
Recommending Texts to Children with an Expert in the Loop

Maria Soledad Pera

People and Information Research Team
Dept. of Computer Science
Boise State University, ID, 83725, USA
solepera@boisestate.edu

Katherine Wright

Dept. of Literacy, Language, & Culture
Boise State University, ID, 83725, USA
katherinewright@boisestate.edu

Michael D. Ekstrand

People and Information Research Team
Dept. of Computer Science
Boise State University, ID, 83725, USA
mekstrand@boisestate.edu

Authors maintain the copyright.
Workshop on International and Interdisciplinary Perspectives on Children
and Recommender Systems (KidRec) at IDC 2018, June 19, 2018,
Trondheim, Norway.

Abstract

In this position paper we discuss a number of open problems we believe the community should address in order to enhance the recommendation task for children. We specifically outline algorithmic and evaluation limitations when it comes to recommending reading materials for children in the classroom setting. Furthermore, we focus on the need to involve an expert (e.g., teacher) as part of the recommendation process to better serve the population under study.

Author Keywords

Expert-in-the-loop; recommender systems

ACM Classification Keywords

• Information systems~Recommender systems • Social and professional topics~Children.

Introduction

Reading remains a foundational skill tied to the development and success of individuals. To champion literacy development, it is essential to encourage children to read beyond the time allotted in the classroom. Unfortunately, for that to happen, children need to be given the right material that not only matches their interests and preferences, but most importantly, they can comprehend. Recommender systems (*RS*) are a promising approach to this problem, but we argue that the state-of-the-art in this

area is not prepared to handle such a task due to resource and algorithmic limitations.

In the remaining of this paper, we outline the needs and limitations we believe have to be addressed in order to enhance RS functionality so that it is capable of handling such a task. We also discuss the central figure an educator plays in this domain, as children will rarely interact with the RS itself, and instead it is the expert (e.g., a teacher in a classroom setting) who will be given the recommendations; which in turn she will later leverage in advising individual students on which texts match their interests and abilities.

Limitations

We discuss below questions we believe should drive the enhancement of the state-of-the-art (both from an algorithmic and user perspectives) when it comes to making reading recommendations for children.

Why are RS algorithms based on explicit or implicit feedback not sufficient?

Sites such as LibraryThing or GoodReads allow users to rate books they read. This explicit user preference information is heavily leveraged by matrix factorization RS algorithms [6]. Unfortunately, based on privacy rules like the Children's Online Privacy Protection Act (COPPA), it is rare that this type of explicit information can be archived and shared for research and development purposes. As Ekstrand [3] observed, this presents challenges not just for the development of RS for children, but also their evaluation, as many common techniques are off the table.

Depending on the application, many of the restrictions surrounding explicit data also apply to implicit data.

Moreover, sites like BiblioNasium allow students to keep track of reading logs—historical data that could be used for personalization [8]. Unfortunately, this only reflects “completed” books, as opposed to the degree to which each child enjoyed the corresponding book.

Why are RS techniques for technology-enhanced learning not always applicable?

RS techniques for technology enhanced learning are meant to support learning practices of both individuals and organizations [7]. To our knowledge, these techniques mostly focus on a more traditional and mature audience, who can explicitly provide ratings/feedback, in turn facilitating development and assessment.

Why are group recommender algorithms not suitable?.

Group recommenders [5] would be one approach to produce recommendations for a class, but group recommenders typically focus on finding *one* item for the entire class; we see significant promise in recommendations that help teachers find a multitude of texts to meet the varied needs of individual students in the classroom. Group recommendation technology may be one piece of the solution, but will require adaptation to the classroom and expert-mediated setting.

Opportunities and Resources

The classroom text recommendation problem poses new challenges, but it also provides distinct resources that can be leveraged to improve students' engagement with texts.

Invested Experts

Teachers have significant expertise in both the general needs of educational texts and the particular needs of



Figure 1. We consider the expert an asset when it comes to recommendations: we argue in favor of including her as part of the recommendation process to facilitate the identification of the right text for the right reader.

their students and classrooms. Our vision is not to replace the teacher as the selector of student readings, but to make them more effective in their work (see Figure 1). This presents challenges — RS in settings where experts mediate the recommendations instead of presenting them directly to end users is not a well-studied problem and requires new developments in RS interfaces and explanations. At the same time, it is also a resource: by focusing on providing meaningful recommendations and decision support to teachers, who make the final judgement about what is useful and appropriate for their classrooms, we can solve a more tractable problem than directly finding the exact right texts for a group of students.

