

Identifying Breakdowns in Conversational Recommender Systems using User Simulation

Nolwenn Bernard
University of Stavanger
Stavanger, Norway
nolwenn.m.bernard@uis.no

Krisztian Balog
University of Stavanger
Stavanger, Norway
krisztian.balog@uis.no

ABSTRACT

We present a methodology to systematically test conversational recommender systems with regards to conversational breakdowns. It involves examining conversations generated between the system and simulated users for a set of pre-defined breakdown types, extracting responsible conversational paths, and characterizing them in terms of the underlying dialogue intents. User simulation offers the advantages of simplicity, cost-effectiveness, and time efficiency for obtaining conversations where potential breakdowns can be identified. The proposed methodology can be used as diagnostic tool as well as a development tool to improve conversational recommendation systems. We apply our methodology in a case study with an existing conversational recommender system and user simulator, demonstrating that with just a few iterations, we can make the system more robust to conversational breakdowns.

CCS CONCEPTS

• Information systems → Recommender systems.

KEYWORDS

Conversational recommender systems; User simulation; Evaluation

ACM Reference Format:

Nolwenn Bernard and Krisztian Balog. 2024. Identifying Breakdowns in Conversational Recommender Systems using User Simulation. In *ACM Conversational User Interfaces 2024 (CUI '24)*, July 8–10, 2024, Luxembourg, Luxembourg. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3640794.3665539>

1 INTRODUCTION

Conversational recommender systems (CRSs) aim to provide personalized recommendations to user through multi-turn conversations [12, 20]. However, ensuring the robustness and effectiveness of these systems under any and all of the possible situations encountered while engaging with users remains a critical challenge [9]. We define a *breakdown* as a moment in the conversation where the flow discontinues or even stops. This includes system failures and unexpected/irrelevant replies from the CRS. Examples include the CRS recommending an item without knowing what the user

wants or prefers (i.e., recommendation before elicitation of the need and preferences) and the CRS repeating itself over and over. Previously, the dialogue breakdown detection challenge [17] motivated research on breakdown detection in chat-oriented dialogues. During the challenge, an annotated dialogue corpora was provided to the participants. Note that the set of labels used does not give information on the specific type of breakdown, despite the availability of at least one detailed breakdown taxonomy [16]. Consequently, most of the proposed solutions were supervised classification approaches determining if a (potential) breakdown occurred or not. Our work focuses on recommendation dialogues, i.e., task-oriented dialogues, that are not annotated. Furthermore, we aim to identify specific breakdowns in dialogues. Therefore, we propose a novel methodology to identify breakdowns and assess the robustness of an existing CRS via user simulation.

Leveraging user simulation has several advantages. First, it is a simple, cost-effective, and efficient solution to test a CRS that supplements human evaluation [38]. Second, it allows for a comprehensive assessment of a CRS's abilities in various scenarios and even allows for the simulation of the behavior of users with different characteristics (e.g., impatient, selective). In particular, our methodology identifies conversational paths (i.e., sequences of intents) that lead to conversational breakdowns. To achieve this, a set of conversations between the CRS and simulated users is analyzed with regards to a set of pre-defined breakdowns. As a starting point, we present detectors for three specific breakdowns to demonstrate our methodology: *system failure* (bugs), *dialogue of the deaf* (communication breakdowns between parties), and *conversational flow discontinuation* (disruptions with regards to a predefined interaction model). The identification of breakdowns in the conversational flow can provide insights to improve the robustness and effectiveness of the CRS. Furthermore, if changes to the behavior of the CRS can be made (e.g., either by updating the source code or providing it with further training examples), then the process can be repeated iteratively.

To demonstrate the usefulness of our approach, we present a case study with an existing CRS and user simulator, in which the conversational breakdown detection is performed on four subsequent versions of the CRS. More specifically, after each breakdown detection, a modification is applied to the CRS based on insights extracted (i.e., conversational patterns) with the goal to reduce the number of breakdowns, hence, improve the CRS. The results show that modifying the CRS with regards to one type of breakdown reduces its presence, while also affecting the number of other types of breakdowns detected. Moreover, we note that some breakdowns stem from imperfections in the user simulator. We demonstrate

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CUI '24, July 8–10, 2024, Luxembourg, Luxembourg

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0511-3/24/07...\$15.00

<https://doi.org/10.1145/3640794.3665539>

that our methodology can help improve the user simulator, parallel to the CRS, thereby making it a more robust and effective tool.

In summary, the main contributions of this work are twofold. First, we propose a methodology to identify conversational breakdowns in CRSs. Second, we present a case study where we apply this approach with an existing open-sourced CRS and user simulator. More specifically, we show how we improve the CRS based on the breakdowns identified.

2 RELATED WORK

Conversational recommender systems can be built following task-oriented dialogue system architectures or as end-to-end trainable systems [12]. In the former case, the *dialogue policy* is a central component, which determines the system’s response and action to a user utterance. Hence, a well-designed dialogue policy allows the CRS to seamlessly handle different situation such as unexpected actions from users and misunderstandings. The dialogue policy may be rule-based [3, 37] or model-based, commonly trained using reinforcement learning [23, 26, 32, 33]. Our methodology remains independent of the specific architecture of the CRS, and can be used with any of the aforementioned solutions.

Breakdowns in conversations have been studied in the field of dialogue systems and conversational analysis. It led to the development of different taxonomies of breakdowns [15, 16, 27, 29] and associated repair strategies [8], i.e., strategies used to overcome a breakdown. We group the breakdowns identified in [15, 16, 27, 29] into four categories which represent different aspects of a conversation (Table 1). In addition to these taxonomies, the BETOLD [34] and Dialogue Breakdown Detection Challenge (DBDC) [17] datasets were released to facilitate the development of breakdown detection methods. It is worth pointing out that these datasets are annotated with a small set of labels, which does not cover all types of breakdown mentioned in previous taxonomies. For example, in DBDC, there are three labels: not a breakdown, possible breakdown, and breakdown, which are given to each system utterance. While, in BETOLD, the labels are binary and at the conversation level. The dialogue breakdown detection challenge [17] ran for several years (between 2016 and 2019) and resulted in the development of different approaches to tackle this tasks in the context of chat-oriented dialogues most of them being based on supervised learning [17, 18], e.g., training of a recurrent neural network or support vector machine. Instead, in this paper, we focus on task-oriented dialogues. Although, chat-oriented and task-oriented dialogues have different characteristics, one can wonder if breakdown detection methods are generalizable. Lopes [25] studied this question by comparing the results of breakdown detection in task- and chat-oriented dialogues using the same set of features, e.g., cosine similarity between turns and bag-of-words representation of utterances. Based on this comparison, the author concludes that the hypothesis of a generic set of features to detect breakdowns “is true to a certain extent” due to the difference in interpretation of the features. In this work, we take a step forward and propose a methodology to detect specific types of breakdown instead of merely identifying the (potential) presence of a breakdown.

