Semantic TrueLearn: Using Semantic Knowledge Graphs in Recommendation Systems

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

Centre for Artificial Intelligence, University College London (UK)

In informational recommenders, many challenges arise from the need to handle the semantic and hierarchical structure between knowledge areas. This work aims to advance towards building a state-aware educational recommendation system that incorporates semantic relatedness between knowledge topics, propagating latent information across semantically related topics. We introduce a novel learner model that exploits this semantic relatedness between knowledge components in learning resources using the Wikipedia link graph, with the aim to better predict learner engagement and latent knowledge in a lifelong learning scenario. In this sense, *Semantic TrueLearn* builds a humanly intuitive knowledge representation while leveraging Bayesian machine learning to improve the predictive performance of educational engagement. Our experiments with a large dataset demonstrate that this new semantic version of TrueLearn algorithm achieves statistically significant improvements in terms of predictive performance with a simple extension that adds semantic awareness to the model.

1 Introduction

Developing artificial intelligence systems that, mildly at least, understand the structure of knowledge is foundational to building an effective recommendation system for education [4, 15], as well as for many other applications [18, 30] related to knowledge management and tracing.

Through this work, we propose *Semantic* TrueLearn, a novel and transparent learner model that incorporates automatic entity linking and Wikipedia (a publicly available, humanly-intuitive, domain-agnostic and ever-evolving) knowledge graph, as a first step towards building an educational recommender that automatically labels materials and embeds the structure of universal knowledge using Wikipedia. Semantic TrueLearn is

a probabilistic graphical model that maintains a symbolic representation of learners' knowledge that allows explanations, rationalisations and scrutinisation. We i) propose a novel approach for modelling semantic relatedness, ii) propose a novel sub-symbolic Bayesian learner model, iii) identify several research questions relating to validating the improvement of this proposed model and iv) evaluate the performance of the proposed model in a large dataset of learners engagement with educational resources.

2 Related work

Knowledge Tracing (KT) [31] is one of the most popular methods for user modelling in educational recommendation contexts. Incorporation of Semantic Relatedness in KT systems has been attempted before in prerequisite modelling [11, 12], exercise similarity [20, 14, 19] and other tasks [4, 29]. However, KT often relies on expert labelling of the *Knowledge Components* (KCs) [27] (sometimes also for knowledge hierarchies [4]), which is not scalable to large-scale lifelong learning applications in practice. *Wikification*, a form of entity linking [5] that expands this idea, has shown substantial progress and great promise for automatically capturing the KCs covered in an educational resource. Educational Recommenders using Wikification, such as TrueLearn [8], has shown promise in building probabilistic graphical models, although one of the assumptions of this model is that Wikipedia concepts are independent and thus unrelated.

Using Wikpedia as an ontology or knowledge graph to understand documents is not a new idea. While Wikipedia itself has been used as an ontology using page link and category links to describe "relates to" and "is type of" relationships respectively [28, 16], other works have pushed further and used the wealth of information in Wikipedia to build downstream knowledge bases and ontologies [1], as well as ontology-driven information retrieval systems [13]. From the early days of Wikipedia, exploiting different aspects of its contents (such as text, link structure, etc.) to model *Semantic Relatedness* (SR) that represents "relates to" links has been attempted. These SR metrics have evolved over time into recent proposals that are diverse and sophisticated metrics highly predictive of concept relatedness [25]. However, the utility in these proposals with graphical models is yet to be explored and is addressed thorugh this work.

Representations from Graphs   The core technical contribution of this work is proposing a method to infer the latent value of an unobserved skill parameter using observed ones via information sharing based on a semantic relatedness graph. Several works have proposed novel ways to use a relationship graph to recover a latent representation for an unknown node using a set of known nodes. Recent Graph Convolutional Neural Networks [17] infer hidden node embeddings ($H^{\ell+1}$) by weighted averaging the embeddings of its neighbours using adjacency matrix A and diagonal D ($H^{\ell+1} = D^{-\frac{1}{2}}AD^{-\frac{1}{2}}H^{\ell}W^{\ell}$). The popularising attention mechanism [2] uses alignment, which is a weighted sum of embeddings. This work puts the foundations to applying these approaches in a educational recommender using 1) a probabilistic graphical model and 2) SR values extracted from Wikipedia.

