
Multi-domain Conversation Quality Evaluation via User Satisfaction Estimation

Praveen Kumar Bodigutla
pbodigut@amazon.com

Lazaros Polymenakos
polym1@amazon.com

Spyros Matsoukas
matsouka@amazon.com

Abstract

An automated metric to evaluate dialogue quality is vital for optimizing data driven dialogue management. The common approach of relying on explicit user feedback during a conversation is intrusive and sparse. Current models to estimate user satisfaction use limited feature sets and employ annotation schemes with limited generalizability to conversations spanning multiple domains. To address these gaps, we created a new *Response Quality* annotation scheme, introduced five new domain-independent feature sets and experimented with six machine learning models to estimate *User Satisfaction* at both turn and dialogue level.

Response Quality ratings achieved significantly high correlation (0.76) with explicit turn-level user ratings. Using the new feature sets we introduced, Gradient Boosting Regression model achieved best (rating [1-5]) prediction performance on 26 seen (linear correlation ~ 0.79) and one new multi-turn domain (linear correlation 0.67). We observed a 16% relative improvement (68% \rightarrow 79%) in binary (“satisfactory/dissatisfactory”) class prediction accuracy of a domain-independent dialogue-level satisfaction estimation model after including predicted turn-level satisfaction ratings as features.

1 Introduction

Automatic turn and dialogue level quality evaluation of end user interactions with Spoken Dialogue Systems (SDS) is vital for identifying problematic conversations and for optimizing dialogue policy using a data driven approach, such as reinforcement learning. One of the main obstacles to designing data-driven policies is the lack of an objective function to measure the success of a particular interaction. Existing methods along with their limitations to measure dialogue success can be categorized into five groups: 1) Using sparse sentiment for end-to-end dialogue system training; 2) Using task success as dialogue evaluation criteria, which does not capture frustration caused in intermediate turns and assumes end user goal is known in advance; 3) Explicitly soliciting feedback from the user, which is intrusive and may cause dissatisfaction; 4) Using a popular approach, such as PARADISE [Walker *et al.*, 2000], to predict dialogue-level satisfaction ratings provided by surveyed users, which has limited generalizability to diverse (novice and experienced) user population [Deriu *et al.*, 2019]; and 5) Estimating per turn dialogue quality using trained Interaction Quality (IQ) [Schmitt *et al.*, 2012] estimation models. Per turn discrete 1-5 scale IQ labels are provided by the annotator while keeping track of dialogue quality till that turn.

Prior attempts to model user satisfaction at turn level as a continuous process evolving over time, employed either users [Engelbrecht *et al.*, 2009] or expert annotators [Higashinaka *et al.*, 2010] to rate each turn in a dialogue. Both approaches suffered from manual feature extraction process, which IQ addressed. Features used in the IQ estimation models were derived from the current turn, the dialogue history, and output from three Spoken Language Understanding (SLU) components, namely: Automatic Speech Recognition (ASR), Natural Language Understanding (NLU), and the

dialogue manager. IQ annotations provided by expert annotators are reliable in comparison to explicit ratings provided by the end users. Various models have been explored to predict IQ, including Hidden Markov Models [Ultes *et al.*, 2014], Support Vector Machines (SVM) [Schmitt *et al.*, 2011], Support Vector Ordinal Regression (SVOR) [Asri *et al.*, 2014], Recurrent Neural Networks (RNN) [Pragst *et al.*, 2017], and Long Short-Term Memory Networks (LSTM) [Rach *et al.*, 2017].

Conversations users have with modern day SDS span multiple domains in a single dialogue session (Table 1). IQ annotation scheme which is developed using within domain conversations has limited generalizability to dialogues that span across multiple domains.

Utterance	Comment
User: Play latest hits. <i>{domain:Music}</i>	<i>User makes a popular request</i>
System: Shuffling from your playlist.	<i>System addresses the request successfully</i>
User: Stop. Cancel my evening appointment. <i>{domain:Calendar}</i>	<i>User barges-in</i>
System: Sorry I don't know that one.	<i>System identifies un-actionable request</i>
User: Cancel my 7pm event if it is raining today. <i>{domain:Calendar, Weather}</i>	<i>User paraphrases</i>
System: Today in Seattle, there is 60% chance of rain.	
Should I cancel your 7pm event - Park Visit?	<i>System Informs and asks a question.</i>
User: Yes please. <i>{domain:Global}</i>	<i>User confirms</i>
System: Alright, 7pm event - Park visit Canceled.	<i>System executes user request</i>

Table 1: Conversation spanning multiple domains.¹

The IQ approach to rate turns in a dialogue incrementally, increases cognitive load on the annotators and lowers inter-annotator reliability. We also hypothesize that the features using in IQ estimation models are limited, and including additional contextual signals improves the performance of dialogue quality estimation models for both single-turn and multi-turn conversations².

