

Towards a Taxonomy of User Feedback Intents for Conversational Recommendations



Wanling Cai and Li Chen

Department of Computer Science, Hong Kong Baptist University, Hong Kong, China
{cswlcai, lichen}@comp.hkbu.edu.hk

Introduction

- **Dialogue-based Conversational Recommender System (DCRS)** allows users to interact with the system in natural language for seeking for recommendations.
- **Current Limitation:** Most of existing systems have provided one-shot recommendation, but little work has investigated users' language-based feedback when they did not get the desired item within a single turn.
- **Research Question:** How do users express their feedback on the unsatisfied recommendation in natural language (e.g., what intents they may have and what kinds of preferences they may express)?

Recommender:	Hi how are you today? I heard you might be interested in a movie. Any particular genre?
Seeker:	Hi, I'm good, just looking for a nice horror movie. Nothing too gory, I liked Beetlejuice.
Recommender:	hmm. I don't know too many horror movies. I did watch The Birds .
Seeker:	Yeah I've seen the birds it was okay but I felt like it was too old for my tastes.
Recommender:	border line with suspense might be something like Hannibal or The Silence of the Lambs .
Seeker:	I didn't like any of those movies, too much talking.
Recommender:	okay. Well, how about Saw ?
Seeker:	Something more like Final Destination.
Recommender:	Do you like any other genres?
Seeker:	The Saw was okay, I felt like it was too violent. I really love like fantasy horror, maybe Ghost.
Recommender:	I've heard that is a good one. Have you seen Signs ?
Seeker:	I heard about that but didn't watch it.
Recommender:	Mel Gibson in it. I've heard it is excellent.
Seeker:	okay, great I will check it out. thank you.

Seekers' feedback on recommendations

An example of human-human dialogue
(source: ReDail [3])


Taxonomy for User Feedback Intents

We established a taxonomy to classify users' feedback intents through theme identification techniques and constant comparison method.

User Feedback Intent (Code)	Description	Example	Percentage
Reject (REJ)	Seeker dislikes the recommended item.	"I hated that movie. I did not even crack a smile once."	19.2%
Seen (SEE)	Seeker has seen the recommended item before.	"I have seen that one and enjoyed it."	16.3%
Critique-Feature (CRI-F)	Seeker makes critique on specific features of the current recommendation.	"That's a bit too scary for me."	11.8%
Provide Details (PRO)	Seeker provides detailed preferences for the item s/he is looking for.	"I usually enjoy movies with Seth Rogen and Jonah Hill."	11.7%
Inquire (INQ)	Seeker wants to know more about the recommended item.	"I haven't seen that one yet. What's it about?"	10.9%
Critique-Add (CRI-A)	Seeker adds further constraints on top of the current recommendation.	"I would like something more recent."	8.5%
Start Over (STO)	Seeker starts a new query.	"Anything that I can watch with my kids under 10."	5.2%
Neutral Response (NRE)	Seeker does not indicate her/his preferences for the current recommendation.	"I have actually never seen that one."	5.1%
Critique-Compare (CRI-C)	Seeker requests something similar to the current recommendation.	"Den of Thieves (2018) sounds amazing. Any others like that?"	2.9%
Answer (ANS)	Seeker answers the question issued by the recommender.	"Maybe something with more action." (Q: "What kind of fun movie you look for?")	2.8%
Ask (ASK)	Seeker asks the recommender's personal opinions.	"I really like Reese Witherspoon. How about you?"	1.6%
Restate with Further Constraints (RES-CO)	Seeker restates her/his query with further constraints.	"Do you have something that is a thriller but not too scary?"	1.6%
Restate (RES)	Seeker completely restates her/his query.	"Maybe I am not being clear. I want something that is in the theater now."	1.5%
Restate with Clarification (RES-CL)	Seeker restates her/his query with clarification.	"I'm fine with any sort of horrors, jump scares, clowns, etc."	0.4%
Others (OTH)	The utterance cannot be categorized into any other categories.	"Sorry about the weird typing."	0.4%


Feedback-Oriented Dialogue Data Analysis

Based on the established taxonomy, we annotated a set of over 200 dialogues [3], and analyzed the data from the following aspects:



Seeker feedback intent co-occurrence.

Reject often co-occurs with *Critique-Add*, *Provide Details*, and *Start Over*.



Seeker preference expression.

Seekers more frequently express their preferences at the **attribute** level (with more **objective** criteria), but express **subjective** opinions on **entity**.

Action (Code)	Description	Percentage
Recommend (REC)	Recommender provides one or more recommendations.	43.8%
Explain (EXP)	Recommender explains why the item is recommended.	30.0%
Respond (RES)	Recommender responds to any other queries by the seeker.	12.4%
Answer (ANS)	Recommender answers the question from the seeker.	10.2%
Request (REQ)	Recommender requests for the seeker's preferences.	3.1%

Recommender reactions to seekers' feedback, and action distribution in our dataset.

Recommenders tend to *recommend* some other items when receiving seekers' negative feedback.

Future Work

- To reveal frequent conversation patterns that may occur between users and recommenders through temporal analysis.
- To develop a dedicated model to predict users' intents given their utterances.
- To eventually build a more effective DCRS that is able to track users' current states, refine their preference model, and then select appropriate actions to respond to users.

References

[1] Jie Kang, Kyle Condiff, Shuo Chang, Joseph A. Konstan, Loren Terveen, and F. Maxwell Harper. 2017. Understanding How People Use Natural Language to Ask for Recommendations. In *Proceedings of the Eleventh ACM Conference on Recommender Systems (RecSys '17)*. 229–237. <https://doi.org/10.1145/3109859.3109873>

[2] Li Chen and Pearl Pu. 2012. Critiquing-based Recommenders: Survey and Emerging Trends. *User Modeling and User-Adapted Interaction* 22, 1-2 (April 2012), 125–150. <https://doi.org/10.1007/s11257-011-9108-6>

[3] Raymond Li, Samira Ebrahimi Kahou, Hannes Schulz, Vincent Michalski, Laurent Charlin, and Chris Pal. 2018. Towards Deep Conversational Recommendations. In *Advances in Neural Information Processing Systems* 32. Curran Associates, Inc., 9748–9758.