

DEEP LEARNING TO PREDICT STUDENT OUTCOMES

Byung-Hak Kim
 Udacity, AI Team
 hak@udacity.com

ABSTRACT

The increasingly fast development cycle for online course contents, along with the diverse student demographics in each online classroom, make real-time student outcomes prediction an interesting topic for both industrial research and practical needs. In this paper, we tackle the problem of real-time student performance prediction in an on-going course using a domain adaptation framework. This framework is a system trained on labeled student outcome data from previous coursework but is meant to be deployed on another course. In particular, we introduce a GritNet architecture, and develop an *unsupervised* domain adaptation method to transfer a GritNet trained on a past course to a new course without any student outcome label. Our results for real Udacity student graduation predictions show that the GritNet not only *generalizes* well from one course to another across different Nanodegree programs, but also enhances real-time predictions explicitly in the first few weeks when accurate predictions are most challenging.

1 INTRODUCTION

With the growing need for people to keep learning throughout their careers, massive open online course (MOOCs) companies, such as Udacity and Coursera, not only aggressively design new and relevant courses, but they also refresh existing course content to keep the material up-to-date. This effort results in a significant increase in student numbers, which make it impractical for even experienced human instructors to assess an individual student and anticipate their learning outcomes. Moreover, students in each MOOC classroom are heterogeneous in their background and intention, which is very different from a classic classroom (Chuang & Ho, 2016; Economist, 2017). Even subsequent offerings of a course within a year can have different sets of students, mentors, instructors, content, and workflows. In this world of MOOCs, an automated machine which reliably forecasts a student’s performance early in their coursework would be a valuable tool for making smart decisions about when (and with whom) to make live educational interventions. These interventions have the aim of increasing engagement, providing motivation, and empowering students to succeed.

2 GRITNET

In that, we first recast the student performance prediction problem as a sequential event prediction problem and introduce the GritNet. The task of predicting student outcomes can be expressed as a sequential event prediction problem: given a past event sequence $\mathbf{o} \triangleq (o_1, \dots, o_T)$ taken by a student, estimate likelihood of future event sequence $\mathbf{y} \triangleq (y_{T+D}, \dots, y_{T'})$ where $D \in \mathbb{Z}_+$. In the form of online classes, each event o_t represents a student’s action (or activities) associated with a time stamp. In other words, o_t is defined as a paired tuple of (a_t, d_t) , where each action a_t represents, for example, “a lecture video viewed”, “a quiz answered correctly/incorrectly”, or “a project submitted and passed/failed”, and d_t states the corresponding (logged) time stamp. Then, the log-likelihood of $p(\mathbf{y}|\mathbf{o})$ can be written approximately as $\sum_{i=T+D}^{T'} \log p(y_i|v)$, given a fixed-dimensional embedding representation v of \mathbf{o} . The goal of each GritNet is, therefore, to compute an individual log-likelihood $\log p(y_i|v)$, and those estimated scores can be simply added up to estimate long-term student outcomes.

In order to feed students’ raw event records into the GritNet, it is necessary to encode the time-stamped logs (ordered sequentially) into a sequence of fixed-length input vectors by *one-hot encoding*. A one-hot vector is a vector $\mathbb{1}(a_t) \in \{0, 1\}^L$, where L is the number of unique actions and only one element could take the value 1 to distinguish each activity a_t from every other. One-hot vectors

of the same student are connected into a long vector sequence to represent the student’s whole sequential activities in \mathbf{o} . To further capture a student’s varying learning speeds, GritNet defines the discretized time difference between adjacent events Δ_t as $d_t - d_{t-1}$ ¹. Then, one-hot encode Δ_t into $\mathbb{1}(\Delta_t)$ and connect them with the corresponding $\mathbb{1}(a_t)$ to represent $\mathbb{1}(o_t)$ as $[\mathbb{1}(a_t); \mathbb{1}(\Delta_t)]$. Lastly, the output sequences shorter than the maximum event sequence length (of a given training set) are pre-padded with all $\mathbf{0}$ vectors.