Breakdowns in conversational agents impact users’ satisfaction and their engagement [19]. Previous work [4, 10, 22] studied solutions to mitigate breakdowns in conversational agents, by focusing on strategies that can be applied by the conversational agent itself during the conversation to overcome breakdowns in a way that the user experience is impacted as little as possible. These strategies include acknowledgment of failure, rephrasing, and clarifying that are common in human communication [4]. Lee et al. [22] investigate the impact of different mitigation strategies, such as forewarning and apologizing, in the context of drink service by a robot. Their findings indicate that, in this particular context, each mitigation strategy may have a different influence on the users’ experience with the robot. In [10], the authors examine users’ feelings regarding different humorous repair strategies when performing a task with a smart speaker as the agent. The results of their user study show that the type of humor used by the repair strategy impacts users’ experience in terms of satisfaction, perceived intelligence, and likeability of the agent. Repair strategies are also studied from a user perspective [2, 5, 7], i.e., the user is the one attempting to overcome the agent’s breakdowns. These types of studies commonly perform user studies to collect conversations that are manually analyzed to identify which repair strategies are used. Unlike previous work on breakdown mitigation, our proposed method aims to automatically identify breakdowns and give insights on how to prevent them from happening, without the involvement of real users. Modifying the conversation agent based on these insights should reduce the use of repair strategies or even eliminate their need altogether.

User simulation can serve different purposes such as the optimization of a dialogue policy using reinforcement learning [31, 32, 35] and evaluation of conversational agents [13, 36, 38]. Different types of user simulators can be identified in the literature including model-based and data-driven [6]. User simulators for conversational recommendation are typically of the former type, commonly agenda-based [30, 33, 38], which allows for an explicit control over user behavior, i.e., what actions to perform with regards to a specific goal. Nevertheless, data-driven approaches [21, 24] have gained attention in the past years mostly due to the increase of large-scale datasets available. Our methodology is simulator-agnostic and can incorporate both solutions. However, due to their reliance on observed data, data-driven simulators have limitations when it comes to detecting breakdowns that happen along conversational paths that are less explored.

Möller et al. [28] proposed the MeMo system to automatically assess the usability of spoken dialogue systems. The system uses mental models to simulate users which can generate different types of errors. The generation of these errors is facilitated by the annotation of errors in historical dialogues. The aim of MeMo is to assess the impact of errors on the user experience rather than their detection. Indeed, the detection of errors is facilitated by the fact that they are generated by the user simulator. In early work, Eckert et al. [11] used user simulation for both automatic debugging and evaluation of conversational agents. In particular, they assess spoken dialogue systems and demonstrate that shortcomings and bugs can be identified by analyzing the length of generated dialogues. Differing from [11], our proposed methodology is more

Table 1: Categorization of breakdowns identified from the references [15, 16, 27, 29]. The categories considered in this work are marked with †.

Category	Description	Breakdowns
Linguistic	Breakdowns occurring due to language issues in an utterance	Syntactic error, semantic error, uninterpretable, wrong information, and command-level error
Contextual [†]	Breakdowns occurring when the utterance is not appropriate with regards to the context	Excess/lack of information, non-understanding, non-relevance, unclear intention, misunderstanding, excess/lack of proposition, contradiction, self-contradiction, non-relevant topic, topic switch error, and repetition
Social	Breakdowns related to social norms	Lack of common ground, lack of common sense, and lack of sociality
Functional [†]	Breakdowns related to the absence of a function or the inability to perform a function	Goal-level error, recognition-level error, and concept-level error

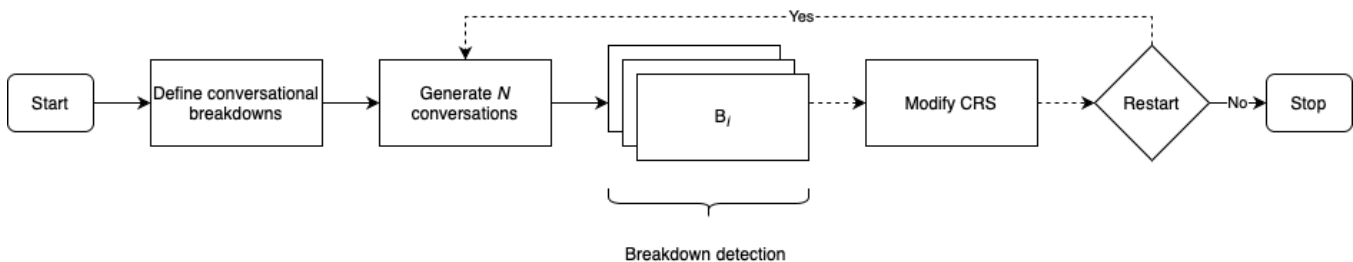


Figure 1: Flowchart of the proposed methodology, with dashed arrows denoting transitions when the methodology is used as a development tool.

generic (can detect multiple types of breakdowns that can be captured heuristically) and provides a more refined analysis of the breakdowns (conversational patterns based on dialogue intents). To the best of our knowledge, no other similar work on the debugging of conversational agents has been proposed more recently.

3 DETECTING CONVERSATIONAL BREAKDOWNS USING USER SIMULATION

This section presents the proposed methodology for the detection of conversational breakdowns along with a description and operationalization of the breakdowns considered.

3.1 Problem Statement

The problem studied in this work considers a conversational recommender system (CRS), also referred to as *agent*, and a user simulator (US), also referred to simply as *user*, engaged in a conversation with the aim of finding items that match the user’s preferences. Both the CRS and US are active *dialogue participants* in the recommendation process. We assume that the CRS aims to provide personalized recommendations by processing each incoming utterance and replying to it. Similarly, the US is expected to communicate the user’s preferences toward specific properties or items in the form of declarations or replies to the CRS. In our methodology, both the CRS and US are treated as ready-to-use components, i.e., no assumptions are made regarding their underlying architectures (e.g., component-based vs. end-to-end). For our diagnostic purposes, a conversation between the CRS and the US is represented as a sequence of turns, each

having an utterance accompanied by a corresponding *dialogue act*. A dialogue act comprises an intent and annotations taking the form of slot-value pairs (e.g., REVEAL (genre=action) is the dialogue act corresponding to the utterance “I like action movies”).

