

# "It doesn't look good for a date": Transforming Critiques into Preferences for Conversational Recommendation Systems

Victor S. Bursztyn<sup>1</sup>, Jennifer Healey<sup>2</sup>, Nedim Lipka<sup>2</sup>,  
Eunyee Koh<sup>2</sup>, Doug Downey<sup>1,3</sup>, and Larry Birnbaum<sup>1</sup>

<sup>1</sup>Department of Computer Science, Northwestern University, Evanston, IL, USA

<sup>2</sup>Adobe Research, San Jose, CA, USA

<sup>3</sup>Allen Institute for Artificial Intelligence, Seattle, WA, USA

v-bursztyn@u.northwestern.edu, {jehealey, lipka, eunyee}@adobe.com  
{d-downey, l-birnbaum}@northwestern.edu

## Abstract

Conversations aimed at determining good recommendations are iterative in nature. People often express their preferences in terms of a critique of the current recommendation (e.g., "It doesn't look good for a date"), requiring some degree of common sense for a preference to be inferred. In this work, we present a method for transforming a user critique into a positive preference (e.g., "I prefer more romantic") in order to retrieve reviews pertaining to potentially better recommendations (e.g., "Perfect for a romantic dinner"). We leverage a large neural language model (LM) in a few-shot setting to perform critique-to-preference transformation, and we test two methods for retrieving recommendations: one that matches embeddings, and another that fine-tunes an LM for the task. We instantiate this approach in the restaurant domain and evaluate it using a new dataset of restaurant critiques. In an ablation study, we show that utilizing critique-to-preference transformation improves recommendations, and that there are at least three general cases that explain this improved performance.

## 1 Introduction

Conversational recommendation systems (CRSs) are dialog-based systems that aim to refine a set of options over multiple turns of a conversation, envisioning more natural interactions and better user modeling than in non-conversational approaches.

However, the resulting dialogs still do not necessarily reflect how real conversations unfold. Most CRSs fall into two categories: they either frame the problem as a slot-filling task within a predefined feature space, such as [Sun and Zhang \(2018\)](#); [Zhang et al. \(2018\)](#); [Budzianowski et al. \(2018\)](#), which is closer to how people make decisions but not as flexible as real conversations; or they elicit preferences by asking users to rate specific items, such as [Christakopoulou et al. \(2016\)](#), which is in-

dependent of a feature space but not as natural to users.

What's it about the menu that you didn't like?

I need something less expensive and more appropriate for kids

I hear you. You prefer a more affordable place for kids.

Check out **M Burger** for that. A previous customer wrote: "This is an affordable, kid-friendly place specialized in burgers and American food."

Figure 1: An example of our system transforming a critique into a positive preference and then using a customer testimonial to sell the user on a new option.

When we examine situations involving real human agents ([Lyu et al., 2021](#)), decisions typically require multiple rounds of recommendations by the agent and critiques by the user, with the agent continuously improving the recommendations based upon user preferences that can be inferred from such critiques.

These inferences can be compared to the types of common sense inferences that have been studied recently with LMs ([Davison et al., 2019](#); [Majumder et al., 2020](#); [Jiang et al., 2021](#)). However, use of LMs for critique interpretation remains underexplored, despite the important role of critiques in communicating preferences—a very natural real-world task. Working in the restaurant domain, we prompt GPT3 ([Brown et al., 2020](#)) to transform a free-form critique (e.g., "It doesn't look good for a date") into a positive preference (e.g., "I prefer more romantic") that better captures the user's needs. Compared with most previous work on common sense inference, which relies on manually-constructed question sets, our task presents an opportunity to study common sense inference within a naturally arising, real-world application.

We test the effect of our novel critique interpretation method on the quality of recommendations using two different methods: one that matches the embedding of an input statement (e.g., "I prefer

more romantic”) to persuasive arguments found in customer reviews (e.g., “Perfect for a romantic dinner”); and another one that fine-tunes BERT (Devlin et al., 2018) in using an input statement to rank a given set of arguments.