To address the aforementioned gaps in existing dialogue quality estimation methods, we propose an end-to-end User Satisfaction Estimation (USE) metric which predicts user satisfaction rating for each dialogue turn on a continuous 1-5 scale. To obtain consistent, simple and generalizable annotation scheme that easily scales to multi-domain conversations, we introduce turn-level Response Quality (RQ) annotation scheme. To improve the performance of the satisfaction estimation models, we design five new feature sets, which are: 1) user request paraphrasing indicators, 2) cohesion between user request and system response, 3) diversity of topics discussed in a session, 4) un-actionable user request indicators, and 5) aggregate popularity of domains and topics across the entire population of the users.

Using RQ annotation scheme, annotators rated dialogue turns from 26 single-turn and multi-turn sampled Alexa domains³ (e.g., *Music, Calendar, Weather, Movie booking*). We trained USE machine learning models using the annotated RQ ratings. To explain model predicted rating using features' values, we experimented with four interpretable models that rank features by their importance. We benchmarked performance of these models against two state-of-the-art dialogue quality prediction models. Using ablation studies, we showed improvement in the best performing USE model's performance using new contextual features we introduced.

In order to test the generalization performance of the model, we reserved one "new" multi-turn Alexa domain in the test set. Furthermore, we trained a dialogue-level satisfaction estimation model with

¹ Due to confidentiality consideration, this example is not a real user conversation but authored to mimic the true dialogue.

² In single-turn conversations the entire context is expected to be present in the same turn, while in a multi-turn case the context is carried from previous turns to address user's request in the current turn.

³ A multi-turn domain supports multi-turn conversations whereas a single-turn domain treats each user query as an independent request where context is not carried forward from previous turns.

explicit session-level ratings provided by both novice and experienced users on multi-turn dialogues from 14 domains. We evaluated the impact of including turn-level satisfaction ratings as features on the performance of this domain-independent dialogue-level satisfaction estimation model.

The outline of the paper is as follows: Section 2 introduces the Response Quality annotation scheme and discusses its effectiveness in terms of predicting turn-level user satisfaction ratings. Section 3 summarizes the Response Quality annotated data, discusses dialogue-level satisfaction estimation using ratings provided by novice & experienced users and presents our experimentation setup. Section 4 provides an empirical study of turn and dialogue-level satisfaction estimation models' performance and discusses results from feature ablation study. Section 5 concludes.

2 Response Quality Annotation

We designed the Response Quality (RQ) annotation scheme to generate training data for our turn-level User Satisfaction Estimation (USE) model.

2.1 RQ annotation Scheme and Comparison with IQ

In RQ, similar to IQ annotations, annotators listened to raw audio and provided per turn's system RQ rating on a 5-point scale. The scale we asked annotators to follow was: 1=Terrible (system fails to understand and fulfill user's request), 2=Bad (understands the request but fails to satisfy it in any way), 3=OK (understands users request and either partially satisfies the request or provides information on how the request can be fulfilled), 4=Good (understands and satisfies the user request, but provides more information than what the user requested or takes extra turns before meeting the request), and 5=Excellent (understands and satisfies user request completely and efficiently). Using a 5 point scale over a binary scale provides an option to the annotators to factor in their subjective interpretation of the extent of success or failure of system's response to satisfy a user's request.

Annotators rated conversations that spanned multiple domains. They were instructed to use the follow-up feedback from the user (e.g., user expresses frustration or rephrases an initial request) in making judgements. Unlike IQ annotation scheme, we removed the constraint on the annotators to keep track of the quality of dialogue so far while determining RQ ratings for a given turn. This relaxation in constraint, coupled with making full conversation context available to the annotators, reduced the cognitive load on them. This simplified annotation scheme not only helped in scaling RQ to multiple domains but also enabled precise identification of defective turns which is not straightforward in the case of IQ where an individual turn's IQ rating depends on the prior turns' ratings.