3 Methodology

Problem Formulation   Consider a learning environment in which a learner ℓ interacts with a set of educational resources $S_\ell \subset \{r_1, \ldots, r_Q\}$ over a period of $T = (1, \ldots, t)$ time steps, Q being the total of resources in the system. Each resource r_i is characterised by the top KCs or topics covered $K_{r_i} \subset \{1, \ldots, N\}$ (N being the total of KCs considered by the system) and the depth of coverage d_{r_i} of those. We represent learner knowledge at time t as a multivariate Gaussian distribution $\boldsymbol{\theta}_\ell^t \sim \mathcal{N}(\boldsymbol{\mu}_\ell^t, \Sigma_\ell^t)$, $\boldsymbol{\mu}_\ell^t \in \mathbb{R}^Q$ being the mean of knowledge and Σ_ℓ^t the covariance matrix. TrueLearn assumed that Σ is a diagonal covariance matrix in all cases and thus knowledge topics are completely independent from each other. The work in this paper builds towards considering a full covariance matrix, assuming that ρ_{ij} (estimated semantic relatedness) is a proxy for Σ_{ij} for topics i and j when $i \neq j$.

The key idea behind TrueLearn [8] is to model the probability of engagement $e_{\ell,r_i}^t \in \{1, -1\}$ between learner ℓ and resource r_i at time t as a function of the learner skills/ knowledge $\boldsymbol{\theta^t}_\ell$ and resource representation d_{r_i} for the top KCs covered K_{r_i}. When a new learner joins the recommender system, TrueLearn sets $\mu_\ell^0 = 0$ and $\Sigma_{ii} = \beta$, where β is a hyperparameter of the system, and $\Sigma_{ij} = 0$, $i \neq j$. Then, when the learner consumes an educational video fragment, TrueLearn updates the learner model/skills accordingly. Every skill that is not updated is set to the value from the last step, meaning at time t there might be many unobserved skills, specially given the amount of topics considered by the system (equal to the number of Wikipedia pages). Thus, TrueLearn assumes that the skill for topics in K_{r_i} can only be obtained through those topics and not semantically related ones.

3.1 Semantic TrueLearn

Extending the TrueLearn model [8, 7], the proposed Semantic TrueLearn, is a learner model that infers the knowledge state of learners in an online fashion. Semantic TrueLearn exploits its current knowledge of observed concepts and their SR to a novel concept to make a better prior estimation as per Figure 1.

Incorporating Semantic Relatedness Between Topics   The main assumption for Semantic TrueLearn, the proposal, is that knowledge can be shared across semantically related topics. By taking inspiration from graph convolutions [17], we formally assume a relationship as per equation 1.

$$\theta_{\ell,i}^t = \frac{1}{|\Omega_{\ell,i}|} \sum_{j \in \Omega_{\ell,i}} \gamma_{ij} \cdot \theta_{\ell,j}^t, \text{ where } \theta_{\ell,i}^t \sim \mathcal{N}(\mu, \sigma^2) \tag{1}$$

where Ω_i represents the set of topics used to infer the representation of topic i (e.g. most semantically related seen topics), where $i \neq j$. The mixing factors γ_{ij} can be set to semantic relatedness ρ_{ij} or to a factor of the standard error of topic j (meaning most observed topics are used). In the TrueLearn [8] model, which we extend, θ is a Gaussian variable. We use equation 2 to calculate unknown parameter $\theta_{\ell,i}^t$.
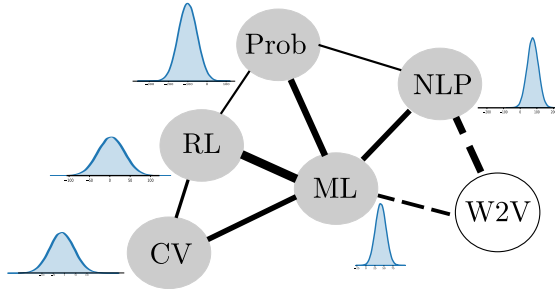
Figure 1: Inferring the knowledge for unseen topic (white circle) based on semantically related and seen ones (grey circles) by transferring knowledge (dotted lines). Topics are: ML (Machine Learning), RL (Reinforcement Learning), Prob (Probability), CV (Computer Vision), NLP (Natural Language Processing) and W2V (Word2Vec).