Our work differs from previous critiquing-based systems that strongly limit the types of critiques that can be used (Chen and Pu, 2012) and aligns with a recent trend in the CRS literature towards more open-ended interactions (Radlinski et al., 2019; Byrne et al., 2019). To the best of our knowledge, Penha and Hauff (2020) are the closest prior work investigating whether BERT can be used for recommendations by trying to infer related items and genres. Here, we focus specifically on critique-to-preference inferences, aiming at more natural dialogs and better recommendations.

Our contributions are the following: 1. We propose a critique interpretation method that does not limit the feature space a priori; 2. We demonstrate that transforming critiques into preferences improves recommendations over two fold when matching embeddings and by 19-59% when fine-tuning an LM to rank recommendations, and present three possible explanations for this; and 3. We release a new dataset of user critiques in the restaurant domain, contributing a new applied task where common sense has great practical value.

## 2 Methods

In this section, we describe three methods: A critique interpretation method (2.1), an embeddings-based recommender (2.2.1), and an LM-based recommender (2.2.2).

### 2.1 Critique Interpretation

Critique interpretation is the task of transforming a free-form critique into a positive preference. Our critique interpretation method uses GPT3 in a few-shot setting similarly to Brown et al. (2020), which can be represented in a 3-shot version as follows:

#### GPT3 Input :

*It looks cheap => I prefer a fancier place.  
Too expensive => I prefer a more affordable place.  
That's so tacky => I prefer a more stylish place.  
That's not good for a date => I prefer*

#### GPT3 Output :

*a more romantic place.*

To prime GPT3 for our task, we include ten examples in its prompt, five related to food and five to the atmosphere.<sup>1</sup> We then append the critique that we would like to transform followed by the string “*I prefer*”, which conditions GPT3 to generate a positive preference. In our experiments, positive preferences are sampled using OpenAI’s Completion API (the DaVinci model, temperature = 0.7, top p = 1.0, response length = 20, and no penalties).

Besides not requiring a hand-crafted feature set, this method is also capable of more flexible interpretation of language, such as transforming “How come they only serve that much?”—with no clearly negative words—into “I prefer larger portions.”

### 2.2 Content-based Recommendations

#### 2.2.1 Recommendation Search

Our embeddings-based recommender,  $f_{cos}$ , takes a preference statement and searches for persuasive arguments in customer reviews. As seen in Figure 1, we can define a persuasive argument as a review sentence that conveys clearly positive sentiment while being as specific as possible w.r.t the user’s preferences.

To incorporate this definition in  $f_{cos}$ , first we parse sentences in customer reviews using spaCy (Honnibal and Montani, 2017) and use EmoNet (Abdul-Mageed and Ungar, 2017) to keep the sentences with at least a minimum amount of “joy” ( $\geq 0.7$ ) as our set of argument candidates  $A$ .

Then we use the Universal Sentence Encoder (Cer et al., 2018) to calculate the similarity of all these argument candidates w.r.t a given user preference. We calculate the cosine similarity between their representations in this embedding space, select the argument with maximum alignment, and recommend the associated restaurant:

$$\begin{aligned} Sim(s_1, s_2) &= \cos(Enc(s_1), Enc(s_2)) \\ f_{cos}(\text{preference}) &= \underset{a \in A}{\text{argmax}} Sim(a, \text{preference}) \end{aligned}$$

As with critique interpretation,  $f_{cos}$  can take any natural language statement as input to search for potential recommendations. We denote  $f_{cos}^{pref}$  when it uses an inferred positive preference as input (“I prefer more romantic”) and  $f_{cos}^{crit}$  when it directly uses a critique (“It doesn’t look good for a date”). In our first ablation study, we use  $f_{cos}^{crit}$  as a baseline to test the efficacy of  $f_{cos}^{pref}$  in retrieving better recommendations.

