

# Towards an Integrative Educational Recommender for Lifelong Learners (Student Abstract)

Sahan Bulathwela, María Pérez-Ortiz, Emine Yilmaz, John Shawe-Taylor

Department of Computer Science, University College London

Gower Street, London WC1E 6BT, UK

{m.bulathwela, maria.perez, emine.yilmaz, j.shawe-taylor}@ucl.ac.uk

## Abstract

One of the most ambitious use cases of computer-assisted learning is to build a recommendation system for lifelong learning. Most recommender algorithms exploit similarities between content and users, overseeing the necessity to leverage sensible learning trajectories for the learner. Lifelong learning thus presents unique challenges, requiring scalable and transparent models that can account for learner knowledge and content novelty simultaneously, while also retaining accurate learners representations for long periods of time. We attempt to build a novel educational recommender, that relies on an integrative approach combining multiple drivers of learners engagement. Our first step towards this goal is TrueLearn, which models content novelty and background knowledge of learners and achieves promising performance while retaining a human interpretable learner model.

## Introduction

As the world population grows, more innovative approaches should be sought to provide high quality lifelong learning education opportunities to people of diverse cultures, languages, age groups and backgrounds. Machine learning now promises to provide such benefits of personalised teaching to anyone in the world cost effectively.

Since learner engagement is a prerequisite for achieving impactful learning outcomes (Lan et al. 2017), we attempt to build a recommender system that models different drivers of engagement, assisting learners on their *personal learning trajectory* to achieve their learning goals. Our approach differs from previous work in that it (i) incorporates different drivers of engagement such as resource quality, novelty, learner knowledge and interests; (ii) matches learners to useful and *engaging fragments* of knowledge, as opposed to lengthy full resources; and (iii) supports a multi-lingual and multi-modal collection of learning resources.

## Related Work

Contrary to conventional recommender systems, educational recommenders face different challenges, that stem from attempting to bring learners closer to their goals effectively.

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

These comprise of: i) Detecting learners interests and goals, as these can significantly affect their motivation (Salehi, Nakhai Kamalabadi, and Ghaznavi Ghouschi 2014); ii) Detecting the knowledge state of learners, the topics covered in a resource and the prerequisites necessary for benefiting from a learning material; iii) recommending novel and relevant educational resources; and iv) accounting for how different resource factors impact how engaging a resource is to the general population (Guo, Kim, and Rubin 2014).

The majority of work in adaptive educational systems builds on Item Response Theory (IRT) (Herbrich, Minka, and Graepel 2007) and Knowledge Tracing (KT) (Yudelson, Koedinger, and Gordon 2013) that focus on estimating learner’s knowledge for a narrow set of skills based on test answers. The work focusing on modelling a wide spectrum of skills over longer periods of time is surprisingly scarce.

While excelling on the personalisation front, there are other features that are often overlooked when designing educational recommendation systems. We design our system with these features in mind: (i) Cross-modality (e.g. video, text, audio etc.) and (ii) cross-linguality are vital to identifying and recommending educational resources across different modalities and languages. (iii) Transparency empowers learners by building trust while supporting the learner’s metacognition processes, such as planning, monitoring and reflection (Bull and Kay 2016). (iv) Scalability and (v) data efficiency allows maintaining the states of large masses of learners over longer periods of time and making best use of available user signals, such as implicit engagement (Salehi, Nakhai Kamalabadi, and Ghaznavi Ghouschi 2014).

## Our Approach

We identify four factors that influence learners’ engagement and develop a probabilistic graphical model that aims to recover those hidden variables using implicit engagement signals. Using a graphical model that learns from implicit engagement allows us to infer these hidden variables without compromising learner experience through explicit user interventions. The identified factors are: i) baseline resource quality ( $Q$ ), how engaging a resource is for the average learner; ii) background knowledge of the learner ( $B$ ); iii) novelty of the learning material ( $N$ ); and iv) curiosity or

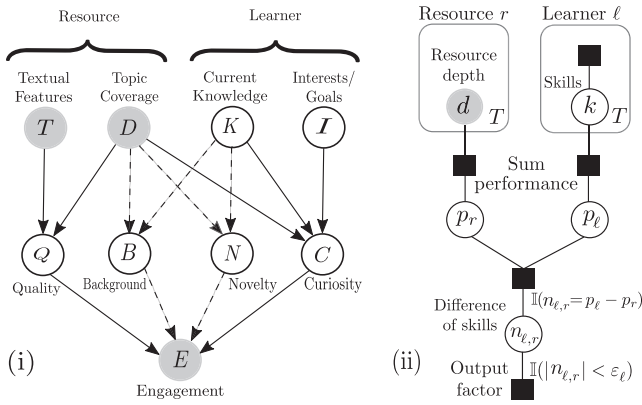


Figure 1: (i) Graphical model representing learner engagement (dashed arrows indicating the components tested) and (ii) TrueLearn factor graph, integrating resource topics ( $d$ ), current knowledge ( $k$ ) and novelty ( $n$ ) to predict engagement (output factor).  $\varepsilon_\ell$  is a dynamic factor of learner  $\ell$  indicating the engagement margin with respect to the amount of novelty. Plates represent  $T$  top ranked Wikipedia topics.