$$\hat{\theta}_{\ell,i}^{t} \sim \mathcal{N} \left(\sum_{j \in \Omega_i} \frac{1}{|\Omega_i|} \cdot \rho_{ij} \cdot \mu_{\ell,j}^{t}, \sum_{j \in \Omega_i} \left(\frac{1}{|\Omega_i|} \cdot \rho_{ij} \right)^2 \cdot \left(\sigma_{\ell,i}^{t} \right)^2 \right) \tag{2}$$

Semantic Relatedness Metric (SR Metric)   As mentioned in section 2, different measures of SR for Wikipedia concepts exist [25]. We empirically evaluate if the predictive performance of an educational recommender can be improved by incorporating 7 different SR Metrics to substitute ρ_{ij} in equation 2. We devise Milne and Witten (`M&W`), Entity Embeddings (`W2V`), Point-wise Mutual Information (`PMI`), Language Model based (`LM`), Jaccard Similarity (`Jaccard`), Conditional Probability (`CP`) and Barabasi and Albert (`BA`) SR Metrics, where SR values are pre-computed and publicly available [23].

4   Experiments and Results

We have four main research questions we verify in this preliminary experiment.

- **RQ1**: Which SR Metric is most suitable?

- **RQ2**: How many related topics should be considered?

- **RQ3**: Does Semantic TrueLearn outperform TrueLearn?

- **RQ4**: Does semantic information contribute to the gains? In which cases?

4.1 Data

We use the PEEK dataset [10], a dataset of more than 20,000 learners consuming video lectures in VideoLectures.Net[1] platform. The dataset provides information about how different users consumed fragments of videos over time [9]. This dataset uses entity linking [5] to associate most related Wikipedia concepts to documents. TagMe WAT API [24] [2] provides the required SR annotations.

To keep the computational complexity lower, a smaller dataset of 20 most active users is used for RQ1 and RQ2. However, the full dataset of 20,000 users is used to validate RQ3 and RQ4, which are our primary research questions.

4.2 Evaluation

A sequential prediction design where engagement at time t is predicted using events 1 to $t-1$. A training set of 70% of the learners is used for hyperparameter tuning and the remainder is used for testing and reporting. Being a binary classification task, precision, recall and F1-measure are evaluated whereas F1-measure is used for overall model selection [10]. The evaluation metrics are computed for each learner separately and the weighted average of the scores based on the number of learner's events is reported. To verify statistical significance of the improvement in RQ3, we use a learner-wise one-tailed paired t-test.

4.2.1 Impact of Semantic Relatedness (RQ4)

We use the topics encountered in user sessions to build a topic relatedness graph and extract a few attributes linked to graph connectedness for each user. Spearman's Rank Order Correlation Coefficient (SROCC) statistic is then used to evaluate the correlation between the extracted features and the predictive performance. User's *number of events*, *number of unique topics*, *topic sparsity rate* [8], *positive label rate*, *Avg. Connectedness*, i.e. average of the degree distribution of the topics, and *Min. Cut Set Size*, i.e. the minimum number of topics that need to be removed to break the graph into more sub-graphs, are analysed. The correlation with the recall score is investigated as the improvement in recall attributes to the performance gains of the proposed model (see Table 1). To validate if semantic relatedness is specifically influential in earlier parts of the user session, we plot the mean recall score of all users at event n, for different number of events (n).

