

An Automated SQL Query Grading System Using An Attention-Based Convolutional Neural Network

Donald R. Schwartz¹  and Pablo Rivas² 

¹ School of Computer Science and Mathematics
Department of Computing Technology
Marist College, New York, USA
Donald.Schwartz@Marist.edu

² School of Engineering and Computer Science
Department of Computer Science
Baylor University, Texas, USA
Pablo.Rivas@Baylor.edu

Abstract. Grading SQL queries can be a time-consuming, tedious and challenging task, especially as the number of student submissions increases. Several systems have been introduced in an attempt to mitigate these challenges, but those systems have their own limitations. This paper describes our novel approach to automating the process of grading SQL queries. Unlike previous approaches, we employ a unique convolutional neural network architecture that employs a parameter-sharing approach for different machine learning tasks that enables the architecture to induce different knowledge representations of the data to increase its potential for understanding SQL statements.

Keywords: Automated query grading; self-attention models; long-short term memory models; language modeling.

1 Introduction

One of the common challenges when grading SQL queries is the large number of possible correct answers. For example, a simple query such as “Name the professors who teach a class called ‘Introduction to Programming’” could be written very many different ways: as a non-nested query with three tables in the FROM clause, as a nested EXISTS query, as a nested IN query, as a nested IN/EXISTS query, as a NATURAL JOIN, as a JOIN ON, as a JOIN USING, with or without aliasing table names, with or without aliasing field names in the answer table, etc., etc., etc. This challenge significantly increases as the complexity of the queries increases. Another challenge arises when considering partial credit for queries that are not correct. Being consistent with partial credit given the wide range of possible answers is very important but is not always easy. And a significant challenge is simply the sheer tedium of grading dozens of queries for scores of students over the course of the semester!

This paper describes our attempt to help mitigate these challenges by automating the grading process. We begin with background material, describing various approaches to automating the grading process, convolutional neural networks, and self-attention mechanisms. Next, we describe the database components within our system, the machine learning system we have developed and the results of our experiments. We conclude with the benefits of our system and future work.

2 Background

Most automated grading systems fall into one of two categories: static analysis systems and dynamic analysis systems.

Static analysis systems generally evaluate SQL queries without actually running them. They do so by comparing the structure of a submitted query with the structure of an answer query. An early example of this approach [1] used a matrix to represent a variety of relational expressions in order to establish equivalence classes for those expressions. Another system developed by Štajduhar et al. compares queries by employing string similarity metrics [24]. Cosette is a system developed by Chu et al. that determines whether queries are logically equivalent by encoding the queries into logical expressions [7]. They advanced this approach by later developing an approach using an unbound semiring to determine the level of equivalence of queries [8]. This improved system allowed for more detailed feedback to the students but was a much more complex approach. It required all possible correct answers to be included in the answer key, so that the system could compare the logical expression for the student’s query with the logical expressions of all possible answers in order to award the student with the highest possible resulting score.

The other approach is dynamic analysis, in which submitted queries are actually run against one or more datasets and the resulting table is compared to an answer key [27]. There are many early examples of this approach, including AsseSQL [18], SQLify [9] and SQLator [20]. SQL Tester is an online tool developed by Kleerekoper et al. This interactive tool does a case-sensitive comparison of the resulting tables, which requires the records to appear in the same order [13]. ASQLAG is a system capable of evaluating queries across a variety of database platforms (including SQL Server, PostgreSQL, and MySQL) by using an object-oriented design approach employing a Model-View-Controller (MVC) framework [22]. Systems that compare the answer table to the submitted table often cannot accurately determine whether the query itself is correct. For example, in the query described above, it might be the case that in the current state of the database, the Professors whose ID number is between 1000 and 2000 might be the ones teaching Introduction to Programming. A query based on ID numbers should certainly not earn credit for the original query. This weakness is especially apparent for queries whose answer is an empty table. However, comparing the answer table to the submitted table can certainly identify those queries that are NOT correct. This information can be very useful, but such systems are not well-equipped to accurately assign partial credit to those queries.

An example of this is when a submitted query results in one extra (or one fewer) field than the answer key.