<sup>1</sup>Fully available at <https://bit.ly/3fnf8v2>

### 2.2.2 Recommendation Ranking

Besides using pretrained embeddings to search for recommendations from customer reviews, we design a more computationally intensive method,  $f_{LM}$ , that fine-tunes BERT to rank a set of arguments  $A$  considering a given input statement.

We use the currently top performing open-source solution (Han et al., 2020; Pasumarthi et al., 2019) on the MSMARCO passage ranking leaderboard<sup>2</sup> to fine-tune three versions of BERT:  $f_{LM}^{pref}$  uses a positive preference as input (“I prefer more romantic”),  $f_{LM}^{crit}$  uses a critique (“It doesn’t look good for a date”), and  $f_{LM}^{concat}$  uses a concatenation of both a critique and a preference (“It doesn’t look good for a date. I prefer more romantic”). Hypothetically, the more powerful LM method could learn to satisfy the user’s preferences without the need of critique interpretation if the performances of  $f_{LM}^{crit} \approx f_{LM}^{pref} \approx f_{LM}^{concat}$ .

In our experiments, BERT-Base is fine-tuned for 10,000 steps, with learning rate =  $10^{-5}$ , maximum sequence length = 512, and softmax loss, using a Nvidia Quadro RTX 8000 for 3-6h per run (when ranking 15 and 30 arguments, respectively) and two runs per model (2-fold cross validation).

## 3 Evaluation

We run two ablation studies to evaluate the hypothesis that critique interpretation would be beneficial to the overall recommendation approach. First, we analyze our embeddings-based recommender,  $f_{cos}$ , to check whether the performance of  $f_{cos}^{pref} > f_{cos}^{crit}$ . Secondly, we analyze our LM fine tuning-based recommender,  $f_{LM}$ , to check if  $f_{LM}^{pref} > f_{LM}^{crit}$  or  $f_{LM}^{concat} > f_{LM}^{crit}$ . Finally, we discuss qualitative differences between the tested arms.

### 3.1 Data

Our methods were instantiated in a system comprising 15 restaurants selected from two of the largest metropolitan areas in the United States, covering a variety of price ranges and cuisines. For each restaurant, up to 100 four- or five-star customer reviews were collected from Google Places. This resulted in a total of 1455 reviews comprising 5744 sentences, 2865 of which pass the threshold for being identified as positive review sentences.

We compiled a set of user critiques from two sources: a set of 46 unique critiques from user

studies that were conducted to test an earlier system prototype (Burszty et al., 2021), and 294 additional critiques adapted from the Circa dataset (Louis et al., 2020). Circa was designed to study indirect answers to yes-no questions, such as “Are you a big meat eater?” answered with “I prefer leafy greens”, from which the critique “I’m not a big meat eater” can be generated. We end with a total of 340 individual critiques after examining 1205 similar examples.

We generated a positive preference for each individual critique using our critique interpretation method in 2.1, without discarding any critiques. Our method yielded accurate preferences for 298 critiques (87.6%). For the remaining 42, we found GPT3 mostly undecided and vague (e.g., “Jalapeños are my limit” generates “I prefer food without jalapeños”). In our experiments, for these edge cases, we kept the best of three trials, but we believe that results using just the first generation would have been qualitatively similar.

The 340 critiques were randomly combined into pairs and triples in order to simulate longer conversations, i.e., two- and three-round critiques. We sampled 340 pairs and 340 triples, substituting only exceptional combinations that contained contradictory statements (e.g., “I’m not a big meat eater.” paired with “I’m not in the mood for vegetables.”), for a total of 1020 critiques. Compound critiques were concatenated into single statements as well as their corresponding preferences.