The objective of the proposed methodology is to identify conversational breakdowns in a CRS, that is, when the flow of a conversation is discontinued or stopped. Our methodology is designed to be generic, therefore, both the CRS and US are treated as ready-to-use components. The detection of conversational breakdowns is performed by a set of pre-defined breakdown detection components. These components indicate whether or not their associated breakdown occurs in a given conversation. Thus, by processing a sample of conversations between the CRS and US, they can identify recurring conversational patterns causing breakdowns. Based on the patterns identified as problematic, improvements may be made to the CRS to make it more robust and reliable.

3.2 Methodology

The proposed methodology considers three main components: the conversational recommender system, the user simulator, and breakdown detection. Our method is divided into four steps with the last one being optional, depending on whether it is used as a diagnostic or a development tool; see Fig. 1:

- (0) Define conversational breakdowns to detect and implement the associated detection components. In this work, we consider three particular conversational breakdowns that are described in Section 3.3.

- (1) Generate a set of N conversations between the CRS and US. The choice of N may depend on the use case or the complexity of the CRS evaluated. For example, interactions with a complex CRS can exhibit more conversational paths compared to a simpler CRS, hence, N should be high to explore as many paths as possible.
- (2) Run breakdown detection on the generated conversations. Each breakdown detection component B_i produces a summary that includes the number of breakdowns detected per conversational pattern.
- (3) (If used as a development tool) Modify the CRS based on insights from the previous step. Optionally, the system developer can repeat from step (1) depending on the number of breakdowns remaining. In this step, it is recommended to address a single breakdown at a time (as opposed to implementing multiple changes simultaneously), as it allows for a more accurate analysis of the effect of a modification on the CRS during the next iteration.

This method can serve as a diagnostic tool when the CRS is a black box that cannot be modified, i.e., step (3) cannot be applied. Alternatively, it might serve as a development tool when the dialogue policy, more generally the CRS, can be modified (e.g., by updating the source code). Indeed, by iteratively executing steps (1)–(3), one can investigate the impact of the modifications on the robustness and effectiveness of the CRS. We note here that step (1) could be ignored if enough historical data (i.e., recorded conversations) are available. Nonetheless, we advocate for the use of user simulation, as it provides the opportunity to quickly, inexpensively, and thoroughly assess the impact of the modifications made in step (3) through the generation of a large number of conversations, exploring diverse conversational paths.

3.3 Conversational Breakdown Detection

In this section, we describe the three specific conversational breakdowns considered in our use case. We note here that this set is by no means exhaustive and that our methodology is generic and extensible with other type of breakdowns that may be added in the future.

B_1 **System failure.** This detector simply verifies that no system errors are thrown during conversation generation between the agent and the user. This failure can help to identify underlying bugs that were not detected during integration tests.

B_2 **Dialogue of the deaf.** It identifies the agent’s utterances and dialogue acts that are identical and consecutive. This failure can indicate a pitfall in the dialogue policy from which the conversational agent cannot escape. Indeed, the goal of recommending an item to the user cannot be achieved if the agent keeps repeating the same utterance for multiple turns, i.e., the conversation does not progress. Moreover, such situations can easily frustrate users, who will most likely stop interacting with the agent. An example of this breakdown is presented in Fig. 2, where the agent fails to fill the slots “type of food” and “price range” with “thai” and “small budget,” respectively. This breakdown can be linked to the context-level error labeled *repetition* in [15], however, the definition of this error only mentions the textual similarity of

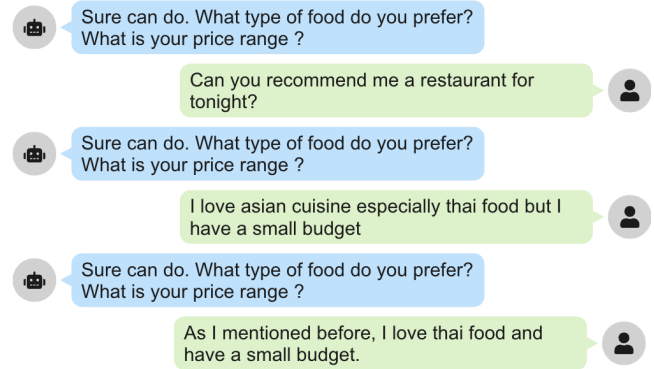


Figure 2: Example of a conversation section illustrating a dialogue of the deaf breakdown. The green and blue bubbles represent user and agent utterances, respectively.




Figure 3: Simplified interaction model for conversational recommendation (inspired by [14]). The blue and green states represent the agent and user, respectively.

utterances. The idea of repetition is also mentioned in [27, 29] but as a consequence of breakdown.

B_3 **Conversation flow discontinuation.** This failure corresponds to an unexpected reply from either of the dialogue participants, more specifically a reply that affects negatively the naturalness of the conversation. It includes delayed replies, i.e., the participant replies to an utterance that was received at least two utterances before, and unexpected responses, i.e., a sequence of utterances with dialogue acts that should not follow each other. For example, considering the simplified interaction model in Fig. 3, which defines the flow of recommendation in terms of user/system intents, it would be unexpected behavior from the CRS to provide information about an item that the user has previously rejected as a recommendation. This breakdown may be classified as a task-level error by the taxonomies proposed in [27, 29], as it indicates that an intent is invalid at a given state of the conversation. Additionally, we can consider the *topic transition* error described in [15] as a specific case of this breakdown because it includes delayed replies.

Some breakdowns might be detected on the system level (*system failure*), while others require semantic annotations on the utterance level (*dialogue of the deaf* and *conversational flow discontinuation*). Further, we categorize *system failure* as a functional breakdown, whereas *dialogue of the deaf* and *conversational flow discontinuation* are contextual breakdowns (cf. Table 1).

Algorithm 1 Detect *dialogue of deaf* breakdown

```
procedure DETECT_FAILURE(dialogue : Dialogue)
  path ← []
  if dialogue ended due to a RecursionError then
    for each utterance in dialogue do
      path.append(utterance.intent)
    end for
    Return path
  end if
  agent_utterances ← all agent utterances in dialogue
  for each agent_utterance in agent_utterances[: -2] do
    b_same_intent ← True if agent_utterance.intent is the
    same intent as the one of next two agent utterances
    b_similar_text ← True if agent_utterance is highly sim-
    ilar to the next two agent utterances
    if b_same_intent and b_similar_text then
      for each utterance in dialogue until agent_utterance
      do
        path.append(utterance.intent)
      end for
      Return path
    end if
  end for
  if no failure is detected, return an empty path
end procedure
```

3.4 Operationalization

In this section, we present the logic behind the implementation of detectors for the selected conversational breakdowns.