4.3 Results

We run experiments to answer the research questions outlined above. To identify the most suitable SR metric (RQ1), we evaluate Semantic TrueLearn model using 7 SR Metrics proposed in section 3.1. The results are outlined in Table 1. To understand the Effect of Ω_i, the Number of Semantically Related Topics (RQ2), we use the identified SR

---

[1] www.videolectures.net
[2] https://sobigdata.d4science.org/web/tagme/wat-api

Table 1: Predictive performance of adding Semantic Relatedness (SR) to TrueLearn Novel algorithm. The different configurations (SR Metric) of Semantic TrueLearn Novel algorithm (our proposal) are evaluated using Precision (Prec.), Recall (Rec.) and F1 Score (F1). The most performant value and the next best value are highlighted in **bold** and *italic* faces respectively. The Semantic TrueLearn algorithms that outperform baseline model in terms of F1 score are underlined.

Algorithm	SR Metric	Prec.	Rec.	F1
TrueLearn Novel	-	0.7667	0.9476	0.8348
	M&W	0.7701	0.9469	*0.8364*
	W2V	0.**7714**	0.9467	**0.8370**
Semantic	PMI	0.7682	*0.9480*	0.8355
TrueLearn	LM	0.7605	**0.9507**	0.8322
Novel	Jaccard	0.7605	**0.9507**	0.8322
	CP	0.7621	**0.9507**	0.8330
	BA	*0.7704*	0.9469	*0.8364*

Metric to experiment with different numbers of semantically related topics. The results of this experiment are reported in Table 2. Finally, we use the full PEEK dataset to validate if the use of SR data improves baseline TrueLearn model (RQ3). The results obtained in this experiment are presented in Table 3. Figure 2 presents the results obtained in investigating the impact of semantic relatedness (RQ4) where (left) the correlation investigation between topic connectivity of users and recall score, and (right) the performance of the model based on different number of events is reported.

Table 2: The performance of Semantic TrueLearn model with W2V SR metric is reported in terms of Precision (Prec.), Recall (Rec.) and F1 Score (F1). The performance of the model is reported when different $\Omega_{\ell,i}$ top semantically related topics are utilised in equation 1. The most performant value and the next best value are highlighted in **bold** and *italic* faces respectively.

Number of Topics ($\Omega_{\ell,i}$)	Prec.	Rec.	F1
Most Related Topic	**0.7717**	0.9431	*0.8359*
3 Most Related Topics	0.7622	*0.9486*	0.8325
5 Most Related Topics	0.7659	**0.9490**	0.8345
10 Most Related Topics	0.7654	**0.9490**	0.8342
All Related Topics	*0.7714*	0.9467	**0.8370**

4.4 Discussion

It is evident from Table 1 that incorporating semantic relatedness leads to improvements in overall F1 score in majority of the SR metrics beating the baseline TrueLearn algorithm. Four Semantic TrueLearn models (ones that use M&W, W2V, PMI and BA) tend

Table 3: Predictive performance of Semantic TrueLearn model (our proposal) using Precision (Prec.), Recall (Rec.) and F1 Score (F1). The most performant value is highlighted in **bold** face. The Semantic TrueLearn Model that outperform baseline model ($p < 0.01$ in a one-tailed paired t-test) are marked with $\cdot^{(*)}$.

	Prec.	Rec.	F1
TrueLearn Novel	**0.5829**	0.7924	0.6471
Semantic TrueLearn	0.5711	**0.8563**$^{(*)}$	**0.6512**$^{(*)}$

to outperform the baseline TrueLearn Novel model in terms of precision and F1. The remainder seem to lead in recall. Entity embedding-based SR metric (`W2V`) leads to the best performing model. This is expected as neural-based semantic relatedness measures often outperform their graph-based counterparts [25]. Our empirical results in Table 2 also showed that using *all* semantically related topics gives best results contrary to restricting to some of them. Finally, Table 3 shows the superiority of Semantic TrueLearn in comparison to the baseline model that does not exploit SR information from Wikipedia. This is a clear indication that a knowledge base such as Wikipedia can be critical to improving the assumptions used for a learner, modelled using a probabilistic graphical models in the education context. Semantic relatedness can be truly valuable in early stages of the user session when the interaction data about the user is limited, thus addressing the cold-start problem.