2.2 Inter Annotator Agreement (IAA) and Correlation with user satisfaction rating

We conducted a user study to verify the accuracy of RQ and IQ [Schmitt *et al.*, 2012] annotation process. In the study, eight users were asked to achieve 30 pre-determined goals. Dialogue sessions to achieve the given goals spanned over six single-turn and two multi-turn Alexa domains. For 15 out of the 30 goals, we asked the users to provide satisfaction rating on a discrete (1-5) scale based on turn's system response. For the remaining 15 goals, we asked the users to rate each turn incrementally based on their perception of interaction so far. Then we sent the same utterances to six annotators (3 annotators per annotation type) for obtaining RQ (950 turns) and IQ (700 turns) annotations. The IQ annotations are less in number in comparison to RQ because we excluded turns where the users were unsure about the per-turn incremental satisfaction ratings they provided. We did not impose any restriction on the number of turns required to achieve a given goal, which led to an unequal number of turns per dialogue session across users who were assigned the same goal.

Since IQ annotation guidelines suggest rating unrelated queries independently, the annotators providing IQ annotations, were also instructed to identify independent interactions within a given dialogue session. Specifically, annotators marked the beginning of a new interaction if they felt that the user's goal in the current turn is unrelated to the one he/she tried to achieve in the previous turn.

We found that the [1–5] scale RQ ratings provided by 3 annotators were highly correlated (Spearman's rho 0.94)⁴ with each other, suggesting high IAA. The mean RQ ratings were significantly (at

⁴ We also asked 3 novice annotators to provide RQ ratings for 400 single and multi-turn conversations' turns, the IAA was still high with mean Spearman's rho of 0.75.

95% confidence interval) correlated (0.76) with surveyed user satisfaction ratings with system’s response. In the case of IQ ratings, mean IAA and correlation with user ratings dropped to 0.20 and 0.32 respectively, suggesting limited generalizability of IQ annotation scheme to multi-domain conversations.

We used cohen’s kappa to measure IAA between IQ annotators on the binary annotations they provided to indicate the beginning of an interaction within a dialogue session. Low IAA (cohen’s kappa 0.16) achieved on these binary annotations, indicate ambiguity faced by the annotators in-terms of identifying end user’s goal from a sequence of turns within a dialogue.

3 Data and Experimental Setup

This section describes our turn-level and dialogue-level dataset, details the list of features derived from various signals, and explains our experimentation setup.

3.1 Turn-level Response Quality Data

To demonstrate that our RQ annotation scheme and predictive models are domain-independent and effective for both single-turn and multi-turn dialogues, we used 30,500 dialogue turns randomly sampled from 26 single-turn (90%) and multi-turn (10%) domains that are representative of end user interactions with Alexa. The imbalance towards single-turn dialogues is due to annotation priority. We also tested model’s generalization performance on 200 dialogue turns sampled from a “new” multi-turn goal oriented application. Figure 1 shows the diversity in rating distribution between the single-turn and multi-turn dialogues.




Figure 1: Distribution of RQ annotation for single-turn, multi-turn domains and new multi-turn application. Exact percentage on y-axis is masked for confidentiality

3.2 Features

To estimate the turn level user satisfaction score, we used features derived from turn, dialogue context, and Spoken Language Understanding (SLU) components’ output similar to turn level IQ prediction models [Ultes *et al.*, 2017]. As shown in Figure 2, we define a dialogue turn at time n as $t_n = (t_n^u, t_n^s)$, where t_n^u and t_n^s represent the user request and system response on turn n . A dialogue session of N turns is defined as $(t_1:t_N)$. To improve the performance of the USE models across single and multi-turn conversations spanning multiple domains, we introduced the following 5 sets of domain-independent features:

1. **User request paraphrasing** – Calculated by measuring syntactic and semantic (NLU predicted intent) similarity between consecutive turns’ user utterances.
2. **Cohesion between response and request** – Cohesiveness of system response with a user request is computed by calculating jaccard similarity score between user request and system response. System response “*Here is a sci-fi movie*” will get higher cohesion score over “*Here is a comedy movie*”, if the user request was “*recommend a sci-fi movie*”.
3. **Aggregate topic popularity** – Usage statistics such as aggregate domain and intent usage count and ratio of usage count to number of customers, provides us a prior on the popularity of a topic across all users of Alexa.

4. **Un-actionable user request** – Identifies if the user request could not be fulfilled, by searching for phrases indicating an apology and negation in system response (e.g., “*sorry I don’t know how to do that*”).
5. **Diversity of topics in a session** – This dialogue level feature is calculated using the percentage of unique intents till the current dialogue turn.