System failure. This detector simply checks the conversation’s error logs. In case an error different than RecursionError¹ is retrieved, the conversational path is flagged as problematic. RecursionError is associated with the *dialogue of the deaf* breakdown, that is, despite a large number of interaction the participants do not understand each other to accomplish the task.

Dialogue of the deaf. Similar to the previous detector, this one begins by checking if the conversation ended due to a RecursionError, in which case, the conversational path is marked as problematic. In the absence of a RecursionError, it examines the agent’s utterances in sequence of length three to verify if their intents are identical and texts are near identical, indicating a breakdown. Note that using a slicing window of three utterances is an arbitrary choice and could be modified. This choice is motivated by the idea that the agent may repeat the same question twice in case of a misunderstanding, but should not persist in doing so. The pseudo-code for the detection of this breakdown is presented in Algorithm 1.

Conversation flow discontinuation. This detector is initialized with an interaction model that can be represented as a directed graph, where nodes correspond to intents and edges to allowed transitions between two intents; see the example in Fig. 3. Then, following the conversation, the detector verifies if the conversational

¹The naming is Python specific, it is also known as Stack Overflow error in other languages such as Java and C++.

Algorithm 2 Detect *flow discontinuation* breakdown

```
procedure DETECT_FAILURE(dialogue : Dialogue)
  path ← []
  for utterance in dialogue do
    path.append(utterance.intent)
    if path is not a valid path in the dialogue flow graph then
      Return path
    end if
  end for
  if no discontinuity is detected, return an empty path
end procedure
```

path up to the current utterance exists in this graph. For example, if the conversation has five utterances, the detector checks if the first two intents are connected, in which case, it continues by adding the third intent to the path and so on, until the last utterance. The pseudo-code in Algorithm 2 illustrates the logic behind this detector.

4 CASE STUDY

To demonstrate the feasibility and value of the proposed methodology, a case study is conducted with an existing open-sourced CRS and user simulator. As the CRS is modifiable, we employ our methodology as a development tool to assess whether the modifications applied can decrease conversational breakdowns.

IAI MovieBot [14] is a conversational recommender system in the movie domain. It has a component-based architecture with four core components (Fig. 4a), namely, natural language understanding (NLU), dialogue manager (comprising dialogue policy and dialogue state tracking), recommendation engine, and natural language generation, that is very similar to the typical architecture found in [20]. The user utterance is processed by a rule-based NLU component. Then, the structured representation of the utterance is received by the dialogue manager that is responsible for deciding the next action of *IAI MovieBot* based on a pre-defined set of rules. It is connected to the recommendation engine, which is responsible for finding relevant items with regards to the dialogue state. Finally, the natural language generation component sends a template-based response to the user based on the next action chosen by the dialogue manager. For user simulation, we use *UserSimCRS* [1] that is designed for evaluation of the CRSs and implements an agenda-based simulator. Its architecture is based on a typical task-oriented dialogue system (Fig. 4b), and thus is relatively similar to that of *IAI MovieBot*, the main difference is the absence of recommendation engine. Two practical differences with *IAI MovieBot* can be pointed out. First, the NLU component is trained on a small corpora of dialogues between users and *IAI MovieBot*. Second, although the NLG component also uses templates, they are extracted from the corpora of dialogues, as opposed to being pre-defined. Agenda-based simulation assumes a Markovian state representation; additionally, the allowed state transitions are specified by an interaction model that is designed specifically for the task of conversational recommendations. The simulator determines the next action based on the agenda and the current context of the conversation.

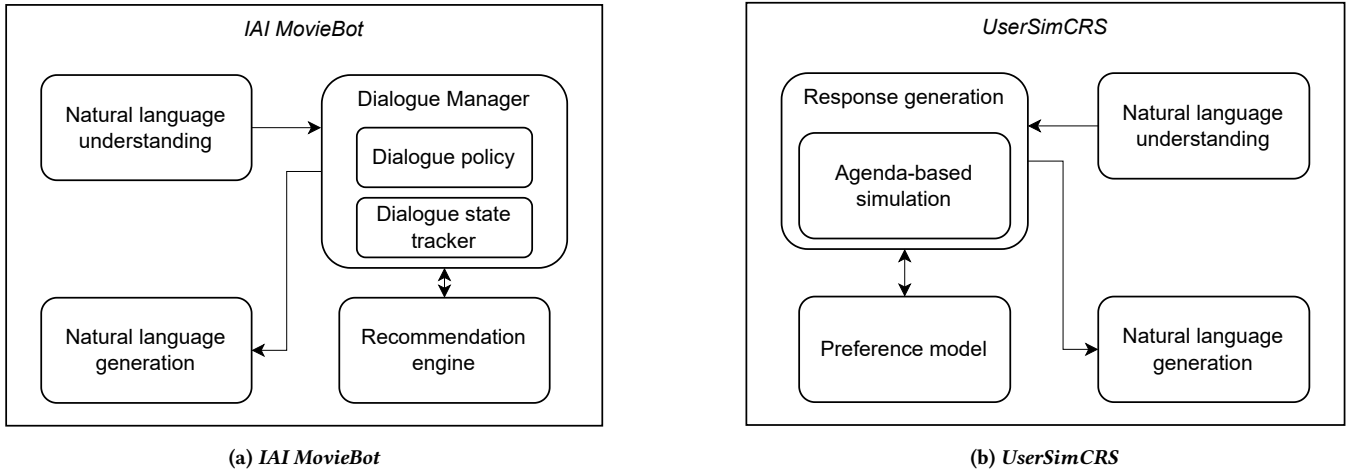


Figure 4: Architecture of dialogue participants: *IAI MovieBot* and *UserSimCRS*.