Some more-recent systems combine the two approaches. Wang et al. have developed a system that identifies incorrect queries by using dynamic analysis, then employs static analysis on those queries to determine how much partial credit to award. Partial credit is determined by comparing the submitted query to multiple correct answers, awarding the highest resulting value [27]. The XData system, developed by Chandra et al., takes a similar approach. This system dynamically analyzes queries by generating specific datasets designed to expose common errors students make when developing those queries. It then uses static analysis to compute how many changes must be made to the submitted query in order to transform it to an equivalent correct query [4,3,5].

With the introduction of attention mechanisms in 2017 by Vaswani et al., researchers have developed robust algorithms to attend to specific portions of sequential data for classification, regression, and other machine learning tasks [25]. These include novel architectures, such as combining an attention mechanism with a convolutional neural network (CNN), in applications in the domain of medical imaging [28,14], sequential data [10], and others [29]. In all cases, the data contains correlated information across a temporal feature space. Researchers that used these architectures have reported improvements over classic approaches.

The success of these kinds of architectures is that CNNs have an inductive bias toward highly correlated data in a defined n-dimensional space, including images, time series, sentences, or SQL statements. The self-attention mechanism provides contextual information about the data within its context, making the architecture very powerful.

Other work has been done in modeling SQL statements [6,2]; however, the authors have significantly different data, aims, and goals. Architecturally, our proposed model, depicted in Fig. 1, is different as it uses a parameter-sharing approach for different machine learning tasks that enable the architecture to induce different knowledge representations of the data to increase its potential for SQL statement understanding. To the best of our knowledge, no other approaches have used this architecture and learning paradigms on SQL statement grading.

3 Description of the Database Components

Our overall database contains all of the tables required for the assignments (which, using the above example, would include at least **PROFESSORS**, **TEACHES** and **COURSES**), along with the following tables, some of which are used to provide the information needed for our proposed system.

The **STUDENT** table contains identification information about each student registered for each section.

The **QUERY** table consists of the **QueryID**, the text/wording of the query, and the number of possible points for that query.

The **ASSIGNMENT** table describes which queries are part of which lab assignment.

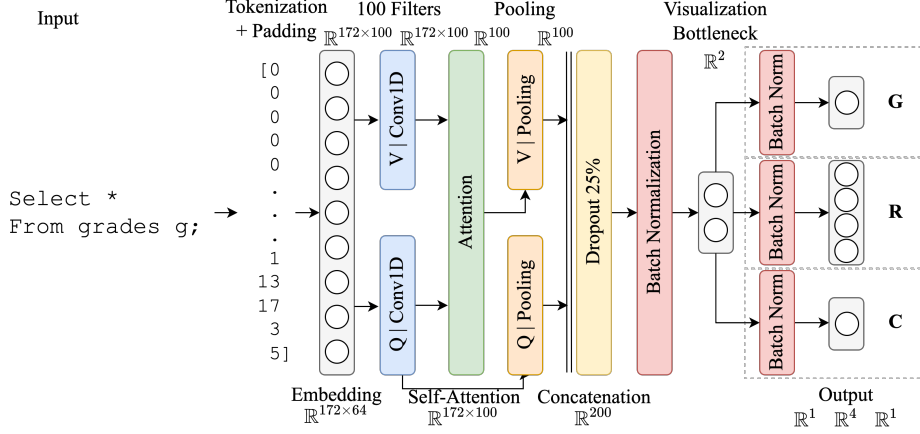


Fig. 1. Proposed neural architecture for SQL Statement grading based on self-attention, CNNs, and parameter sharing.

The **ANSWER** table holds the **AnswerID**, the **QueryID**, and a correct SQL answer for that query. Naturally, there are multiple **AnswerID**s associated with each **QueryID**.

The largest table is the **STUDENTSUBMISSION** table. This table holds the **StudentID**, **QueryID**, **SubmittedAnswer** (the SQL query submitted by that student), a flag indicating whether the **SubmittedAnswer** **IsCorrect**, a **Remark** field that describes whether the submission was **correct**, **partially correct**, **uninterpretable**, and the number of **Points** earned.

