reduce	292	228	0	87	1369	4	42

Table 4: Breakdown of most common Javascript and MongoDB errors by question

Question	Median Time (HH:MM)	# Submissions	# Attempted Questions	# Completed Questions
(1) Selection and projection using find() (Basic)	00:32	2788	527	525
(2) Selection and projection using find() (Advanced)	11:15	6563	524	524
(3) Aggregation pipeline with group and project	18:24	7899	525	525
(4) Aggregation pipeline with match, group, sort and project	10:39	7421	523	522
(5) Aggregation pipeline with unwind operator	16:13	10014	521	518
(6) Querying arrays	08:01	10951	521	517
(7) Querying embedded documents	07:09	10528	517	515
(8) Querying linked documents using cursor	10:29	9811	514	510
(9) Map Reduce	09:35	10216	502	499

Table 5: The Breakdown of Concepts and attempts to apply them with median time

7 Conclusion

In this paper we shed a light on students experiences when learning a NoSQL Database, MongoDB, using over 76 thousands MongoDB queries written by 527 students. Our work shows that students take comparable time and number of submissions on advanced concepts. Also, we found that students make more Reference errors in advanced queries than other queries. In addition, we found that students seem to make more JavaScript Errors when their queries have more complex MongoDB queries. Finally, we noticed that when students learn new operators of MongoDB. These results should help instructors to enhance their curriculum by allotting students more time for the MongoDB concepts discussed above and giving novices better expectations about learning MongoDB.

Acknowledgments

Redacted for anonymization.

References

- [1] 2020. Map-Reduce. <https://docs.mongodb.com/manual/core/map-reduce/>
- [2] 2020. MongoDB vs MySQL. <https://www.mongodb.com/compare/mongodb-mysql>
- [3] 2020. The most popular database for modern apps. <https://www.mongodb.com/>
- [4] 2020. Operators. <https://docs.mongodb.com/manual/reference/operator/>
- [5] Alireza Ahadi, Vahid Behbood, Arto Vihavainen, Julia Prior, and Raymond Lister. 2016. Students' Syntactic Mistakes in Writing Seven Different Types of SQL Queries and Its Application to Predicting Students' Success. In *Proceedings of the 47th ACM Technical Symposium on Computing Science Education (SIGCSE '16)*. ACM, New York, NY, USA, 401–406. <https://doi.org/10.1145/2839509.2844640>
- [6] Alireza Ahadi, Julia Prior, Vahid Behbood, and Raymond Lister. 2015. A Quantitative Study of the Relative Difficulty for Novices of Writing Seven Different Types of SQL Queries. In *Proceedings of the 2015 ACM Conference on Innovation and Technology in Computer Science Education (ITICSE '15)*. ACM, New York, NY, USA, 201–206. <https://doi.org/10.1145/2729094.2742620>

- [7] Anonymous. 2015. Paper Describing an Online Learning System.
- [8] Anonymous. 20XX. Paper analyzing student's SQL homework solutions.
- [9] Brad Fowler, Joy Godin, and Margaret Geddy. [n.d.]. Teaching Case: Introduction to NoSQL in a Traditional Database Course. <http://jise.org/Volume27/n2/JISEv27n2p99.html>
- [10] Minzhe Guo, Kai Qian, and Li Yang. 2016. Hands-on labs for learning mobile and NoSQL database security. In *2016 IEEE 40th Annual Computer Software and Applications Conference (COMPSAC)*, Vol. 2. IEEE, 606–607.
- [11] Lei Li, Kai Qian, Qian Chen, Ragib Hasan, and Guifeng Shao. 2016. Developing Hands-on Labware for Emerging Database Security. In *Proceedings of the 17th Annual Conference on Information Technology Education (SIGITE '16)*. Association for Computing Machinery, New York, NY, USA, 60–64. <https://doi.org/10.1145/2978192.2978225>
- [12] Sriram Mohan. 2018. Teaching NoSQL Databases to Undergraduate Students: A Novel Approach. In *Proceedings of the 49th ACM Technical Symposium on Computer Science Education (SIGCSE '18)*. Association for Computing Machinery, New York, NY, USA, 314–319. <https://doi.org/10.1145/3159450.3159554>
- [13] Neo4j, Inc. 2019. Neo4j. <https://neo4j.com/>
- [14] Oracle Corporation. 2019. MySQL. <https://www.mysql.com/>
- [15] Stack Overflow. 2019. Stack Overflow Developer Survey 2019. <https://insights.stackoverflow.com/survey/2019/> [Online; accessed 10-January-2020].
- [16] Toni Taipalus and Piia Perälä. 2019. What to Expect and What to Focus on in SQL Query Teaching. In *Proceedings of the 50th ACM Technical Symposium on Computer Science Education (SIGCSE '19)*. Association for Computing Machinery, New York, NY, USA, 198–203. <https://doi.org/10.1145/3287324.3287359>
- [17] Toni Taipalus, Mikko Siponen, and Tero Vartiainen. 2018. Errors and Complications in SQL Query Formulation. *ACM Trans. Comput. Educ.* 18, 3, Article 15 (Aug. 2018), 29 pages. <https://doi.org/10.1145/3231712>