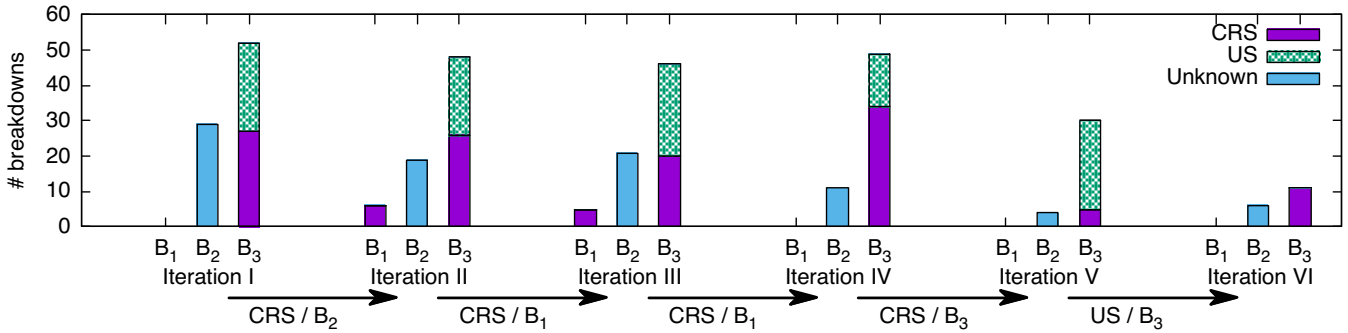


Figure 5: Conversational breakdowns per type for each iteration (groups). B_1 , B_2 , and B_3 represent system failure, dialogue of the deaf, and flow discontinuation, respectively. Colors indicate whether the breakdowns can be attributed to the CRS (purple), the US (green), or cannot be attributed (blue). Our aim in this work is to decrease the number of breakdowns especially those attributed to the CRS. The arrows between the iterations indicate which participant (CRS/US) was modified and what breakdown was targeted.

For this case study, we generate $N = 100$ conversations per iteration² as the action space of *IAI MovieBot* is relatively simple (i.e., 10 possible actions [14]). We perform five iterations; in each, we update *IAI MovieBot* by addressing one specific type of breakdown. Additionally, we conduct an extra iteration where we use an updated version of *UserSimCRS*. Fig. 5 presents a summary of the different types of breakdowns detected in each iteration. Moreover, the figure provides an overview of the breakdowns addressed in each iteration. We expect the number of breakdowns related to the one in focus to decrease. However, it is possible that the modifications applied lead to an increase or decrease in the other types of breakdowns. Additionally, we attempt to attribute each breakdown to either the CRS or the US as indicated by color-coding in Fig. 5. By analyzing the logs, we can find the source of a B_1 breakdown, while for B_3 , we investigate the conversational path, with regards to intents, leading to a breakdown to identify which participant

made a “forbidden” transition, i.e., one that is not allowed by the interaction model. For example, the US accepts a recommendation (i.e., intent ACCEPT) while the CRS is eliciting the US’s preferences (i.e., intent ELICIT). The responsibility for B_2 cannot be attributed to a specific participant as both participants fail to escape the misunderstanding loop they are in despite rephrasing utterances. We note here that the *UserSimCRS* is imperfect as illustrated in the responsibility distribution of B_3 breakdowns.

- (I) First, we perform conversational breakdown detection with the original version of *IAI MovieBot*. For this iteration, we notice that both participants (i.e., *IAI MovieBot* and *UserSimCRS*) are free of system failure, while a significant number of B_2 and B_3 breakdowns are detected. Therefore, we first decide to focus on reducing the number B_2 breakdowns by modifying *IAI MovieBot*. During the analysis of the simulated conversations, we note that this breakdown often appears when *IAI MovieBot* does not understand *UserSimCRS* and does not manage to find a way to pick up the conversation, as

²The generated conversations are made publicly available at: <https://github.com/NoB0/crs-breakdown-detection>

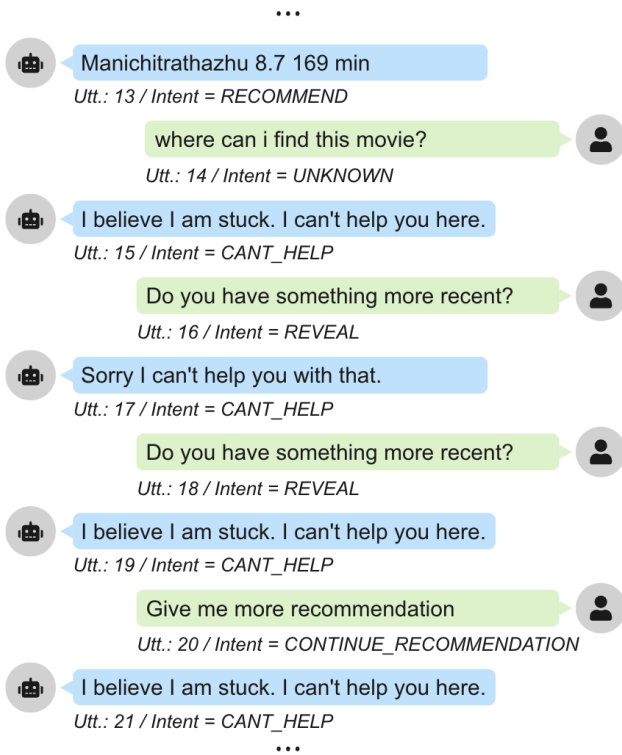


Figure 6: Excerpt of a conversation between *UserSimCRS* (in green) and *IAI MovieBot* (in blue) illustrating a B_2 breakdown, i.e., *IAI MovieBot* does not understand that *UserSimCRS* asks for an additional recommendation. In the entire conversation, we count the CANT_HELP intent 12 times in a row.

illustrated in Fig. 6. Hence, we apply a straightforward modification that consists in *IAI MovieBot* proactively restarting the conversation, using a predefined message, when such situation happens.

- (II) As expected, during the second iteration, we observe a decrease in the number of B_2 breakdowns. However, we note that some B_1 breakdowns are detected in *IAI MovieBot*. One possible explanation is that new conversational paths can be explored as *IAI MovieBot* is more robust to B_2 breakdowns. After looking at the summary of the breakdowns and the logs, we decide to fix a bug that appears when *IAI MovieBot* tries to conclude the conversation.
- (III) In this iteration, we observe a decrease in the number of B_1 breakdowns caused by *IAI MovieBot*. Interestingly, we find that the number of B_3 breakdown also decreases, while B_2 breakdowns slightly increase. We continue to focus on B_1 breakdown, hence, we fix a new bug discovered when *IAI MovieBot* wants to remove a preference from the user model.
- (IV) During this iteration, we do not detect B_1 caused by *IAI MovieBot*. This shows that the previous modifications helped to improve its the robustness. Unlike the previous iteration, the number of B_3 breakdowns increases, while we detect less B_2 breakdowns. For the last modification of *IAI MovieBot*, we