We populated our database with queries written by our students over the past couple of years. This required us to regrade the assignments to remove penalties that didn't necessarily affect the queries themselves, for example, late penalties, deductions for not following directions, and penalties for using non-descriptive aliases for tables and fields. Our datasets included roughly 2700 individual student queries.

The following data is fed into our system: **SubmissionID** (which is based on the **studentID** and **queryID**), **SubmittedAnswer**, **IsCorrect**, **Remark** and **Grade** (the number of points earned divided by the possible points for that query).

4 Description of the Machine Learning System

The neural architecture displayed in Fig. 1 will be discussed in detail next.

4.1 Embedding Layer

The embedding layer is a trainable set of neurons that take a token in the context of the entire sequence [11,19]. Our sequence has size 172 and returns a vector of size 64 to represent a token in the sequence.

4.2 Convolutional Self-Attention Layer

The self-attention CNN is composed of three major parts. First, we have two convolutional encoding layers that model a query Q and value V . These are classic elements of attention mechanisms [25]. For self-attention, the input to both Q and V is the same embedding. Second, the self-attention layer uses a dot product similarity [16]. Third, there is a pooling strategy that implements a global average, reducing the dimensionality of the problem down to two vectors of size 100 each [15]. During training, this layer learns optimal embeddings based on self-attention.

4.3 Dropout Layer

To regularize the model and reduce the chances of overfitting we used a 25% dropout [23]. This mechanism further promotes a holistic learning from the attention embeddings themselves.

4.4 Batch normalization Layers

To improve numerical stability during learning and help regularize the learning problem, we used batch normalization layers. Batch normalization calculates optimal means and variances for data batches reducing, vanishing or exploding gradient problems [21].

4.5 Bottleneck Layer

The model has a bottleneck dense layer with hyperbolic tangent activations that enables us to visualize the learned representation space. We use two neurons only, allowing us to display the learned representations in two dimensions if needed.

4.6 Outputs C , R , and G

The right-most portion of Fig. 1 displays three different groups of outputs, each of which are different models that share the same parameters of the prior layers, up to the bottleneck. These models each have a batch normalization layer and dense neural units. Model C is for predicting correctness, model R is for predicting the grader’s remark, and model G is for predicting the grade. Each of these models is trained iteratively as we discuss in the next section.

4.7 Experiments, Evaluation, and Results

The model in Fig. 1 was implemented in Python using the Tensorflow platform with Keras libraries. The training discussed below was performed on an NVIDIA P100 GPU system with 25 GB of RAM and 166 GB of storage.




Fig. 2. ROC curves and corresponding AUC across the 10 different folds. The average AUC is 0.64. In all folds, the model performs better than random chance.

We performed two major experiments: 1) a jointly trained model with the purpose of validating correctness on k-fold cross-validation; and 2) a jointly trained model that validates all models using leave-one-out cross-validation.

For our experiments the three models **C**, **R**, and **G** are trained one at a time with different loss functions optimized using gradient descent with an RMSprop optimizer and learning rate of 0.001 [17].

Jointly trained model cross-validated performance. During training in this experiment, the gradient is calculated using a single loss: the binary cross entropy loss. The length of the output is a vector of six elements: one item for the correctness, four items for the remarks, and one item for the grade. We measure the performance using the area under the receiving operating characteristic (ROC) curve (AUC) using 10-fold cross validation. The motivation for measuring this at this point is to make an early assessment before launching a full-scale performance evaluation.

The results of using cross validation are displayed in Fig. 2. From the figure we can assert that the model, in average across folds, performs better than random chance in all cases. The maximum AUC reported was 0.72 and the minimum was 0.57, with an average of **0.64** and very low variance, 0.04. To improve the estimation of the generalization error we also performed a full cross-validation analysis using a leave-one-out (LOO) strategy across all models.




Fig. 3. Confusion matrix for the model that predicts correctness.