This curated dataset of 1020 restaurant critiques and inferred preferences is made available to the research community.<sup>3</sup>

### 3.2 Measurements

For evaluating our embedding-based methods  $f_{cos}$ , we use critiques as input to  $f_{cos}^{crit}$  and their positive preferences as input to  $f_{cos}^{pref}$ . For each query we retrieve the top 3 arguments, which are labeled as accurate or inaccurate by a human judge (illustrated in Table 1). To measure labeling consistency, a second human annotator redundantly labeled a sample of 100 arguments resulting in a Cohen’s Kappa of 0.71, which indicates strong agreement.

We then measure Precision@1, Precision@2, and Precision@3 in Table 2 for the embeddings-based method with ( $f_{cos}^{pref}$ ) and without critique interpretation ( $f_{cos}^{crit}$ ).

<sup>2</sup><https://microsoft.github.io/msmarco/>

<sup>3</sup><https://github.com/vburszty/critique-to-preference-emnlp2021>

Test case	Positive preference	Without critique interpretation ( $f_{cos}^{crit}$ )			With critique interpretation ( $f_{cos}^{pref}$ )		
		Rank #1	Rank #2	Rank #3	Rank #1	Rank #2	Rank #3
It looks too casual.	I prefer a fancier place.	Very cheesy, very fresh!	Very kid friendly.	Awesome ambiance!	<b>Elegant, upscale and classy place for a special occasion.</b>	<b>The best restaurant around here.</b>	<b>Superior restaurant, the only place I will have a dim sum.</b>
It has a freaking band!	I prefer a more quiet place.	It has an awesome atmosphere.	It has an awesome atmosphere.	It has a great atmosphere.	<b>Excellent spot to spend time alone or talk business.</b>	Good ambiance.	Great place to be at night.
I don't really like seafood.	I prefer beef or chicken.	Everything delicious with an exception of the shrimps.	I found that I do not enjoy tuna, but my mom thought it was excellent.	For dinner, I enjoyed the scallops one night and the sea bass the second.	<b>I only eat Beef Brisket here because it is delicious!</b>	Chicken flautas are always delish.	Chicken moist and tender.

Table 1: Three test cases with the top 3 arguments from  $f_{cos}^{crit}$  and  $f_{cos}^{pref}$  (accurate marked in bold).

	Precision@1	Precision@2	Precision@3
$f_{cos}^{crit}$	0.256	0.251	0.250
$f_{cos}^{pref}$	<b>0.574</b>	<b>0.546</b>	<b>0.525</b>

Table 2: Precision@1, 2, and 3 for  $f_{cos}^{crit}$  and  $f_{cos}^{pref}$ .

	model	nDCG1	nDCG3	nDCG5	nDCG10
$task_1$	$f_{LM}^{crit}$	0.617	0.674	0.723	0.811
	$f_{LM}^{pref}$	<b>0.731</b>	<b>0.753</b>	<b>0.773</b>	<b>0.858</b>
	$f_{LM}^{concat}$	0.726	0.740	0.773	0.844
$task_2$	$f_{LM}^{crit}$	0.676	0.754	0.805	<b>0.865</b>
	$f_{LM}^{pref}$	0.729	0.761	0.774	0.856
	$f_{LM}^{concat}$	<b>0.805</b>	<b>0.772</b>	<b>0.808</b>	0.863
$task_3$	$f_{LM}^{crit}$	0.498	0.537	0.605	0.660
	$f_{LM}^{pref}$	<b>0.790</b>	<b>0.754</b>	<b>0.758</b>	<b>0.791</b>
	$f_{LM}^{concat}$	0.686	0.663	0.685	0.746

Table 3: nDCG1, 3, 5, and 10 for  $f_{LM}$  on each task.

To train and evaluate the BERT-based method  $f_{LM}$ , we retrieve the top 15 arguments from  $f_{cos}^{crit}$  and the top 15 arguments from  $f_{cos}^{pref}$  for 100 queries. Each argument receives a score from 3 (very relevant) to 1 (irrelevant). Again, a second human annotator relabeled 100 arguments for a Cohen’s Kappa of 0.73, also indicating strong agreement.

