Deep Knowledge Tracing and Engagement in MOOCs

Klint Kanopka, David Lang, and Kritphong Mongkhonvanit

Stanford University

Abstract

MOOCs and online courses have notoriously high attrition. One challenge is that it can be difficult to tell if students fail to complete because of disinterest or because of course difficulty. Starting from the Deep Knowledge Tracing framework, we account for student engagement by including course interaction covariates. With these, we find that we can predict a student's next item response with over 88% accuracy. Based on these predictions, targeted interventions can be offered to students and courses can be improved.

Introduction

Deep Knowledge Tracing is a way to estimate student knowledge through a course and has been explored with MOOCS in several ways:

- Predict and identify latent skills pertaining to items and tasks [1]
- Clickstream interactions e.g.(pause,plays, and video interactions) are a feature that improves grade prediction [2]

We are combining past item response information with video engagement interactions to predict the probability a student will answer the next item correctly.

Data

anon_name	feature	index1	index2	···	index103
38fqh9dy	items attempted	103	···		
38fqh9dy	correct	0	1	···	1
38fqh9dy	playback_speed	1	1.25	···	0
38fqh9dy	pauses	1	0	···	0
38fqh9dy	seek_back	0	0	···	0
38fqh9dy	seek_forward	0	0	···	0
38fqh9dy	video_completed	0	0	···	0
38fqh9dy	attempt	1	1	···	0
38fqh9dy	quiz	0	0	···	0

Figure 1: Example entry in the data set for one student

- Data comes from 12,007 students in a MOOC course on statistics with item responses and video interactions for each student.
- Video interaction data contains seeks, pauses, average playback speed, and video completion.
- Individual items are directly mapped to specific instructional videos.

Model

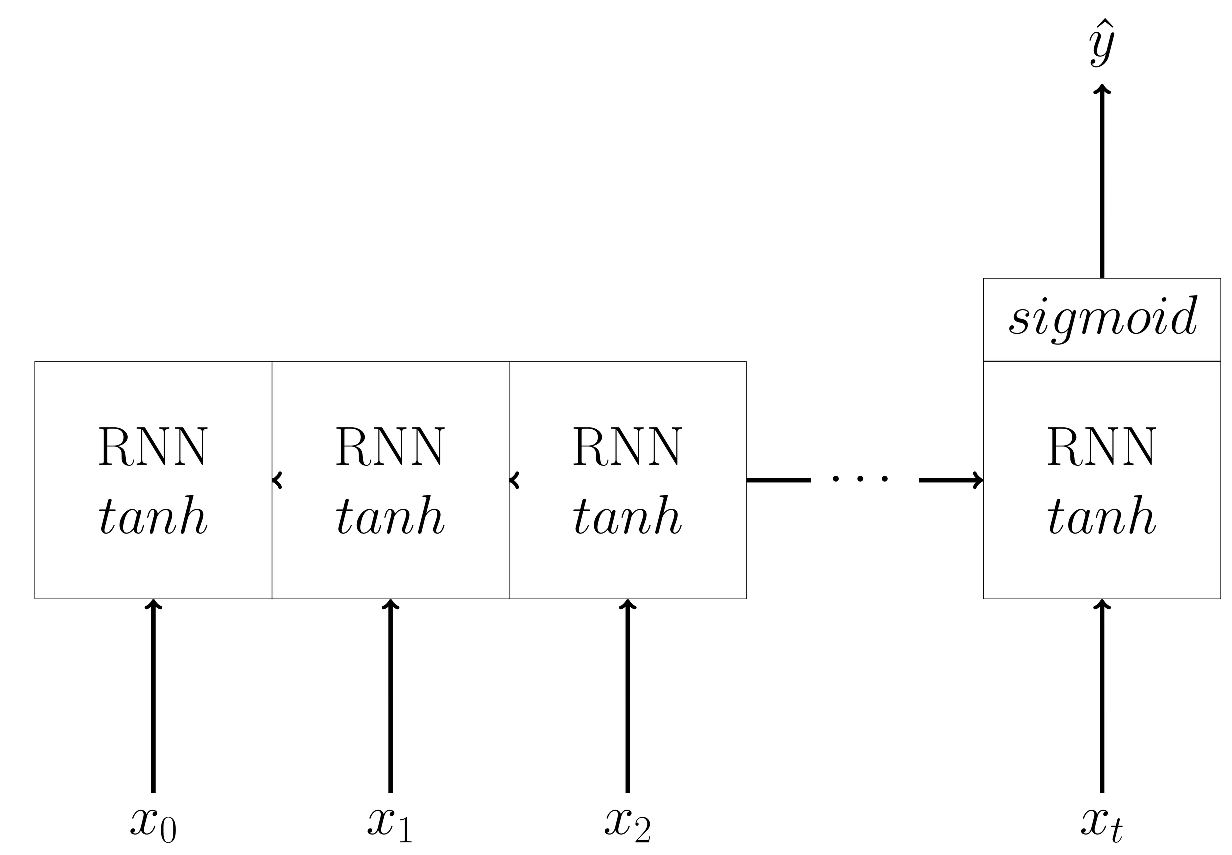

- Each RNN block contains 128 hidden units. Learning was done using binary cross-entropy loss with L2 regularization and an Adam optimizer.
- Each x_t contains the item response and interaction data for the video associated with item t.
- $\hat{y}$ is the predicted probability of a correct item response for item t.
- After training for 150 epochs, predictive accuracy was 0.8834

Output

- Predicted item response can be plotted as a function of item number.
- Regions of low correct response probability can identify struggling students or areas to improve the course.