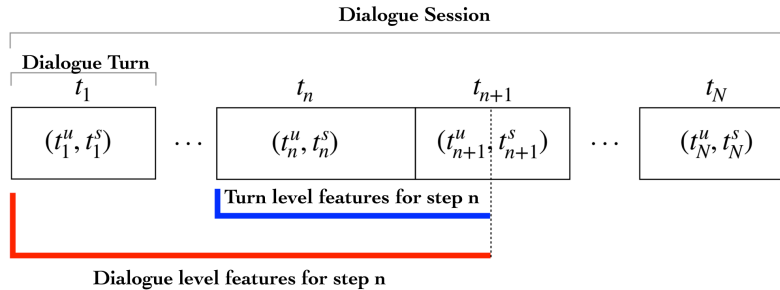


Figure 2: dialogue and turn definitions. The blue and red lines indicate the history used for generating turn-based and dialogue-based features for the user satisfaction estimation on turn t_n .

3.3 Dialogue-level satisfaction estimation using predicted turn-level ratings

To show that predicted turn-level satisfaction ratings are effective in estimating overall dialogue level satisfaction we trained dialogue-level quality estimation models to predict explicit dialogue-level ratings provided by users. Dialogue-level labels were obtained from ratings provided by 10 users (5 novice and 5 experienced) who were asked to achieve 40 multi-turn goals spanning 14 {commercial conversation agent’s} domains (example goals in Appendix Table 5).

Since earlier attempts to estimate explicit dialogue-level satisfaction ratings did not generalize to different user population (see section 1), we collected data from users belonging to both “novice” and “experienced” groups. A novice user has minimal experience conversing with the SDS and he/she has never used the functionality provided by the 14 domains prior to the study. An experienced user is a seasoned user of Alexa and its domains.

Users provided their satisfaction rating with the dialogue on a discrete $[1 - 5]$ scale at the end of each session, irrespective of the outcome. The rating scale we asked the annotators to follow was 1=Very dissatisfied, 2=Dissatisfied, 3=Moderately Satisfied (or Slightly dissatisfied), 4=Satisfied and 5=Extremely Satisfied. We collected 1, 042 dialogue-level satisfaction ratings in total. 45% of these ratings came from novice users and the rest were provided by the experienced users. Figure 3 shows the distribution of dialogue level ratings from both sets of users.

Features to train the dialogue level satisfaction model included turn-level features computed on the last turn (t_N), aggregate statistics (e.g., average ASR confidence score, number of barge-ins) computed over all turns ($t_1:t_N$) in a dialogue session and average estimated turn-level satisfaction rating calculated across all turns in the same session.

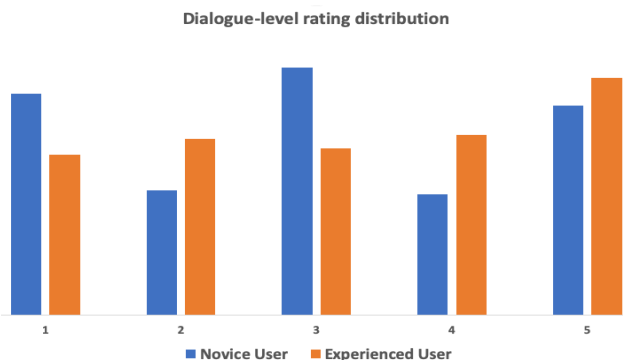


Figure 3: Distribution of dialogue level satisfaction ratings. Exact percentage on y-axis is masked for confidentiality

3.4 Experimental Setup

This section describes the experimental setup we used for training and evaluating turn and dialogue level satisfaction estimation models.

3.4.1 Turn-level satisfaction estimation

We obtained a turn’s RQ rating by averaging the discrete 1-5 labels provided by 3 annotators. Hence, we considered regression models for experimentation which predicted satisfaction rating on a continuous 1-5 scale. To achieve interpretability, in our experiments we selected four models - LASSO [Tibshirani, 1996], Decision Tree Regression [Agrawal *et al.*, 1993], Random Forest Regression [Breiman, 2001] and Gradient Boosting Regression [Friedman, 2001] that rank features by their importance. For benchmarking we used Multi-layer Perceptron (MLP) [Gardner and Dorling, 1998] and Support Vector Regression (SVR) ⁵. Data was randomly split into training (60%), validation (20%) and test (20%) sets, however turns from the “new” multi-turn application were included only in the test set.