The correlation evaluation presented in Figure 2 (left) shows the lack of correlation between Positive Label Rate and recall score across both TrueLearn models. Although it
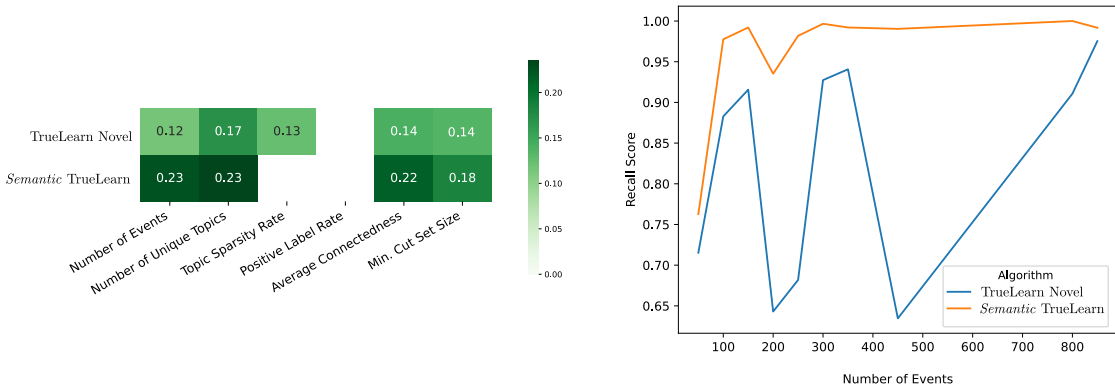

Figure 2: (Left) Relationship between different behavioural characteristics of user profiles and model recall performance presented using SROCC. The numbers and the intensity of each cell corresponds to the Spearman R coefficient where significant correlation is present ($p < 0.01$). Empty cells represent lack of significant correlation.(Right) The average recall performance of the two models for the learner population at different number of events.

has been demonstrated by the original authors that TrueLearn algorithm capitalises on recall, there is no information in the work regarding the positive label rate in the datasets. This observation confirms that TrueLearn family of algorithms find true patterns in learner data rather than merely capitalising on the positive labels to boost performance.

Multiple observations in Figure 2 (left) give evidence of the superiority of Semantic TrueLearn exploiting the semantic relatedness between topics to boost recall. The main two observations are the new model's stronger Spearman's rank correlation with learner *Avg. Connectedness* and *Min. Cut Set Size*. This is a strong indication that the Semantic TruLearn model is exploiting the topic correlations. The correlation between the number of events, number of unique topics and topic connectedness cause the higher correlation between these features and Semantic TrueLearn model. The Figure 2 (right) clearly shows how the recall score of predictions is much larger in Semantic TrueLearn algorithm regardless in early or latter stage of the learner session. Linking this to results in Table 1 shows that this impressive gain of recall score is achieved with a much smaller sacrifice of precision score.

Limitations   Amid the significant gains, we observe that most KCs encountered by the model in a session are highly correlated to each other. This leads to overlapping information being propagated repeatedly when using equation 1 which may lead to over-estimation of knowledge of unseen KCs. Equation 1 only acknowledges the existence of correlations between seen and unseen topics (dotted lines in Figure 1) not the correlation within seen topics (solid lines in Figure 1) which is an issue. As the proposed method primarily infers *unobserved* skills, its use diminishes over time when the user session matures (as new topics are encountered less often). Mechanisms to keep using semantic awareness to refine representations is a much needed improvement to the proposed method.