The complete GritNet architecture is the following. The first embedding layer (Bengio et al., 2001) learns an embedding matrix $\mathbf{E}^o \in \mathbb{R}^{E \times |O|}$, where E and $|O|$ are the embedding dimension and the number of unique events (i.e., input vector $\mathbb{1}(o_t)$ size). This embedding can be used to convert an input vector $\mathbb{1}(o_t)$ onto a low-dimensional embedding v_t defined as $\mathbf{E}^o \mathbb{1}(o_t)$. This (dense) event embedding v_t is then passed into the bidirectional long short term memory (BLSTM) (Graves & Schmidhuber, 2005) and the output vectors are formed by concatenating each forward and backward direction outputs. Next, a global max pooling (GMP) layer (Collobert & Weston, 2008) is added to form a fixed dimension vector (independent of the input event sequence length) by taking the maximum over time (over the input event sequence). This GMP layer is able to capture the most relevant signals over the sequence and is also able to deal with the imbalanced nature of data. One can view this GMP layer as a hard self-attention layer or can consider the GMP layer output generates a sequence level embedding of the whole input event sequence. The GMP layer output is, ultimately, fed into a fully-connected (FC) layer, and a softmax (i.e., sigmoid) layer subsequently, to calculate the log-likelihood $\log p(y_i|v)$.

In particular, GritNet is the first deep learning architecture which successfully advances the state of the art by demonstrating substantial prediction accuracy improvements (particularly pronounced in the first few weeks when predictions are extremely challenging). In contrast to other works (Mi & Yeung, 2015; Piech et al., 2015; Whitehill et al., 2017a; Wang et al., 2017), the GritNet does not need any feature engineering (it can learn from raw input) and it can operate on any (raw) student event data associated with a time stamp even when highly imbalanced.

3 DOMAIN ADAPTATION WITH GRITNET

Though GritNet’s superior prediction accuracy was reported in Kim et al. (2018), it was not fully addressed whether GritNet models transfer well to different courses or if they could be deployed for real-time prediction with ongoing courses. With an increase in new courses and the fast pace of content revisions of MOOC classes to meet students’ educational needs, the GritNet model is required to generalize well to unseen courses and to predict performance in real-time. This prediction happens for students who have not yet finished the course (i.e., new data without prior knowledge of the students’ outcomes).

3.1 ORDINAL INPUT ENCODING

Specifically, in order to leverage the unlabeled data in the target course, we propose an ordinal input encoding procedure as a basis for GritNet to provide transferable sequence-level embeddings across courses. For any MOOC coursework, each content exhibits a natural order, following a contents tree. For the sake of simplicity, we group content IDs into subcategories and convert each content ID to corresponding ordinal ID by exploiting actual ordering implicit in the content paths within each category. Then each encoded MOOC contains the same number of unique actions as the original. For example, the Udacity data in Table 1 contains three subcategories - content pages, quizzes, and projects. One Nanodegree program dataset has three sequences of ordinal IDs - content-1 to content- i , quiz-1 to quiz- j and project-1 to project- k . Note that with this encoding, the number of unique actions L turns out to be equivalent to $i + 2 \times j + 2 \times k$, since quizzes and projects allow two potential actions (i.e. a quiz answered correctly/incorrectly or a project submitted and passed/failed). It is also worth mentioning that different courses may have different i , j , k and L values.

3.2 PSEUDO-LABELS AND TRANSFER LEARNING

To deal with unlabeled target course data in *unsupervised* domain adaptation (Pan & Yang, 2010) from source to target course, we harness the trained GritNet on a source data to assign pseudo labels to the target data. This is a common real-world scenario when releasing new, never-before consumed educational content. In other words, we use a trained GritNet on a source course to assign

¹For the Udacity data described in Section 4.1, we use day to represent inter-event time intervals.

Table 1: Udacity data characteristics

Dataset	Enrolled From	Students	Contents (i)	Quizzes (j)	Projects (k)	AvgSeq. Length	Graduates	Grad. Rates
ND-A v1	2/3/2015	5,626	471	168	4	421	1,202	21.4%
ND-A v2	4/1/2015	2,230	471	168	4	881	453	20.3%
ND-B	3/22/2017	4,377	568	84	10	675	1,726	39.4%
ND-C	1/14/2017	4,761	346	50	5	430	2,198	46.2%

pseudo outcome labels to the target course, then continue training the GritNet on the target course (as precisely described in Algorithm 1).