3.4.2 Dialogue-level satisfaction estimation

We trained the same four interpretable regression models (LASSO, Decision Tree, Random Forest and Gradient Boosting) to predict explicit dialogue level ratings on a continuous [1 – 5] scale. Since ratings obtained at dialogue level from users are sparse (3.5% of total number of turn-level RQ annotations), we performed 9-fold cross validation on randomly sampled 90% of data and we evaluated the performance of the trained model on remaining 10% test data.

3.4.3 Evaluation Criteria

We used Pearson’s linear correlation coefficient (r) for evaluating each model’s 1-5 prediction performance. For the use case to identify problematic turns from an end user’s perspective, it is sufficient to identify *satisfactory* (rating ≥ 3) and *dissatisfactory* (rating < 3) interactions. For evaluating turn-level satisfaction models, we used F-score for the dissatisfactory class as the binary classification metric. Identifying dissatisfactory turns is of more importance and is in general a difficult task as majority of turns belong to the satisfactory class (Figure 1). We used *satisfactory* and *dissatisfactory* binary class prediction accuracy as an additional metric to evaluate the performance of dialogue-level satisfaction estimation models, since unlike turn-level data, the dialogue ratings are not concentrated around one class (Figure 3).

4 Results and Analysis

In this section we present a performance analysis of both turn-level and dialogue-level satisfaction estimation models.

4.1 Turn-level satisfaction estimation results

Model/Metric	Single-turn		Multi-turn		Multi-turn new application	
	<i>Correlation</i>	<i>F – dissatisfactory</i>	<i>Correlation</i>	<i>F – dissatisfactory</i>	<i>Correlation</i>	<i>F – dissatisfactory</i>
Lasso	0.70 ± 0.01	0.69 ± 0.02	0.68 ± 0.05	0.74 ± 0.06	0.62 ± 0.07	0.74 ± 0.06
Decision Tree	0.73 ± 0.01	0.70 ± 0.03	0.67 ± 0.05	0.71 ± 0.06	0.58 ± 0.07	0.73 ± 0.05
Random Forest	0.77 ± 0.01	0.74 ± 0.02	0.74 ± 0.05	0.75 ± 0.05	0.61 ± 0.07	0.76 ± 0.05
G.Boost	0.80 ± 0.01	0.77 ± 0.02	0.79 ± 0.04	0.78 ± 0.05	0.67 ± 0.06	0.79 ± 0.05
SVR	0.73 ± 0.01	0.72 ± 0.02	0.75 ± 0.04	0.71 ± 0.06	0.64 ± 0.07	0.76 ± 0.05
MLP	0.75 ± 0.01	0.73 ± 0.02	0.74 ± 0.04	0.73 ± 0.05	0.67 ± 0.06	0.79 ± 0.05

Table 2: Six machine learning regression models performance on turns from single-turn, multi-turn conversations and new multi turn test application. Each cell shows the mean and 95% bootstrap confidence interval with the highest mean in bold.

⁵ Recurrent Neural Networks and Long Short-term Memory Networks [Rach *et al.*, 2017] [Drucker *et al.*, 1997] showed non-significant improvement (≤ 0.02 difference in Spearman’s rho) over state-of-the-practice MLP and SVM models in predicting IQ.

Amongst the six models we experimented with, Gradient Boosting Regression achieves superior performance on single-turn domains and multi-turn domains ($r = \sim 0.79$ and F-dissatisfaction = 0.77). On the new multi-turn application, both Gradient Boosting Regression and MLP models achieve better performance ($r = 0.67$ and F-dissatisfaction = 0.79) in comparison to the other four models we experimented with (results in Table 2).

Features Removed\Metric	Single-turn		Multi-turn		Multi-turn new application	
	<i>Correlation</i>	<i>F - dissatisfaction</i>	<i>Correlation</i>	<i>F - dissatisfaction</i>	<i>Correlation</i>	<i>F - dissatisfaction</i>
None	0.80 ± 0.01	0.77 ± 0.02	0.79 ± 0.04	0.78 ± 0.05	0.67 ± 0.06	0.79 ± 0.05
Aggregate topic Popularity	0.74 ± 0.02	0.72 ± 0.02	0.75 ± 0.04	0.72 ± 0.06	0.62 ± 0.07	0.74 ± 0.06
Un-actionable Request identifier	0.76 ± 0.01	0.74 ± 0.02	0.78 ± 0.05	0.77 ± 0.06	0.50 ± 0.08	0.70 ± 0.05
Cohesion between request and response	0.78 ± 0.01	0.77 ± 0.02	0.76 ± 0.04	0.77 ± 0.05	0.63 ± 0.07	0.77 ± 0.06
Topic diversity	0.79 ± 0.01	0.77 ± 0.01	0.77 ± 0.04	0.75 ± 0.05	0.67 ± 0.05	0.78 ± 0.05
User request paraphrasing	0.78 ± 0.01	0.76 ± 0.02	0.78 ± 0.04	0.76 ± 0.05	0.65 ± 0.05	0.76 ± 0.05