5 Conclusions

Leveraging semantic relatedness between Wikipedia topics has demonstrated promise to improving the predictive performance of informational recommenders such as TrueLearn that are built on Wikipedia ontology and probabilistic graphical models. In addition, we identify that restricting the number of related topics leads to degraded performance, suggesting the use of all available knowledge components extracted from Wikification. Our analysis also shows that topic connectedness within learner sessions is positively correlated with the performance gains of Semantic TrueLearn, giving clearer evidence of the positive impact of incorporating this aspect when modelling learners and their journey within an education setting.

The proposed model is a stepping stone to accounting for semantic relatedness. However, it still disregards the correlation among the observed topics. To address this, we propose the following research avenues: 1) using algorithms such as PageRank [6] to derive uncorrelated skill parameters, 2) accounting for inter-skill correlation in equation 2 and 3) building a hierarchical representation of knowledge [21] consisting

of mutually exclusive (uncorrelated) Wikipedia concepts. It may also be fruitful to consider richer ontologies [1] that contain more fine-grained relationships, entity definitions/categorisations and constrains in the place of raw Wikipedia graph to incorporate finer grain semantic-awareness to the learner model. Mechanisms to continuously utilise SR information should be identified and investigated in future work.

Moreover, semantic relatedness measures are not usually built and validated with educational datasets or topics, which is a limitation. In the future, we also aim to validate the SR metrics with education and informational recommender focused datasets. As Semantic TrueLearn builds a sub-symbolic representation that is humanly-intuitive, it is possible to create narratives and intelligent user interfaces (e.g. [9, 22]) used to interpret and rationalise [26] the learnings of the model leading towards more *human-in-the-loop* artificial intelligence that allows verification and scrutinisation of the models [3].

6 Acknowledgments

Bibliography

[1] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. Dbpedia: A nucleus for a web of open data. In *The semantic web*. Springer, 2007.

[2] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.

[3] Krisztian Balog, Filip Radlinski, and Shushan Arakelyan. Transparent, scrutable and explainable user models for personalized recommendation. In *Proc. of the 42nd Int. ACM SIGIR Conf. on Research and Development in Information Retrieval (SIGIR '19)*, 2019.

[4] Konstantin Bauman and Alexander Tuzhilin. Recommending remedial learning materials to students by filling their knowledge gaps. *MIS Quarterly*, 42(1):313–332, 2018.

[5] Janez Brank, Gregor Leban, and Marko Grobelnik. Annotating documents with relevant wikipedia concepts. In *Proc. of Slovenian KDD Conf. on Data Mining and Data Warehouses (SiKDD)*, 2017.

[6] Sergey Brin and Lawrence Page. The anatomy of a large-scale hypertextual web search engine. In *Proc. of Int. Conf. on World Wide Web*, 1998.

[7] S. Bulathwela, M. Perez-Ortiz, E. Yilmaz, and J. Shawe-Taylor. Towards an integrative educational recommender for lifelong learners. In *AAAI Conference on Artificial Intelligence*, 2020.

[8] S. Bulathwela, M. Perez-Ortiz, E. Yilmaz, and J. Shawe-Taylor. Truelearn: A family of bayesian algorithms to match lifelong learners to open educational resources. In *AAAI Conference on Artificial Intelligence*, 2020.

[9] Sahan Bulathwela, Stefan Kreitmayer, and María Pérez-Ortiz. What's in it for me? augmenting recommended learning resources with navigable annotations. In *Proc. of Int. Conf. on Intelligent User Interfaces Companion*, IUI '20, page 114–115, 2020.

[10] Sahan Bulathwela, Maria Perez-Ortiz, Erik Novak, Emine Yilmaz, and John Shawe-Taylor. Peek: A large dataset of learner engagement with educational videos, 2021.

[11] C. Carmona, E. Millán, J. L. Pérez-de-la Cruz, M. Trella, and R. Conejo. Introducing prerequisite relations in a multi-layered bayesian student model. In *Proc. of the Int. Conf. on User Modeling*, UM'05, page 347–356, 2005.

[12] Penghe Chen, Yu Lu, Vincent W Zheng, and Yang Pian. Prerequisite-driven deep knowledge tracing. In *2018 IEEE Int. Conf. on Data Mining (ICDM)*, pages 39–48. IEEE, 2018.