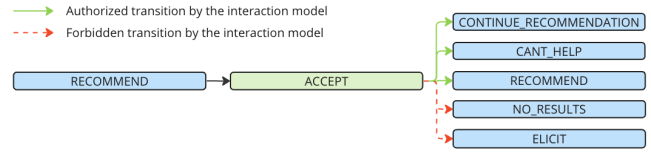


Figure 7: Partial conversational flows observed in simulated conversations after a recommendation is made by *IAI MovieBot*. The intent of utterances from *IAI MovieBot* and *UserSimCRS* are represented in blue and green, respectively.

shift our focus towards B_3 breakdowns. The analysis of the simulated conversations from the previous iterations reveals that this type of breakdown can appear when *IAI MovieBot* fails to appropriately reply to *UserSimCRS*'s utterance accepting a recommendation that is different from the predefined message "I like this recommendation." Some examples of such utterances are: "Sounds good I like it," "Sounds good" and "Ok I like this recommendation." Fig. 7 shows partial conversation flows after a recommendation is made; three out of the five flows observed are not allowed in the interaction model. Our approach to reduce B_3 breakdowns in the aforementioned situation is to improve the recognition of the ACCEPT intent by the natural language understanding component of *IAI MovieBot*.

- (V) The addition of new rules, i.e., patterns such as "like it," "love it," and "i like this," to improve the recognition of the ACCEPT intent results in a major drop of B_3 breakdowns caused by *IAI MovieBot*. However, the number of breakdowns caused by *UserSimCRS* increases. Therefore, we update *UserSimCRS* with a focus on B_3 breakdowns, more especially on *UserSimulation*'s agenda update. Originally, an agenda is initialized for the entire conversation and the next action is pulled off when the CRS replies in an expected manner, otherwise it is sampled [1]. However, we find that this could lead to B_3 , even if *IAI MovieBot* responds appropriately. For example, the intent REVEAL may be followed by ACCEPT in the agenda; yet, it may happen that no results were found after REVEAL, in which case it would not make sense to have ACCEPT as the user response. Therefore, we update *UserSimCRS* such that the the next action is always sampled.
- (VI) In this extra iteration, we observe a minor increase of B_2 and B_3 breakdowns. Interestingly, B_3 breakdowns caused by *UserSimCRS* are eliminated entirely after our last update.

Table 2 provides an analysis of the conversational paths explored across different iterations. We notice that in iterations I–IV there is a large number of number of unique conversational paths explored, which reduces later in iterations V and VI. Moreover, we see that the length of the conversations follows a similar pattern, i.e., it decreases after iteration IV. While there is no direct explanation for this, one hypothesis is that the modifications applied lead to more direct conversations, that is, less back and forth between the participants to give feedback for a recommendation. Additionally, the high standard deviation of conversation length hints at the diversity of the conversational paths explored. When examining the success of these conversational paths, i.e., both existing

Table 2: Analysis of conversational paths explored during each iteration. A conversation is deemed successful if at least one recommendation made by the CRS is accepted. Existing and new conversational paths are identified by comparison of conversational paths from the previous and current iterations. The average length of the conversations is calculated based on the number of utterances.

	# Unique conv. path	Avg. conv. length	Existing conv. path		New conv. path	
			Success	Not Success	Success	Not Success
Iteration I	96	19.84 ± 14.45			66	34
Iteration II	92	19.34 ± 11.02	7	2	59	32
Iteration III	93	19.29 ± 10.87	5	7	55	33
Iteration IV	98	19.42 ± 12.97	5	3	69	23
Iteration V	78	17.6 ± 10.77	17	7	56	20
Iteration VI	76	15.04 ± 9.02	26	4	31	39

and new paths, we make three main observations. First, there is a certain overlap between the conversational paths explored from one iteration to the next, but the majority of paths explored are new. Second, in most cases, more successful conversational paths are kept from one iteration to the next. Indeed, there is a higher proportion of successful paths within the existing ones. Third, we see that the number of new and unsuccessful paths decreases up until the last iteration. It indicates that the modifications applied to *IAI MovieBot* improve its ability to satisfy *UserSimCRS*'s needs. However, this is not the case after the update of *UserSimCRS* in the last iteration. The aforementioned observations support the idea that *IAI MovieBot* is becoming more robust and better at satisfying *UserSimCRS*'s needs with each iteration. Indeed, the modifications applied to *IAI MovieBot* lead to the exploration of diverse and new conversational paths, while providing better recommendations in different contexts.

In summary, this case study demonstrates that our methodology can be used to trace down different types of conversational breakdowns and to make *IAI MovieBot* more robust to them in a few iterations. Indeed, in two iterations (iterations II and III), we eliminate the *system failures* caused by *IAI MovieBot*. Furthermore, the modifications applied during iterations I and IV illustrate that *IAI MovieBot* can be improved in a more straightforward manner thanks to the identification of problematic conversational flows. It is noteworthy that each modification might also affect other types of breakdowns. For example, the modification applied to tackle *system failures* in iteration III increases the number of *flow discontinuation* detected in iteration IV. This can be explained by the fact that after a fix other conversational paths might be explored and new breakdowns can be uncovered. The analysis of the conversational patterns leading to *flow discontinuation* breakdowns shows that *UserSimCRS* is responsible for a large number of breakdowns. This highlights that despite being used to evaluate one CRS, the proposed methodology may also benefit the user simulator in parallel to improve its robustness and expose its limitations. Indeed, the last iteration shows that improving the quality of *UserSimCRS* affects the results. We note here that *flow discontinuation* breakdowns caused by *IAI MovieBot* are still detected, mostly in the situation where *UserSimCRS* acknowledges after an utterance from *IAI MovieBot*.

5 DISCUSSION

One main concern related to the use of user simulation for the evaluation of CRSs regards its quality; that is, the user simulator is most likely imperfect (e.g., overly simplified user model or context). Indeed, in the case of evaluation, the quality of the user simulation can impact the interpretation of the results obtained as demonstrated in the case study. Assuming that the user simulator can be controlled or modified (e.g., by updating its source code), it can be developed in parallel with the CRS to reduce the number of breakdowns and enhance its overall quality as previously argued in [1]. Having a more robust user simulator can lead to the exploration of new conversational paths, therefore, ensuring a better coverage of breakdown-free interaction space. We also note that our method would greatly benefit from future advances related to the open challenge of fairly representing all types of users in simulation [6]. Indeed, the analysis of conversational breakdown detection would be more comprehensive by simulating users outside the target user population as well.