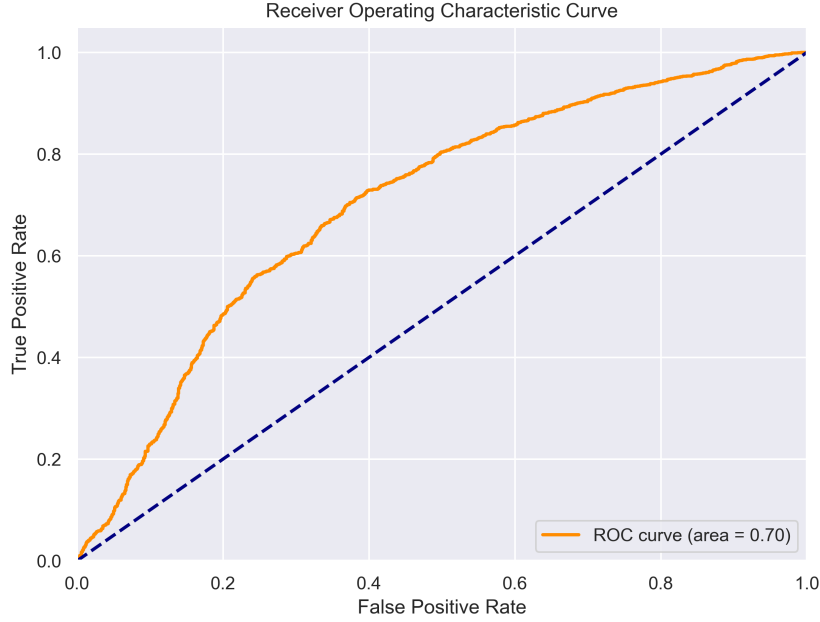


Fig. 4. ROC and AUC for the model that predicts correctness.

LOO cross-validation on all models. LOO cross-validation is well-known as the most robust generalization error estimator, particularly when it comes to small datasets such as the one we use [12,26]. For the correctness model, **C**, after training under a LOO analysis, we obtain the results shown in Fig. 3 and 4. The confusion matrix shown in Fig. 3 indicates the slightly higher false negative (FN) count on the correctness predictions; however, the class imbalance of the dataset may be affecting this outcome. This imbalance obligates us to consider the balanced accuracy instead of the traditional accuracy score. In this case the

balanced accuracy is 66.5%. We also calculate the LOO cross-validated ROC and AUC, shown in Fig. 4. From the figure we observe that the model performs far beyond random chance and reaches an AUC of **0.70**.

Now we briefly look at mistakes, specifically beginning with an example of the *best mistakes*, i.e., mistakes that were so close to the actual prediction but did not make the classification threshold by two or three decimal points, or in other words, mistakes with high uncertainty. The *best mistake* was:

```
SELECT id, first_name,last_name
FROM person p
WHERE p.year_born = (SELECT MIN(p1.year_born)
                     FROM person p1);
```

The sample above is labeled as correct=0, also remark= partially correct, and with grade=40. However, our model predicted it incorrectly as a correct=1 SQL statement with the highest uncertainty, which can be explained by the grade and partial credit given.

Now let's inspect the *worst mistake*, i.e., a mistake that was made with highest confidence. Such a sample is:

```
SELECT DISTINCT ztv.show_name, SUM(zprb.prod_salary) +
SUM(zply.actor_salary) AS
Total_Price
FROM ztvshow ztv, zprodby zprb, zplay zply, zactor zact
WHERE zprb.show_num = ztv.show_num
AND zply.show_num = ztv.show_num
AND zact.actor_num = zply.actor_num
GROUP BY ztv.show_name;
```

The sample above is labeled as correct=0, also remark=partially correct, and with grade=66.66. Once again this was predicted incorrectly as a correct=1 but this time with the highest degree of certainty. Again, notice the higher grade that it receives and that it is partially correct, regardless of the lack of aesthetics in the SQL statement.

For the remarks model, **R**, after performing LOO, we obtain the results shown in Fig. 5. The figure shows that the model associated with the remarks about correctness, i.e., Correct and Partially Correct, exhibit better performance in comparison to the remarks associated with negative feedback. These results are consistent with our previous finding that correctness can be better predicted than non-correctness. In Fig. 5, for reference, the average precision (AP) score is listed, which is a weighted mean of precisions; it also shows the F_1 curves for reference. In Fig. 5 and the confusion matrix in Fig. 6 we can also observe that the lowest-performing class was Cheating; this is due to the limited number of samples available.