[13] Gregory Grefenstette and Karima Rafes. Transforming wikipedia into an ontology-based information retrieval search engine for local experts using a third-party taxonomy. *arXiv preprint arXiv:1511.01259*, 2015.

[14] Zhenya Huang, Yu Yin, Enhong Chen, Hui Xiong, Yu Su, Guoping Hu, et al. Ekt: Exercise-aware knowledge tracing for student performance prediction. *IEEE Transactions on Knowledge and Data Engineering*, 2019.

[15] Weijie Jiang, Zachary A. Pardos, and Qiang Wei. Goal-based course recommendation. In *Proceedings of International Conference on Learning Analytics & Knowledge*, 2019.

[16] Tokio Kawakami, Takeshi Morita, and Takahira Yamaguchi. Building wikipedia ontology with more semi-structured information resources. In *Semantic Technology - 7th Joint International Conference, JIST 2017, Proceedings*. Springer Verlag, 2017.

[17] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2017.

[18] Patrick Lewis, Ethan Perez, Aleksandara Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *arXiv preprint arXiv:2005.11401*, 2020.

[19] Hiromi Nakagawa, Yusuke Iwasawa, and Yutaka Matsuo. Graph-based knowledge tracing: Modeling student proficiency using graph neural network. In *IEEE/WIC/ACM International Conference on Web Intelligence*, WI '19, 2019.

[20] Shalini Pandey and Jaideep Srivastava. Rkt: Relation-aware self-attention for knowledge tracing. *arXiv preprint arXiv:2008.12736*, 2020.

[21] Radek Pelánek. Managing items and knowledge components: domain modeling in practice. *Educational Technology Research and Development*, 68(1):529–550, 2020.

[22] María Pérez-Ortiz, Claire Dormann, Yvonne Rogers, Sahan Bulathwela, Stefan Kreitmayer, Emine Yilmaz, Richard Noss, and John Shawe-Taylor. X5learn: A personalised learning companion at the intersection of ai and hci. In *26th International Conference on Intelligent User Interfaces*, pages 70–74, 2021.

[23] Francesco Piccinno. *Algorithms and data structures for big labeled graphs*. PhD thesis, Universitade Pisa, 2017.

[24] Francesco Piccinno and Paolo Ferragina. From tagme to wat: A new entity annotator. In *Proc. of the First Int. Workshop on Entity Recognition & Disambiguation*, ERD '14, 2014.

[25] Marco Ponza, Paolo Ferragina, and Soumen Chakrabarti. On computing entity relatedness in wikipedia, with applications. *Knowledge-Based Systems*, 188, 2020.

[26] Mark Owen Riedl and Vadim Bulitko. Interactive narrative: An intelligent systems approach. *Ai Magazine*, 34(1):67–67, 2013.

[27] Douglas Selent, Thanaporn Patikorn, and Neil Heffernan. Assistments dataset from multiple randomized controlled experiments. In *Proc. of the Conf. on Learning @ Scale*, 2016.

[28] Zareen Syed, Tim Finin, and Anupam Joshi. Wikipedia as an ontology for describing documents. In *Proc. of Int. Conf. on Weblogs and Social Media*. AAAI Press, March 2008.

[29] Khushboo Thaker, Lei Zhang, Daqing He, and Peter Brusilovsky. Recommending remedial readings using student knowledge state. In *Proc. of Int. Conf. on EDM*, 2020.

[30] Tae Yano and Moonyoung Kang. Taking advantage of wikipedia in natural language processing. Technical report, Carnegie Mellon University Language Technologies Institute, 2016.

[31] Michael V. Yudelson, Kenneth R. Koedinger, and Geoffrey J. Gordon. Individualized bayesian knowledge tracing models. In H. Chad Lane, Kalina Yacef, Jack Mostow, and Philip Pavlik, editors, *Proc. of Artificial Intelligence in Education*, 2013.