

Predicting Student Math Assignment Performance

Quinn Straus

In this project, I developed a model to predict how a student would perform on a given math assignment given their assignment history. I leveraged the ASSISTments 2020-2021 dataset, a large dataset containing all student logs for problems solved on the ASSISTments math problem bank. I processed this data into sequences for each student, with all relevant assignment information encoded into each point in the sequence. I then fed these sequences into a gated recurrent unit (GRU) model to predict the difficulty the student would have on each assignment. This difficulty metric was derived from a combination of correctness, number of attempts, and time spent on each problem in the assignment, and was normalized to a 0-1 scale. My model was able to obtain a mean absolute error of 0.169, compared to an error of 0.235 achieved by a naive predictor.