A discerning reader may also ask the following question: if the dialogue policy is trained using reinforcement learning (RL), which relies on user simulation, what benefits does simulation offer for diagnosis? Even if the very same user simulator is employed in both cases, it can still serve as a diagnostic tool, ensuring that the dialogue policy has been trained sufficiently to handle various conversational paths effectively. Nonetheless, there is no requirement to use an identical user simulator. RL often employs data-driven user simulators, whereas we argue that a model-based simulator is more suitable for our objectives as it provides greater control over user behavior.

The conversational breakdown detectors described in this work are rather generic and are meant to provide a starting point to illustrate and validate our methodology. Indeed, these detectors can be used to test any type of CRS agent (i.e., component-based and end-to-end) as long as semantic annotations are available for user and agent utterances. One could also design new detectors that target more refined breakdowns, do not need semantic annotations, or are more specific to recommendation goals (e.g., recommending a single item vs. a set of items). A possible approach to develop these new detectors is to exploit previous observations/feedback from experts, beta testers, or users (if the CRS is deployed). Then, one can devise heuristics to capture the identified issues automatically. For example, if a CRS works with slot-value pairs to represent

preferences and the natural language understanding component shows some imperfections, a naive conversational breakdown could be the CRS asking preferences for the same slot over and over.

A major advantage of our methodology is that it is generic and can be used with any CRS architecture. There are, however, differences across CRSs in how the detected breakdowns may be addressed. In case of component-based systems, the dialogue policy (or its equivalent that is responsible for the conversational behavior of the system) needs to be changed directly; in a rule-based system this entails updating the source code, while in an RL-based system with a learned policy, it requires a re-training of the policy with an updated user simulator. As for end-to-end systems, where the system's behavior cannot be controlled directly but only via the training examples it is exposed to, addressing breakdowns might require the collection of new training examples (possibly with the help of user simulation) with the desired system behavior.

6 CONCLUSION

In this work, we have proposed a methodology that leverages user simulation to identify breakdowns in conversational recommender systems. It consists of the definition of conversational breakdowns and the implementation of their associated detectors. These detectors extract conversational paths (i.e., sequence of dialogue intents) leading to the related breakdowns. The detection of breakdowns is performed on a sample of conversations between the conversational recommender system (CRS) and user simulator (US). We capitalize on the fact that user simulation is a fast and cheap way of collecting a large sample of conversations to be analyzed. The methodology can be used as a diagnostic or a development tool. In the latter case, one can iteratively improve the CRS to reduce the number of breakdowns, thereby making it more robust. We have demonstrated this in a case study with an existing CRS and US. Indeed, in five iterations, we have successfully removed all *system failures* caused by the CRS and drastically reduced the number of other types of breakdowns. The case study also highlights the fact that the US is imperfect and may be responsible for some portion of the conversational breakdowns detected. In an extra iteration, we have shown that the methodology can also be used to improve the robustness of the US in parallel with the CRS.

Future directions include the application of our method to other CRSs and USs to ensure its adaptability. Furthermore, new conversational breakdown detectors can be created to provide a more extensive test of the CRS. This could include, for example, the detection of off-topic utterances that can disrupt the flow of the conversation.

ACKNOWLEDGMENTS

This research was supported by the Norwegian Research Center for AI Innovation, NorwAI (Research Council of Norway, project number 309834).

REFERENCES

[1] Jafar Afzali, Aleksander Mark Drzewiecki, Krisztian Balog, and Shuo Zhang. 2023. UserSimCRS: A User Simulation Toolkit for Evaluating Conversational Recommender Systems. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining (WSDM '23)*. 1160–1163.

[2] Francesca Alloatti, Francesca Grasso, Roger Ferrod, Giovanni Siragusa, Luigi Di Caro, and Federica Cena. 2024. A tag-based methodology for the detection of

user repair strategies in task-oriented conversational agents. *Computer Speech & Language* 86 (2024), 101603.

[3] Liliana Ardissono, Giovanna Petrone, and Marino Segnan. 2004. A Conversational Approach to the Interaction With Web Services. *Computational Intelligence* 20, 4 (2004), 693–709.

[4] Zahra Ashktorab, Mohit Jain, Q. Vera Liao, and Justin D. Weisz. 2019. Resilient Chatbots: Repair Strategy Preferences for Conversational Breakdowns. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. 1–12.

[5] Mirzel Avdic and Jo Vermeulen. 2021. Intelligibility Issues Faced by Smart Speaker Enthusiasts in Understanding What Their Devices Do and Why. In *Proceedings of the 32nd Australian Conference on Human-Computer Interaction (OzCHI '20)*. 314–328.

[6] Krisztian Balog and ChengXiang Zhai. 2023. User Simulation for Evaluating Information Access Systems. *arXiv cs.HC/2306.08550* (2023).

[7] Erin Beneteau, Olivia K. Richards, Mingrui Zhang, Julie A. Kientz, Jason Yip, and Alexis Hiniker. 2019. Communication Breakdowns Between Families and Alexa. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. 1–13.

[8] Dennis Benner, Edona Elshan, Sofia Schöbel, and Andreas Janson. 2021. What do you mean? A review on recovery strategies to overcome conversational breakdowns of conversational agents. In *International Conference on Information Systems (ICIS '17)*. 1–17.

[9] Hayet Brabra, Marcos Báez, Boualem Benatallah, Walid Gaaloul, Sara Bouguelia, and Shayan Zamanirad. 2022. Dialogue Management in Conversational Systems: A Review of Approaches, Challenges, and Opportunities. *IEEE Transactions on Cognitive and Developmental Systems* 14, 3 (2022), 783–798.

[10] Mikkel Clausen, Mikkel Peter Kyhn, Eleftherios Papachristos, and Timothy Merritt. 2023. Exploring Humor as a Repair Strategy During Communication Breakdowns with Voice Assistants. In *Proceedings of the 5th International Conference on Conversational User Interfaces (CUI '23)*. 1–9.

[11] W. Eckert, E. Levin, and R. Pieraccini. 1997. User modeling for spoken dialogue system evaluation. In *1997 IEEE Workshop on Automatic Speech Recognition and Understanding Proceedings*. 80–87.

[12] Chongming Gao, Wenqiang Lei, Xiangnan He, Maarten de Rijke, and Tat-Seng Chua. 2021. Advances and challenges in conversational recommender systems: A survey. *AI Open* 2 (2021), 100–126.

[13] David Griol, Javier Carbó, and José M. Molina. 2013. An Automatic Dialog Simulation Technique to Develop and Evaluate Interactive Conversational Agents. *Applied Artificial Intelligence* 27, 9 (2013), 759–780.

[14] Javeria Habib, Shuo Zhang, and Krisztian Balog. 2020. IAI MovieBot: A Conversational Movie Recommender System. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management (CIKM '20)*. 3405–3408.