We design three ranking tasks:  $task_1$  consists of ranking the 15 arguments originally retrieved with  $f_{cos}^{crit}$ , hence closer to critiques in the embedding space;  $task_2$  consists of ranking the 15 arguments originally retrieved with  $f_{cos}^{pref}$ , hence closer to preferences; and  $task_3$  consists of ranking both sets, i.e., 30 arguments. For each task we train  $f_{LM}^{pref}$ ,  $f_{LM}^{crit}$ , and  $f_{LM}^{concat}$ . We then measure nDCG@1, nDCG@3, nDCG@5, and nDCG@10 in Table 3 averaged after 2-fold cross validation.

### 3.3 Results

We found that using the positive preferences yields substantial improvements in information retrieval. For  $f_{cos}$ , in Table 2,  $f_{cos}^{pref}$  increases Precision@1 by **124%**, Precision@2 by **118%**, and Precision@3 by **110%**. This gap is also present, with marginal variations, when separately analyzing single-, two-, and three-round critiques. For  $f_{LM}$ , in Table 3,  $f_{LM}^{pref}$  outperforms  $f_{LM}^{crit}$  by **19%** on

nDCG@1 even at  $task_1$ , where  $f_{LM}^{crit}$  could have an edge. This gap persists for  $task_2$  ( $f_{LM}^{concat}$  outperforms by **19%**), increases for  $task_3$  ( $f_{LM}^{pref}$  outperforms by **59%**), and tends to narrow towards nDCG@10. Overall, we found strong evidence in support of our hypothesis.

Table 1 shows three examples in which the use of positive preferences was clearly beneficial. These examples represent three critique patterns that cause systematic errors if critique interpretation is turned off: 1. When the user implies a preference for a feature using the polar opposite (e.g., “It looks too casual” implying “I prefer a fancier place”); 2. When the user draws on common sense to express a preference (“It has a freaking band!” implying “I prefer a more quiet place”); and 3. When the user implies a filter within a set of related features (e.g., “I don’t really like seafood” implying preference for alternatives in the meat category).

Analyzing the results of  $f_{cos}^{pref}$  and  $f_{cos}^{crit}$  for the 340 single-round critiques, we found 170 cases where  $f_{cos}^{pref}$  outperformed  $f_{cos}^{crit}$ . Within these, 40 belong to the first pattern (24%), 78 to the second (46%), and 38 to the third (22%).<sup>4</sup> A common trait behind the three patterns is that critiques can be lexically very distinct from their corresponding preference statements, and critique interpretation helps to bridge this gap.

## 4 Conclusion & Future Work

In this paper, we presented an open-ended approach to content-based recommendations for CRS. We developed a novel critique interpretation method that uses GPT3 to infer positive preferences from free-form critiques. We also developed two methods for retrieving recommendations: one that matches embeddings and another that fine-tunes BERT for the task. We ran two ablation studies to test if transforming critiques into positive preferences would yield better recommendations, confirming that it improves performance across both methods.

<sup>4</sup>Fully available at <https://bit.ly/33SCKva>

Finally, we described three critique patterns that cause systematic errors in recommendation search if critique interpretation is turned off.

For future work, we will strive to use critiques to identify and remove unsuitable restaurants; we speculate that the sparsity of customer reviews generally makes it harder to “rule out” than to “rule in.” We will also study other issues such as when to ask clarification questions to resolve ambiguity in the scope of a critique.

## Acknowledgements

We would like to thank reviewers for their helpful feedback. This work was supported in part by gift funding from Adobe Research and by NSF grant IIS-2006851.

## References

Muhammad Abdul-Mageed and Lyle Ungar. 2017. Emonet: Fine-grained emotion detection with gated recurrent neural networks. In *Proceedings of the 55th annual meeting of the association for computational linguistics (volume 1: Long papers)*, pages 718–728, Vancouver, Canada. Association for Computational Linguistics.

Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.

Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gašić. 2018. MultiWOZ - a large-scale multi-domain Wizard-of-Oz dataset for task-oriented dialogue modelling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5016–5026, Brussels, Belgium. Association for Computational Linguistics.

Victor S. Bursztyn, Jennifer Healey, Eunyee Koh, Nedim Lipka, and Larry Birnbaum. 2021. Developing a conversational recommendation system for navigating limited options. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, New York, NY, USA. Association for Computing Machinery.

Bill Byrne, Karthik Krishnamoorthi, Chinnadhurai Sankar, Arvind Neelakantan, Ben Goodrich, Daniel Duckworth, Semih Yavuz, Amit Dubey, Kyu-Young Kim, and Andy Cediñik. 2019. Taskmaster-1: Toward a realistic and diverse dialog dataset. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4516–4525, Hong Kong, China. Association for Computational Linguistics.

Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St. John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, Brian Strope, and Ray Kurzweil. 2018. Universal sentence encoder for English. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 169–174, Brussels, Belgium. Association for Computational Linguistics.

Li Chen and Pearl Pu. 2012. Critiquing-based recommenders: survey and emerging trends. *User Modeling and User-Adapted Interaction*, 22(1-2):125–150.

Konstantina Christakopoulou, Filip Radlinski, and Katja Hofmann. 2016. Towards conversational recommender systems. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 815–824, New York, NY, USA. Association for Computing Machinery.

Joe Davison, Joshua Feldman, and Alexander Rush. 2019. Commonsense knowledge mining from pre-trained models. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1173–1178, Hong Kong, China. Association for Computational Linguistics.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Shuguang Han, Xuanhui Wang, Mike Bendersky, and Marc Najork. 2020. Learning-to-rank with bert in tf-ranking. *arXiv preprint arXiv:2004.08476*.

Matthew Honnibal and Ines Montani. 2017. spacy 2: Natural language understanding with bloom embeddings, convolutional neural networks and incremental parsing. *To appear*, 7(1).

Liwei Jiang, Antoine Bosselut, Chandra Bhagavatula, and Yejin Choi. 2021. “i’m not mad”: Commonsense implications of negation and contradiction. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Online. Association for Computational Linguistics.

Annie Louis, Dan Roth, and Filip Radlinski. 2020. “i’d rather just go to bed”: Understanding indirect answers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7411–7425.

Shengnan Lyu, Arpit Rana, Scott Sanner, and Mohamed Reda Bouadjenek. 2021. A workflow analysis of context-driven conversational recommendation. In *Proceedings of The Web Conference 2021*.

Bodhisattwa Prasad Majumder, Harsh Jhamtani, Taylor Berg-Kirkpatrick, and Julian McAuley. 2020. *Like hiking? you probably enjoy nature: Persona-grounded dialog with commonsense expansions*. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9194–9206, Online. Association for Computational Linguistics.

Rama Kumar Pasumarthi, Sebastian Bruch, Xuanhui Wang, Cheng Li, Michael Bendersky, Marc Najork, Jan Pfeifer, Nadav Golbandi, Rohan Anil, and Stephan Wolf. 2019. Tf-ranking: Scalable tensorflow library for learning-to-rank. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 2970–2978.

Gustavo Penha and Claudia Hauff. 2020. What does bert know about books, movies and music? probing bert for conversational recommendation. In *Proceedings of the 14th ACM Conference on Recommender systems*, New York, NY, USA. Association for Computing Machinery.

Filip Radlinski, Krisztian Balog, Bill Byrne, and Karthik Krishnamoorthi. 2019. Coached conversational preference elicitation: A case study in understanding movie preferences.

Yueming Sun and Yi Zhang. 2018. Conversational recommender system. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 235–244, New York, NY, USA. Association for Computing Machinery.

Yongfeng Zhang, Xu Chen, Qingyao Ai, Liu Yang, and W Bruce Croft. 2018. Towards conversational search and recommendation: System ask, user respond. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 177–186, New York, NY, USA. Association for Computing Machinery.