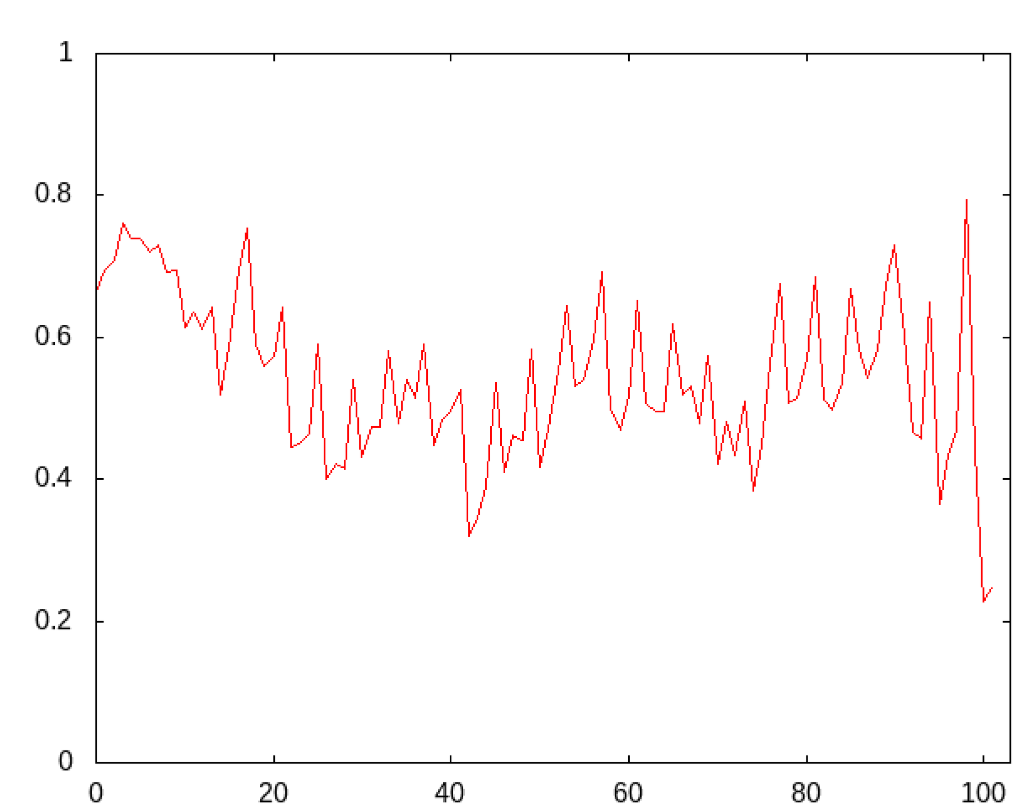

Figure 2: Example output graph for a single student

Discussion

- Adding clickstream data as a covariate increases prediction accuracy.
- For the course we studied, simple RNNs outperform GRU and LSTM based models.
- Our results likely understate the accuracy gains of video interaction data as, by the end of the course, almost three-quarters of users did not interact with course videos.
- Our method functions independent of the content of the course and requires no additional qualitative coding of items.

Applications

This work has applications at the whole-course level and at the individual student level. At the student level:

- Drops in predicted item response probability quickly identify students who need extra support or are at risk of dropping out of the course.
- Video interaction data can help inform appropriate intervention strategy for each student.

At the course level:

- Coordinated drops in predicted item response probability for all students can identify challenging content areas.
- Trends in video interaction can inform the improvement of communication, course materials, test items, and scaffolding methods.

Future Work

- Apply model to new courses and differently structured items.
- Model similarity to "ideal engagement" to identify interventions.

References

[1] Chris Piech, Jonathan Bassen, Jonathan Huang, Surya Ganguli, Mehran Sahami, Leonidas J Guibas, and Jascha Sohl-Dickstein.
Deep knowledge tracing.
In *Advances in Neural Information Processing Systems*, pages 505–513, 2015.

[2] Tsung-Yen Yang, Christopher G Brinton, Carlee Joe-Wong, and Mung Chiang.
Behavior-based grade prediction for moocs via time series neural networks.
IEEE Journal of Selected Topics in Signal Processing, 11(5):716–728, 2017.