Algorithm 1 Domain Adaptation with GritNet

Require: Source course data \mathcal{X}_{source} , Source label \mathcal{Y}_{source} , Target data \mathcal{X}_{target} , Threshold θ

- 1: Set source training set as $\mathcal{T}_{source} = (\mathcal{X}_{source}, \mathcal{Y}_{source})$
 - 2: Train $GritNet_{source}$ with \mathcal{T}_{source}
 - 3: Evaluate on \mathcal{X}_{target} : $\hat{\mathcal{Y}}_{pred} := GritNet_{source}(\mathcal{X}_{target})$
 - 4: Assign pseudo-labels to \mathcal{X}_{target} : $\mathcal{Y}_{label} := \mathbb{1}(\hat{\mathcal{Y}}_{pred} \geq \theta)$
 - 5: Update target training set as $\mathcal{T}_{adapt} = (\mathcal{X}_{target}, \mathcal{Y}_{label})$
 - 6: Freeze all the $GritNet_{source}$ but the last FC layer and continue training with \mathcal{T}_{adapt}
-

This method was motivated by the observation that the GMP layer output captures generalizable sequential representation of an input event sequence, whereas the high-level FC layer learns the course-specific features. It is worth mentioning that the last FC layer limits the number of parameters to learn in Step 6. So this enables the entire Algorithm 1 to be a very practical solution for faster adaptation time (and in fact deployment cycle) and applicability for even a (relatively) smaller size target course. To our knowledge, this is the first work to effectively answer the source and target course distribution mismatch encountered in real-time student performance prediction, which had been considered as a challenging open problem in prior works (Boyer & Veeramachaneni, 2015; Whitehill et al., 2017a;b; Dalipi et al., 2018).

4 GENERALIZATION PERFORMANCE

4.1 UDACITY DATA

To demonstrate the benefits of Algorithm 1, we benchmark the student graduation prediction task on various Udacity dataset (see Table 1 for more information of each Udacity dataset used for this study). In all programs, graduation is defined as completing each of the required projects in a program’s curriculum with a passing grade. Each program has a designated calendar start and end date for each cohort of students that passes through, and users have to graduate before the official end of their cohort’s term to be considered successfully graduated. Each program’s curriculum contains a mixture of video contents, written contents, quizzes, and projects. Note that it is not required to interact with every piece of content or complete every quiz to graduate. For all datasets, an event represents a user taking a specific action (e.g., watching a video, reading a text page, attempting a quiz, or receiving a grade on a project) at a certain time stamp. It should be noted that no personally identifiable information is included in this data and student equality is determined via opaque unique ids.

4.2 PERFORMANCE RESULTS

In order to evaluate the transferability of the GritNet across different courses, we consider three different scenarios from one course to another on the setups below (see Figure 1 for the results):

- **Vanilla Baseline** using the same logistic regression based baseline model in Kim et al. (2018) is trained on a source course and evaluated on a target one.
- **GritNet Baseline** is trained on a source course and evaluated on a target course. Note that there is no training data or labels from target course for training the GritNet.
- **Domain Adaptation** follows the steps specified in Algorithm 1. Note that for pseudo-labels generation (Step 4), a hyper-parameter θ threshold of [0.1, 0.2, 0.3, 0.4] is tried.
- **Oracle** performance bound of domain adaption is computed as the same procedures in Algorithm 1 while skipping Step 3 and Step 4 to use oracle target labels \mathcal{Y}_{target} instead of assigned pseudo-labels \mathcal{Y}_{label} .




Figure 1: Real-time student graduation prediction accuracy comparisons of models (a) from ND-A v1 to the next version ND-A v2 and (b-c) from earlier ND-A v1 program to two later Udacity Nanodegree programs, ND-B and ND-C. (a): GritNet baseline shows only 1.01% abs accuracy loss ($< 5.30\%$ abs loss of vanilla baseline) in average over eight weeks as compared to oracle bound. The accuracy gains of GritNet baseline suggest that sequence-level embeddings trained with GritNet result in more transferable features as compared to features learned with a Vanilla baseline model. (b-c): Being consistent with (a) results, the clear wins of the GritNet baseline over the Vanilla baseline reaffirm that the sequence-level embedding trained with GritNet is more robust to source and target course distribution mismatch as compared to features learned with the Vanilla baseline model. Furthermore, domain adaptation provides 70.60% accuracy recovery in average during first four weeks (up to 84.88% at week 3) for ND-B dataset and 58.06% accuracy recovery during the same four weeks (up to 75.07% at week 3) for ND-C dataset from GritNet baseline performances.

For all models across the setups, we trained different models of a GritNet architecture of BLSTM with forward and backwards LSTM layers containing 256 cell dimensions per direction and embedding layer dimension of 512 for different weeks. Each model is based on students’ week-by-week event records, to predict whether each student was likely to graduate. Further hyper-parameter optimization could be done for the optimal accuracy at each week.