Table 3: Ablation study results on turn-level satisfaction estimation performance. Each cell shows the mean and 95% bootstrap confidence interval with the highest mean in bold. Performance is measured using Linear correlation between predicted and annotated RQ ratings and F-score on dissatisfaction class (F-dissatisfaction). Bold highlights indicate statistically significant difference in performance.

Based on ablation study using Gradient boosting regression model, we found that the new features improved every single metric on the test set. On single-turn dialogues, features corresponding to “aggregate topic popularity” caused largest statistically significant $\sim 7\%$ relative improvement in linear correlation (0.741 \rightarrow 0.796) and F-dissatisfaction (0.72 \rightarrow 0.77) scores. On the new multi-turn application, $\sim 35\%$ relative improvement in linear correlation (0.496 \rightarrow 0.67) shows significant impact of “Un-actionable request” feature on generalization performance (Table3). The five new feature sets we introduced occur in the top 10 sets of important features (Appendix Table 7) returned by model based on their computed importance score.

4.2 Dialogue-level satisfaction estimation results

Model\Metric	Without Average-estimated turn-level rating feature			With Average-estimated turn-level rating feature		
	<i>Accuracy</i>	<i>F - dissatisfaction</i>	<i>Correlation</i>	<i>Accuracy</i>	<i>F - dissatisfaction</i>	<i>Correlation</i>
Lasso	0.73 ± 0.05	0.64 ± 0.09	0.45 ± 0.13	0.78 ± 0.07	0.71 ± 0.09	0.58 ± 0.11
Decision Tree	0.64 ± 0.11	0.56 ± 0.12	0.31 ± 0.11	0.73 ± 0.06	0.63 ± 0.08	0.52 ± 0.10
Random Forest	0.67 ± 0.06	0.61 ± 0.12	0.47 ± 0.11	0.71 ± 0.07	0.64 ± 0.11	0.58 ± 0.10
G.Boost	0.68 ± 0.05	0.61 ± 0.12	0.58 ± 0.11	0.79 ± 0.07	0.73 ± 0.08	0.60 ± 0.09

Table 4: Dialogue level satisfaction estimation performance obtained by four machine learning models with and without Average estimated turn-level satisfaction rating as feature. Each cell shows the mean and 95% bootstrap confidence interval with the highest mean in bold. Wider confidence intervals are due to small sample size.

Across all four models we experimented with, we observed an improvement in dialogue level satisfaction estimation results by including average estimated turn-level satisfaction rating as a feature (results in Table 4). Turn-level satisfaction ratings were estimated using the Gradient Boosting Regression model which is described in the previous section (Section 4.1). Best performing dialogue-level Gradient Boosting Regression model achieved 16% relative improvement (68% \rightarrow 79%) in binary “satisfactory/dissatisfactory” class prediction accuracy and $\sim 3.5\%$ relative improvement (0.58 \rightarrow 0.60) in correlation between predicted and true dialogue-level labels. Based on feature importance scores returned by the Gradient Boosting regression model, average predicted turn-level satisfaction rating score is the most important feature for predicting dialogue-level satisfaction rating, followed by average ASR scores and average aggregate domain popularity (top 10 features in Appendix Table 9).

5 Conclusion

In this paper, we described a user-centric and domain-independent approach for evaluating user satisfaction in multi-domain conversations users have with an AI assistant. We introduced Response Quality (RQ) annotation scheme which is highly correlated ($r = 0.76$) with explicit turn level user satisfaction ratings. By designing five additional new features, we achieved a high linear correlation of ~ 0.79 between annotated RQ and predicted User Satisfaction ratings with Gradient Boosting Regression as the User Satisfaction Estimation (USE) model, for both single-turn and multi-turn dialogues. Gradient Boosting Regression and Multi Layer Perceptron (MLP) models generalized to new domain better ($r = 0.67$) than other models. By including turn-level satisfaction

prediction as features, we observed a relative 16% improvement (68% \rightarrow 79%) improvement in binary “dissatisfactory/satisfactory” class prediction accuracy of a domain-independent dialogue-level satisfaction estimation model. The Gradient Boost Regression based dialogue-level model estimated explicit dialogue-level ratings provided by novice and experienced users on conversations spanning 14 domains and achieved significant (at 95% confidence interval) 0.60 correlation between predicted and true user ratings.