[15] Ryuichiro Higashinaka, Masahiro Araki, Hiroshi Tsukahara, and Masahiro Mizukami. 2021. Integrated taxonomy of errors in chat-oriented dialogue systems. In *Proceedings of the 22nd Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL '21)*. 89–98.

[16] Ryuichiro Higashinaka, Kotaro Funakoshi, Masahiro Araki, Hiroshi Tsukahara, Yuka Kobayashi, and Masahiro Mizukami. 2015. Towards Taxonomy of Errors in Chat-oriented Dialogue Systems. In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL '15)*. 87–95.

[17] Ryuichiro Higashinaka, Kotaro Funakoshi, Yuka Kobayashi, and Michimasa Inaba. 2016. The dialogue breakdown detection challenge: Task description, datasets, and evaluation metrics. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC '16)*. 3146–3150.

[18] Chiori Hori, Julien Perez, Ryuichiro Higashinaka, Takaaki Hori, Y-Lan Boureau, Michimasa Inaba, Yuiko Tsunomori, Tetsuro Takahashi, Koichiro Yoshino, and Seokhwan Kim. 2019. Overview of the sixth dialog system technology challenge: DSTC6. *Computer Speech & Language* 55 (2019), 1–25.

[19] Mohit Jain, Pratyush Kumar, Ramachandra Kota, and Shwetak N. Patel. 2018. Evaluating and Informing the Design of Chatbots. In *Proceedings of the 2018 Designing Interactive Systems Conference (DIS '18)*. 895–906.

[20] Dietmar Jannach, Ahtsham Manzoor, Wanling Cai, and Li Chen. 2021. A Survey on Conversational Recommender Systems. *ACM Computing Survey* 54, 5 (2021), 1–36.

[21] Florian Kreyszig, Inigo Casanueva, Pawel Budzianowski, and Milica Gašić. 2018. Neural User Simulation for Corpus-based Policy Optimisation of Spoken Dialogue Systems. In *Proceedings of the 19th Annual SIGDial Meeting on Discourse and Dialogue (SIGDIAL '18)*. 60–69.

[22] Min Kyung Lee, Sara Kielser, Jodi Forlizzi, Siddhartha Srinivasa, and Paul Rybski. 2010. Gracefully Mitigating Breakdowns in Robotic Services. In *Proceedings of the 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI '10)*. 203–210.

[23] Wenqiang Lei, Xiangnan He, Yisong Miao, Qingyun Wu, Richang Hong, Min-Yen Kan, and Tat-Seng Chua. 2020. Estimation-Action-Reflection: Towards Deep Interaction Between Conversational and Recommender Systems. In *Proceedings of the 13th International Conference on Web Search and Data Mining (WSDM '20)*. 304–312.

- [24] Hsien-chin Lin, Nurul Lubis, Songbo Hu, Carel van Niekerk, Christian Geishauer, Michael Heck, Shutong Feng, and Milica Gasic. 2021. Domain-independent User Simulation with Transformers for Task-oriented Dialogue Systems. In *Proceedings of the 22nd Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL '21)*. 445–456.
- [25] José Lopes. 2017. How generic can dialogue breakdown detection be? the kth entry to dbdc3. In *Proceedings of dialog system technology challenge (DSTC6)*.
- [26] Tariq Mahmood and Francesco Ricci. 2009. Improving Recommender Systems with Adaptive Conversational Strategies. In *Proceedings of the 20th ACM Conference on Hypertext and Hypermedia (HT '09)*. 73–82.
- [27] Sebastian Möller, Klaus-Peter Engelbrecht, and Antti Oulasvirta. 2007. Analysis of Communication Failures for Spoken Dialogue Systems. In *INTERSPEECH 2007, International Speech Communication Association (ISCA) (INTERSPEECH '07)*.
- [28] Sebastian Möller, Roman Englert, Klaus-Peter Engelbrecht, Verena Vanessa Hafner, Anthony Jameson, Antti Oulasvirta, Alexander Raake, and Norbert Reithinger. 2006. Memo: Towards automatic usability evaluation of spoken dialogue services by user error simulations.. In *INTERSPEECH 2006 - ICSLP, Ninth International Conference on Spoken Language Processing (INTERSPEECH '06)*.
- [29] Antti Oulasvirta, Klaus-Peter Engelbrecht, Anthony Jameson, and Sebastian Möller. 2006. The relationship between user errors and perceived usability of a spoken dialogue system (*PQS '06*). 61–67.
- [30] Jost Schatzmann, Blaise Thomson, Karl Weilhammer, Hui Ye, and Steve Young. 2007. Agenda-Based User Simulation for Bootstrapping a POMDP Dialogue System. In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Companion Volume, Short Papers (NAACL '07)*. 149–152.
- [31] Jost Schatzmann, Karl Weilhammer, Matt Stuttle, and Steve Young. 2006. A survey of statistical user simulation techniques for reinforcement-learning of dialogue management strategies. *The Knowledge Engineering Review* 21, 2 (2006), 97–126.
- [32] Konrad Scheffler and Steve Young. 2001. Corpus-based dialogue simulation for automatic strategy learning and evaluation. In *Proceedings of NAACL Workshop on Adaptation in Dialogue Systems (NAACL '01)*. 64–70.
- [33] Yueming Sun and Yi Zhang. 2018. Conversational Recommender System. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval (SIGIR '18)*. 235–244.
- [34] Silvia Terragni, Bruna Guedes, Andre Manso, Modestas Filipavicius, Nghia Khau, and Roland Mathis. 2022. BETOLD: A Task-Oriented Dialog Dataset for Breakdown Detection. In *Proceedings of the Second Workshop on When Creative AI Meets Conversational AI (CAI '22)*. 23–34.
- [35] Bo-Hsiang Tseng, Yinpei Dai, Florian Kreyszig, and Bill Byrne. 2021. Transferable Dialogue Systems and User Simulators. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. 152–166.
- [36] Xiaolei Wang, Xinyu Tang, Wayne Xin Zhao, Jingyuan Wang, and Ji-Rong Wen. 2023. Rethinking the Evaluation for Conversational Recommendation in the Era of Large Language Models. *arXiv cs.CL/2305.13112* (2023).
- [37] Pontus Wärnestål. 2005. Modeling a dialogue strategy for personalized movie recommendations. In *Beyond Personalization Workshop*. 77–82.
- [38] Shuo Zhang and Krisztian Balog. 2020. Evaluating Conversational Recommender Systems via User Simulation. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '20)*. 1512–1520.