Finally, for the grades model, **G**, the LOO cross-validated performance analysis can be summarized in Fig. 7 and 8. A regression comparison between the predicted grade and the actual grade suggests that our model disfavors grades that are




Fig. 5. Precision-Recall curves and average precision scores per remark.

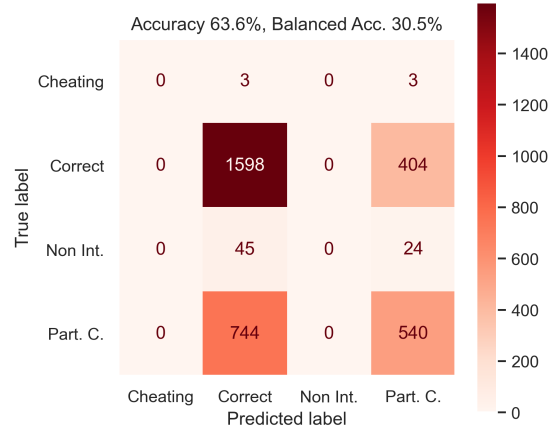


Fig. 6. Confusion matrix for the model that predicts remarks.

exactly 100% and favors a more conservative prediction in a wider range. This can be seen in the histogram of the residuals in Fig. 8, which shows a slight positive skew, while the ideal residuals should be centered around zero. However, the regression performance metrics still capture the predictive ability of the model, which is beyond simply predicting the expected value of the variable Grade.




Fig. 7. True vs predicted regression plot for the model that predicts grades.




Fig. 8. Histogram of the residuals $y - \hat{y}$ for the model that predicts grades.

5 Conclusions and Future Work

As machine learning-based language modeling applications continue to increase in number and variety, we anticipate more research in the general areas of SQL modeling for different analysis and automated decision-making tasks. In this paper

we presented a light-weight language model based on convolutional self-attention and parameter-sharing trained across different tasks on SQL statement data. The presented model can be used within another system that processes SQL-based assignment analysis as downstream tasks. We believe this type of system can effectively address the problem of automatic grading of SQL statements.

Future research will experiment with large language models for an assessment of performance as a function of hyperparameter size; further, the collection of more labeled data with a wider variety of submissions and assignments can be analyzed to distinguish differences in performance according to assignment types, e.g., assignments that vary in difficulty and performance analysis by difficulty.

Acknowledgements

The ML model is based upon work supported in part by the National Science Foundation under Grants CHE-1905043, CNS-2136961, and CNS-2210091.

References

1. Aho, A.V., Sagiv, Y., Ullman, J.D.: Equivalences among relational expressions. *SIAM Journal on Computing* **8**(2), 218–246 (1979)
2. Bogin, B., Berant, J., Gardner, M.: Representing schema structure with graph neural networks for text-to-sql parsing. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 4560–4565 (2019)
3. Chandra, B., Banerjee, A., Hazra, U., Joseph, M., Sudarshan, S.: Automated grading of sql queries. In: *2019 IEEE 35th International Conference on Data Engineering (ICDE)*, pp. 1630–1633. IEEE (2019)
4. Chandra, B., Banerjee, A., Hazra, U., Joseph, M., Sudarshan, S.: Edit based grading of sql queries. In: *8th ACM IKDD CODS and 26th COMAD*, pp. 56–64 (2021)
5. Chandra, B., Joseph, M., Radhakrishnan, B., Acharya, S., Sudarshan, S.: Partial marking for automated grading of sql queries. *Proceedings of the VLDB Endowment* **9**(13), 1541–1544 (2016)
6. Chen, Z., Chen, L., Zhao, Y., Cao, R., Xu, Z., Zhu, S., Yu, K.: Shadowgnn: Graph projection neural network for text-to-sql parser. In: *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 5567–5577 (2021)
7. Chu, S., Li, D., Wang, C., Cheung, A., Suciu, D.: Demonstration of the cosette automated sql prover. In: *Proceedings of the 2017 ACM International Conference on Management of Data*, pp. 1591–1594 (2017)
8. Chu, S., Murphy, B., Roesch, J., Cheung, A., Suciu, D.: Axiomatic foundations and algorithms for deciding semantic equivalences of sql queries. *Proceedings of the VLDB Endowment* **11**(11) (2018)
9. Dekeyser, S., de Raadt, M., Lee, T.Y.: Computer assisted assessment of sql query skills. In: *Proceedings of the 18th Australasian Database Conference (ADC 2007)*, vol. 63, pp. 53–62. Australian Computer Society Inc. (2007)
10. Fahim, S.R., Sarker, Y., Sarker, S.K., Sheikh, M.R.I., Das, S.K.: Self attention convolutional neural network with time series imaging based feature extraction for transmission line fault detection and classification. *Electric Power Systems Research* **187**, 106,437 (2020)