With statistically significant $\sim 7\%$ and $\sim 35\%$ relative improvement in linear correlation on existing domains and new multi-turn application respectively, our ablation study supported our hypothesis that the new features improve model prediction performance. On multi-domain conversations, we plan to explore the use of Deep Neural Net models to reduce handcrafting of features [Rach *et al.*, 2017] and to jointly estimate end user satisfaction at turn and dialogue level, though these models reduce interpretability. To learn dialogue policies using reinforcement learning, we plan to experiment with proposed RQ based User Satisfaction metric as an alternative for reward modeling.

Acknowledgments

We thank Alborz Geramifard and NeurIPS Conversation AI workshop reviewers for their insightful feedback and guidance. We also thank Longshaokan Wang, Swanand Joshi and Joshua Levy for helping with data procurement and turn level satisfaction estimation setup. Finally, we thank Alison Bauter Engel, Kate Ridgeway and Alexa Data Services-RAMP team for their significant help with user studies and data annotations.

References

- Rakesh Agrawal, Tomasz Imieliński, and Arun Swami. Mining association rules between sets of items in large databases. *SIGMOD Rec.*, 22(2):207–216, June 1993.
- L. E. Asri, H. Khouzaimi, R. Laroche, and O. Pietquin. Ordinal regression for interaction quality prediction. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3221–3225, May 2014.
- Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- Jan Deriu, Álvaro Rodrigo, Arantxa Otegi, Guillermo Echevoyen, Sophie Rosset, Eneko Agirre, and Mark Cieliebak. Survey on evaluation methods for dialogue systems. *CoRR*, abs/1905.04071, 2019.
- Harris Drucker, Christopher J. C. Burges, Linda Kaufman, Alex J. Smola, and Vladimir Vapnik. Support vector regression machines. In M. C. Mozer, M. I. Jordan, and T. Petsche, editors, *Advances in Neural Information Processing Systems 9*, pages 155–161. MIT Press, 1997.
- Klaus-Peter Engelbrecht, Florian Gödde, Felix Hartard, Hamed Ketabdar, and Sebastian Möller. Modeling user satisfaction with hidden Markov models. In *Proceedings of the SIGDIAL 2009 Conference*, pages 170–177, London, UK, September 2009. Association for Computational Linguistics.
- Jerome H Friedman. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pages 1189–1232, 2001.
- MW Gardner and SR Dorling. Artificial neural networks (the multilayer perceptron)—A review of applications in the atmospheric sciences. *Atmospheric Environment*, 32(14-15):2627–2636, 1998.
- Ryuichiro Higashinaka, Yasuhiro Minami, Kohji Dohsaka, and Toyomi Meguro. Issues in predicting user satisfaction transitions in dialogues: Individual differences, evaluation criteria, and prediction models. In *IWSDS*, 2010.
- Louisa Pragst, Stefan Ultes, and Wolfgang Minker. *Recurrent Neural Network Interaction Quality Estimation*, pages 381–393. Springer Singapore, Singapore, 2017.
- Niklas Rach, Wolfgang Minker, and Stefan Ultes. Interaction quality estimation using long short-term memories. In *SIGDIAL Conference*, 2017.
- Alexander Schmitt, Benjamin Schatz, and Wolfgang Minker. Modeling and predicting quality in spoken human-computer interaction. In *SIGDIAL Conference*, 2011.
- Alexander Schmitt, Stefan Ultes, and Wolfgang Minker. A parameterized and annotated spoken dialog corpus of the cmu let’s go bus information system. In Nicoletta Calzolari (Conference Chair), Khalid Choukri, Thierry Declerck, Mehmet U?ur Do?an, Bente Maegaard, Joseph Mariani, Asuncion Moreno, Jan Odijk, and Stelios Piperidis, editors, *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC’12)*, Istanbul, Turkey, may 2012. European Language Resources Association (ELRA).
- Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1):267–288, 1996.
- Stefan Ultes, Robert ElChab, and Wolfgang Minker. Application and evaluation of a conditioned hidden markov model for estimating interaction quality of spoken dialogue systems. In Joseph Mariani, Sophie Rosset, Martine Garnier-Rizet, and Laurence Devillers, editors, *Natural Interaction with Robots, Knowbots and Smartphones*, pages 303–312, New York, NY, 2014. Springer New York.
- Stefan Ultes, Pawel Budzianowski, Iñigo Casanueva, Nikola Mrksic, Lina Maria Rojas-Barahona, Pei hao Su, Tsung-Hsien Wen, Milica Gasic, and Steve J. Young. Domain-independent user satisfaction reward estimation for dialogue policy learning. In *INTERSPEECH*, 2017.
- Marilyn Walker, Candace Kamm, and Diane Litman. Towards developing general models of usability with paradise. *Nat. Lang. Eng.*, 6(3-4):363–377, September 2000.