learning goals ( $C$ ) as outlined in Figure 1. As a first step, we reformulate the IRT TrueSkill algorithm (Herbrich, Minka, and Graepel 2007), to model learner knowledge and novelty as a function of engagement (dashed arrows in Figure 1 (i)). TrueSkill has several features that make it an excellent starting point. It is a scalable online method that shares similarities with our problem and provides a good framework for embedding novelty and a dynamic learner factor. We name the reformulation *TrueLearn* (Bulathwela et al. 2020).

**Data:** We construct a dataset from the popular video lectures repository VideoLectures.Net (VLN). Since handcrafting the *Knowledge Components (KCs)* in a resource is not scalable, we use an automatic entity linking algorithm, known as Wikification (Brank, Leban, and Grobelnik 2017). The English transcription of the lecture (or the English translation) is used to annotate the lecture with the 5 most relevant knowledge components using a Wikipedia text ontology through Wikifier (Brank, Leban, and Grobelnik 2017). This allows us to work with multiple languages and modalities and automatise the extraction of KCs. We divide the lecture text into multiple fragments of approximately 5,000 characters (equivalent roughly to 5 minutes of lecture) before Wikification. The engagement label is computed by calculating the normalised watch time (Guo, Kim, and Rubin 2014). The final dataset consists of 18,933 unique learners.

**Models:** We implement four baseline models to compare against: i) *Naïve persistence*, which assumes a static behaviour for all users, i.e. if the learner is engaged, they will remain engaged and vice versa; ii) *Naïve majority*, which predicts future engagement based solely on mean past engagement of users; iii) KT model (*Multi-Skill KT*) according to (Bishop, Winn, and Diethe 2015); and iv) *Vanilla TrueSkill* (Herbrich, Minka, and Graepel 2007).

Table 1: Mean F1-Score with the full VLN dataset

Algorithm	F1-Score
Naïve persistence	0.629
Naïve majority	0.640
Vanilla TrueSkill	0.400
Multi skill KT	0.259
TrueLearn	<b>0.677</b>

**Conclusions:** The results in Table 1 show evidence that *TrueLearn* outperforms the baselines while retaining a transparent learner model. The model is run per learner and trained in an online fashion, thus being scalable. The next step is to model content quality and learner curiosity within the same framework. Exploration into future user interfaces for learning with lecture fragments and ways of plan learning trajectories and recommendng material are also timely.

## Acknowledgments

This research is funded from the EU’s Horizon 2020 research and innovation programme under grant 761758 and partially funded by the EPSRC Fellowship titled ”Task Based Information Retrieval”, under grant EP/P024289/1.

## References

- Bishop, C.; Winn, J.; and Diethe, T. 2015. *Model-Based Machine Learning*. <http://www.mmlbook.com/>. Accessed: 2019-05-23.
- Brank, J.; Leban, G.; and Grobelnik, M. 2017. Annotating documents with relevant wikipedia concepts. In *Proc. of Slovenian KDD Conf. on SiKDD*.
- Bulathwela, S.; Perez-Ortiz, M.; Yilmaz, E.; and Shawe-Taylor, J. 2020. Truelearn: A family of bayesian algorithms to match lifelong learners to open educational resources. In *Proc. of the 2020 AAAI Conf. on Artificial Intelligence*.
- Bull, S., and Kay, J. 2016. Smili: a framework for interfaces to learning data in open learner models, learning analytics and related fields. *IJAIED* 26(1):293–331.
- Guo, P. J.; Kim, J.; and Rubin, R. 2014. How video production affects student engagement: An empirical study of mooc videos. In *Proc. of the First ACM Conf. on L@S*.
- Herbrich, R.; Minka, T.; and Graepel, T. 2007. Trueskill(tm): A bayesian skill rating system. In *Advances in NIPS 20*, 569–576.
- Lan, A. S.; Brinton, C. G.; Yang, T.-Y.; and Chiang, M. 2017. Behavior-based latent variable model for learner engagement. In *Proc. of Int. Conf. on EDM*.
- Salehi, M.; Nakhai Kamalabadi, I.; and Ghaznavi Ghouschi, M. B. 2014. Personalized recommendation of learning material using sequential pattern mining and attribute based collaborative filtering. *Education and Information Technologies* 19(4):713–735.
- Yudelson, M. V.; Koedinger, K. R.; and Gordon, G. J. 2013. Individualized bayesian knowledge tracing models. In *IJAIED*, 171–180. Springer.