11. Ghazanfari, B., Afghah, F.: Multi-level feature learning on embedding layer of convolutional autoencoders and deep inverse feature learning for image clustering. arXiv preprint arXiv:2010.02343 (2020)
12. Kearns, M., Ron, D.: Algorithmic stability and sanity-check bounds for leave-one-out cross-validation. *Neural computation* **11**(6), 1427–1453 (1999)
13. Kleerekoper, A., Schofield, A.: Sql tester: an online sql assessment tool and its impact. In: *Proceedings of the 23rd annual ACM conference on innovation and technology in computer science education*, pp. 87–92 (2018)
14. Li, M., Hsu, W., Xie, X., Cong, J., Gao, W.: Sacnn: Self-attention convolutional neural network for low-dose ct denoising with self-supervised perceptual loss network. *IEEE transactions on medical imaging* **39**(7), 2289–2301 (2020)
15. Li, Z., Wang, S.H., Fan, R.R., Cao, G., Zhang, Y.D., Guo, T.: Teeth category classification via seven-layer deep convolutional neural network with max pooling and global average pooling. *International Journal of Imaging Systems and Technology* **29**(4), 577–583 (2019)
16. Luong, M.T., Pham, H., Manning, C.D.: Effective approaches to attention-based neural machine translation. In: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pp. 1412–1421 (2015)
17. Mukkamala, M.C., Hein, M.: Variants of rmsprop and adagrad with logarithmic regret bounds. In: *International conference on machine learning*, pp. 2545–2553. PMLR (2017)
18. Prior, J.C., Lister, R.: The backwash effect on sql skills grading. *ACM SIGCSE Bulletin* **36**(3), 32–36 (2004)
19. Rivas, P.: *Deep Learning for Beginners: A beginner’s guide to getting up and running with deep learning from scratch using Python*. Packt Publishing Ltd (2020)
20. Sadiq, S., Orlowska, M., Sadiq, W., Lin, J.: Sqlator: an online sql learning workbench. In: *Proceedings of the 9th annual SIGCSE conference on Innovation and technology in computer science education*, pp. 223–227 (2004)
21. Santurkar, S., Tsipras, D., Ilyas, A., Madry, A.: How does batch normalization help optimization? *Advances in neural information processing systems* **31** (2018)
22. Singporn, P., Vichianroj, P., Trongratsameethong, A.: Asqlag-automated sql assignment grading system for multiple dbms. *Journal of Technology and Innovation in Tertiary Education* **1**(1), 41–59 (2018)
23. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* **15**(1), 1929–1958 (2014)
24. Štajduhar, I., Mauša, G.: Using string similarity metrics for automated grading of sql statements. In: *2015 38th International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*, pp. 1250–1255. IEEE (2015)
25. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
26. Wang, J.: *Dataset: combining dynamic and static analysis for automated grading sql statements* (2020). URL <https://doi.org/10.5281/zenodo.6526769>
27. Wang, J., Zhao, Y., Tang, Z., Xing, Z.: Combining dynamic and static analysis for automated grading sql statements. *J Netw Intell* **5**(4), 179–190 (2020)
28. Wu, Y., Ma, Y., Liu, J., Du, J., Xing, L.: Self-attention convolutional neural network for improved mr image reconstruction. *Information sciences* **490**, 317–328 (2019)
29. Zeng, W., Li, M.: Crop leaf disease recognition based on self-attention convolutional neural network. *Computers and Electronics in Agriculture* **172**, 105,341 (2020)