A Appendices

Domain/Application	Goal
Multi-turn conversational bot for movies	Get ratings of two other movies directed by the director of a movie you liked watching recently.
Multi-turn conversational bot for movies	Ask for movie recommendations and navigate through them until you got a recommendation you liked.
Multi-turn conversational bot for movies → Music	Find details about the cast of your favorite movie and play the movie’s soundtrack.
Ticket-booking third party application	Get a list of theaters near you and find out what is playing in at least 2 theaters returned by the application for the same query.
Music→Knowledge	Find song by artist and find out in which movie it was used
Music→Music	Find a popular album by an artist and play it
Knowledge→Local Search	Ask for a fact about a historical monument and find its address
Recipe	Ask for a general type of recipe (e.g. cookies, pasta dishes, etc.), choose a recipe from the list using your voice, and then add the ingredients to your shopping list.
Ticket-booking third party application	Try purchasing ticket for a movie playing in theaters near you.
Ticket-booking third party application	Find what’s playing in a theater outside your current city (or outside the default city from your device)
Notification	Add an appointment to the calendar and then change the time. (requires the user to first cancel and then create a new appointment)

Table 5: Example Multi-turn goals for user study

Model	Hyper parameter - Optimal value
Lasso	alpha: 0.001
Decision Trees	max-depth: 33, min-samples-leaf: 31, min-samples-split: 23
Random Forest	max-depth: 49, min-samples-leaf: 11, min-samples-split: 27
Gradient Boosting Decision Trees	max-depth: 23, min-samples-leaf: 17, min-samples-split: 59
SVR	c: 2, gamma: 0.024
MLP	n-layers: 3, batch-size: 128, hidden size : 100, solver: ‘sgd’ activation: ‘relu’

Table 6: Optimal Hyper parameter values used for training turn-level User Satisfaction Estimation (USE) models.

Feature set description	Turn(s) the feature set is computed on
ASR & NLU Confidence scores	t_n^u
Length of user request and system Response	t_n^u, t_n^s
Time between consecutive user requests	$t_n^u - t_{n+1}^u$
Aggregate - Intent Popularity	t_n^u
Un-actionable user request	t_n^s
Cohesion between system response and user request	t_n^u, t_n^s
Length of dialogue	$t_0 - t_n$
Topic diversity	$t_0^u - t_n^u$
User paraphrasing his/her request	$t_n^u - t_{n+1}^u$

Table 7: Based on turn-level Gradient Boosting Regression model’s output, top 10 feature sets ⁶ranked by Importance score. In bold are the new feature sets we introduced.

Model	Hyper parameter - Optimal value
Lasso	alpha: 0.01
Decision Trees	max-depth: 2, min-samples-leaf: 5, min-samples-split: 2
Random Forest	max-depth: 4, min-samples-leaf: 8, min-samples-split: 13
Gradient Boosting Decision Trees	max-depth: 2, min-samples-leaf: 8, min-samples-split: 17

Table 8: Optimal Hyper parameter values used for training dialogue-level satisfaction Estimation models.

Feature set description	Turn(s) the feature set is computed on
Average estimated turn-level satisfaction ratings	$t_0 - t_N$
Average ASR confidence	$t_0 - t_N$
Last turn’s ASR confidence	t_N
Average aggregate - domain popularity	$t_0 - t_N$
Number of question prompts from the system side	$t_0^s - t_N^s$
Average time difference between consecutive utterances	$t_0^u - t_N^u$
Average NLU confidence	$t_0 - t_N$
Length of the dialogue	$t_0 - t_N$
Last turn’s aggregate intent popularity	t_N
Average aggregate - intent popularity	$t_0 - t_N$

Table 9: Based on dialogue-level Gradient Boosting Regression model output, top 10 features, ⁷ranked by Importance score.

⁶ For confidentiality we are not mentioning the entire list of feature sets used in the model.

⁷ For confidentiality we are not mentioning the entire list of features used in the model.