The accuracy of each system’s prediction was measured by the area under the Receiver Operating Characteristic curve (AUC) which scores between 0 and 100% (the higher, the better)². We used 5-fold student level cross-validation, while ensuring each fold contained roughly the same proportions of the two groups (graduate and non graduate) of students. To compare domain adaption performance with the upper bound assessed with oracle setup, we consider mean AUC recovery rate (ARR), as a metric to evaluate AUC improvements from *unsupervised* domain adaptation against oracle training. This measure is defined as absolute AUC reduction that *unsupervised* domain adaptation produces divided by the absolute reduction that oracle (*supervised*) training produces. We assume that the upper bound for *unsupervised* domain adaptation would be the performance of the oracle training on the same data, which has a AUC recovery of 100%.

5 CONCLUSION

In this paper, we have proposed an extremely practical approach for predicting real-time student performance, which has been considered as of the utmost importance in MOOCs (but under-explored), as this enables student outcomes prediction while a course is on-going. We introduced a novel domain adaptation algorithm with GritNet and demonstrated that GritNet can be transferred to new courses without labels and provides a substantial AUC recovery rate. This method is effective in the sense that it works across different courses varying in lengths, format and contents and does not require custom feature engineering or additional target-course data or labels. Encouraged by this result, many future directions are feasible to explore. One potential direction is to look into a GritNet pretraining across more diverse courses. Given a pretrained GritNet and fine-tuning on the target course, it would boost performance further. We hope that our results will catalyze new developments of transfer learning for the real-time student performance prediction problem.

²Since the true binary target label (1: graduate, 0: not graduate) is imbalanced (i.e., number of 0s outweighs number of 1s), accuracy is not an appropriate metric.

REFERENCES

- Yoshua Bengio, Réjean Ducharme, and Pascal Vincent. A neural probabilistic language model. In *Advances in Neural Information Processing Systems 13 (NIPS 2000)*, pp. 932–938, 2001.
- Sebastien Boyer and Kalyan Veeramachaneni. Transfer learning for predictive models in massive open online courses. In *16th International Conference on Artificial Intelligence in Education (AIED 2015)*, pp. 54–63. Springer, 2015.
- Isaac Chuang and Andrew Ho. HarvardX and MITx: Four years of open online courses, Fall 2012–Summer 2016. *SSRN Electronic Journal*, 2016.
- Ronan Collobert and Jason Weston. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *25th International Conference on Machine Learning (ICML’08)*, pp. 160–167. ACM, 2008.
- Fisnik Dalipi, Ali Shariq Imran, and Zenun Kastrati. MOOC dropout prediction using machine learning techniques: Review and research challenges. In *9th IEEE Global Engineering Education Conference (EDUCON 2018)*, 2018.
- Economist. Equipping people to stay ahead of technological change - Lifelong learning. *A Special Report On Lifelong Learning: How to Survive in the Age of Automation*, 2017.
- Alex Graves and Jürgen Schmidhuber. Framewise phoneme classification with bidirectional LSTM networks. In *2005 International Joint Conference on Neural Networks (ICJNN’05)*, pp. 23–43, 2005.
- Byung-Hak Kim, Ethan Vizitei, and Varun Ganapathi. GritNet: Student performance prediction with deep learning. In *11th International Conference on Educational Data Mining (EDM 2018)*, pp. 625–629, 2018.
- Fei Mi and Dit-Yan Yeung. Temporal models for predicting student dropout in massive open online courses. In *Proceedings of 15th IEEE International Conference on Data Mining Workshop (ICDMW 2015)*, pp. 256–263, Atlantic City, New Jersey, 2015.
- Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359, 2010.
- Chris Piech, Jonathan Bassen, Jonathan Huang, Surya Ganguli, Mehran Sahami, Leonidas Guibas, and Jascha Sohl-Dickstein. Deep knowledge tracing. In *Advances in Neural Information Processing Systems 28 (NIPS 2015)*, pp. 505–513, 2015.
- Wei Wang, Han Yu, and Chunyan Miao. Deep model for dropout prediction in MOOCs. In *Proceedings of the 2nd International Conference on Crowd Science and Engineering (ICCSE 2017)*, pp. 26–32, Beijing, China, 2017.
- Jacob Whitehill, Kiran Mohan, Daniel Seaton, Yigal Rosen, and Dustin Tingley. Delving deeper into MOOC student dropout prediction. *arXiv preprint arXiv:1702.06404*, 2017a.
- Jacob Whitehill, Kiran Mohan, Daniel Seaton, Yigal Rosen, and Dustin Tingley. MOOC dropout prediction: How to measure accuracy? In *Proceedings of the Fourth ACM Conference on Learning @ Scale (L@S 2017)*, pp. 161–164, Cambridge, Massachusetts, USA, 2017b.