These experts can also provide more direct and useful feedback for evaluating and improving the system.

Structured Information Needs and Criteria

There is significant prior work in literacy and education research on what makes texts pedagogically useful and accessible to children [4]. Operationalizing these concepts for use as information filters is a challenge, but the fundamental concepts exist.

Educational settings, particularly in the United States, are also highly structured with respect to the topics and learning outcomes that need to be covered or attained. Standards such as Common Core and the Next Generation Science Standards provide a common reference point for identifying topics of importance.

Challenges

United States K-12 classrooms are becoming increasingly diverse, and teachers can expect to teach language learners, learners with disabilities, and

gifted/talented students in one classroom. This makes selecting appropriate texts challenging, as students bring a variety of skills, interests, and schemas to each lesson [4]. Teachers must be empowered, as experts in content pedagogy and their students, to retrieve appropriate texts to meet the needs of their diverse learners and engage all students in both content and literacy development. RS can help do so, but the fundamental premise (i.e., serve a direct user) must change, as during formative years the recommendations are for *the teacher*, but the *students* remain the major *stakeholders*. This will require a holistic perspective for coming up with solutions that match the requirements of this domain and, most importantly, are useful for the classroom setting.

This problem poses important human-computer interaction challenges that will expand our understanding of how to build recommendation and information retrieval interfaces. Supporting an expert decision-maker in selecting texts for a classroom will require novel explanation techniques for relating the text to the search query, curriculum needs, and the interests and abilities of a group of students.

Solutions may also require new developments in eliciting the information need, as classroom readings are a complex information need that cannot be expressed with a textual query alone. We will need to determine a balance that collects sufficient information to guide the search and identify useful texts without making the search process itself too cumbersome to be of benefit. Most theories of information-seeking behavior [2] focus on settings with a single user's information need, so there is limited literature to work from in addressing these needs. Another challenge

from an information retrieval perspective will be adapting the idea of relevance. In this particular case, the traditional notion of relevance with respect to the needs of an individual user or group of users will not be sufficient. Furthermore, the context, whether that be the class curriculum needs, grade level, reading abilities of the particular child, as well as teachers' expectation of what material is indeed suitable for the particular child will have to be factored into in relevance rankings that can inform the recommendation. Furthermore, techniques that can analyze reading materials from multiple perspectives (text complexity, topic suitability, etc.) while at the same time depending upon openly available resources (e.g., due to copyright, text samples are not always accessible when it comes to books) will also inform the recommendation process.

Acknowledgements

Work partially funded by a grant from the College of Education at Boise State University.

References

1. Gediminas Adomavicius and Alexander Tuzhilin. 2015. Context-aware recommender systems. In *Recommender systems handbook* (pp. 191-226). Springer, Boston, MA.
2. Donald O. Case and Lisa M. Given. 2016. *Looking for Information: A Survey of Research on Information Seeking, Needs, and Behavior*. Emerald Group Publishing.
3. Michael D. Ekstrand. 2017. Challenges in Evaluating Recommendations for Children. In *KidRec Workshop*.
4. Melissa L. Farrall. 2012. *Reading assessment: Linking language, literacy, and cognition*. John Wiley & Sons.

Recommendation technology will also need to be adapted to function well in this context. Most techniques optimize utility functions learned from user interactions, e.g. to rank things by user preference [9]; in this setting, however, the user is not the end beneficiary of the recommendations, so their opinion and actions is only one piece of the input needed for the algorithm to form a complete picture. Further, meeting the curricular and pedagogical needs of a classroom in a manner compatible with the abilities and interests of its students is a complex multi-objective optimization problem that blends the complexity of context-aware recommendation [1] with group recommendation [5] in a higher-dimensional setting. We expect solving this problem to be non-trivial.

5. Alexander Felfernig, Ludovico Boratto, Martin Stettinger, and Marko Tkalčič. 2018. *Group Recommender Systems: An Introduction*. Springer.
6. Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix Factorization Techniques for Recommender Systems. *Computer* 42, 8: 30–37.
7. Manouselis, N., Drachsler, H., Vuorikari, R., Hummel, H. and Koper, R., 2011. Recommender systems in technology enhanced learning. In *Recommender systems handbook* (pp. 387-415). Springer, Boston, MA.
8. Maria S. Pera and Yiu-Kai Ng. 2014. Automating readers' advisory to make book recommendations for k-12 readers. In *ACM RecSys*, pp. 9-16.
9. Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. In *UAI*, 452–461.