

Leveraging Interesting Facts to Enhance User Engagement with Conversational Interfaces

Nikhita Vedula Giuseppe Castellucci Eugene Agichtein
Oleg Rokhlenko and Shervin Malmasi
Amazon.com, Inc. Seattle, WA, USA
{veduln, giusecas, eugeneag, olegro, malmasi}@amazon.com

Abstract

Conversational Task Assistants (CTAs) guide users in performing a multitude of activities, such as making recipes. However, ensuring that interactions remain engaging, interesting, and enjoyable for CTA users is not trivial, especially for time-consuming or challenging tasks. Grounded in psychological theories of *human interest*, we propose to engage users with contextual and interesting statements or *facts* during interactions with a multi-modal CTA, to reduce fatigue and task abandonment before a task is complete. To operationalize this idea, we train a high-performing classifier (82% F1-score) to automatically identify relevant and interesting facts for users. We use it to create an annotated dataset of *task-specific interesting facts*¹ for the domain of cooking. Finally, we design and validate a dialogue policy to incorporate the identified relevant and interesting facts into a conversation, to improve user engagement and task completion. Live testing on a leading multi-modal voice assistant shows that 66% of the presented facts were received positively, leading to a 40% gain in the user satisfaction rating, and a 37% increase in conversation length. These findings emphasize that strategically incorporating interesting facts into the CTA experience can promote real-world user participation for guided task interactions.

1 Introduction

Conversational Task Assistants (CTAs) are a class of conversational agents that guide human users step-by-step in performing a multitude of activities, like cooking or home improvement tasks (Vtyurina and Fourney, 2018; Strathearn and Gkatzia, 2022; Choi et al., 2022; He et al., 2023; Agichtein et al., 2023). CTAs help users search and find the right task for their needs, inform about the task execution (e.g., tools to use), and answer questions about the task (e.g., substituting an ingredient).

¹www.github.com/vnik18/cta-interesting-facts

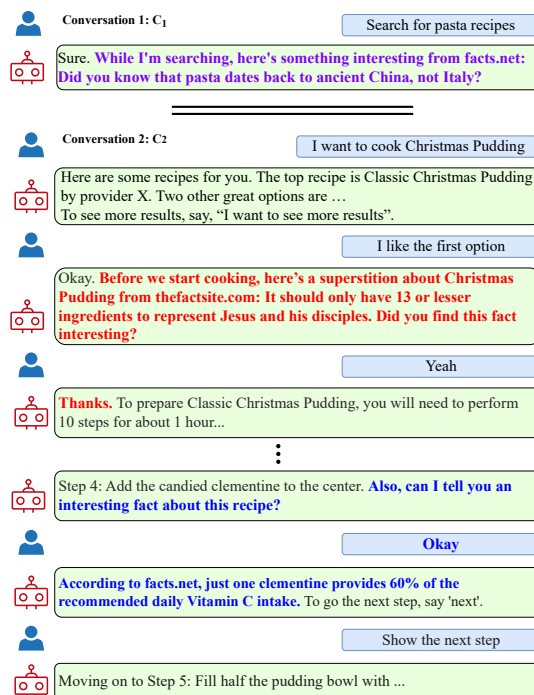


Figure 1: Conversations C_1 and C_2 between two users and our CTA. C_1 shows an interesting fact during search. C_2 has two interesting facts during task execution.

Besides ensuring that CTAs provide useful, accurate content, it is crucial to encourage users to consistently use CTAs to accomplish long and complex tasks. In this setting, users can easily get bored and abandon tasks, so maintaining interesting and entertaining conversations is key (Følstad and Brandtzaeg, 2020; Ceha et al., 2021). Understanding and predicting human interest is crucial to engage and retain users (Constantin et al., 2019).

We hypothesize that incorporating relevant, entertaining pieces of information into conversations is a powerful tool to maintain and sustain user attention, satisfaction, and participation. For example, while cooking with a CTA, users are likely to interact with it longer, frequently and repeatedly, if their experience is enriched with intriguing facts about the recipe's history, nutritional benefits of its ingre-

dients, or useful tips about its cooking techniques (see Figure 1 for an example). Such information can increase user knowledge, satisfaction and trust in the CTA by establishing it as a knowledgeable domain expert. It transforms users’ dialogue with the CTA into more than just task-oriented, resulting in an improved user experience, particularly during long, tedious, or monotonous activities.

We propose a practical approach to enhance user engagement with CTAs. Grounded in socio-psychological theories on human *interest* (Berlyne, 1949, 1970; Silvia, 2005), we design a novel schema to comprehensively and objectively examine the *interestingness* of pieces of information (henceforth called ‘*facts*’) in a CTA setting. We study if such facts can spark and foster user interest, and thus boost user satisfaction and engagement during a conversation with a CTA. We also explore multiple techniques to share interesting facts with users. We perform live testing of our framework in a leading multi-modal voice assistant. 66% of our facts were received positively by real-world users, resulting in a 40% gain in average user satisfaction ratings, and a 37% increase in average conversation length. These results demonstrate the practical benefits of enriching the CTA experience with interesting facts, and showcase our approach’s potential to enhance the adoption of CTAs for complex real-world tasks requiring sustained user engagement.

We also create for public release a dataset of 1,379 interesting facts associated with cooking, extracted from online sources and cleaned for grammar, style, relevance, interestingness, and appropriateness. We believe that this can be a valuable resource to drive future research in CTAs, and user engagement in task-assistance conversations.

2 Related Work

Interest and Interestingness: Our work on enhancing user engagement with CTAs by increasing their level of interest in the conversation is inspired by multiple psychological theories conceptualizing the human behavioral dimensions of interest and interestingness. One of the earliest theories is that of Berlyne (Berlyne, 1949, 1970), who viewed interest as a defining factor of motivation, influenced by several factors including uncertainty, complexity and novelty of the subject at hand. Tomkins (1962); Silvia (2005) defined interest as an appealing effect of an activity on a human caused by novelty,

exploration, learning, creativity and affect, leading to increased interaction with the subject matter.

Geng and Hamilton (2006) surveyed multiple features relevant to interestingness such as conciseness, coverage, reliability, novelty, surprise and utility in the context of pattern mining and association rules. Carvalho et al. (2005) studied their correlation with real human interest. Wu and Miao (2013); Constantin et al. (2019) observed a negative correlation between interest and low-arousal emotions like memorability, familiarity, and boredom. Clark et al. (2019); Ceha et al. (2021) employed humor, jokes and fun facts within their conversation agents to demonstrate that providing such entertaining information to users can increase user interest, engagement and retention, and reduce monotony.

In our work, we model several aspects contributing to human interest in Section 3, to identify facts that can interest and engage users interacting with a CTA, and reduce monotony. We show in Section 4 that a CTA without interest enhancing signals can cause boredom and reduce engagement by users.

Conversational Task Assistants: CTAs are a special type of conversational assistants (Agichtein et al., 2023), different from general purpose intelligent assistants (e.g., Amazon’s Alexa, Apple’s Siri). CTAs guide users to perform tasks in a step-wise manner, also answering user questions along the way. CTAs differ from conversational search agents (Vtyurina et al., 2017) and task-oriented conversational agents (Zhang et al., 2020; Strathern and Gkatzia, 2022), where the agent and not the user performs a task based on user input (e.g., booking a ticket). User satisfaction and engagement has been widely studied in the areas of search, task-oriented dialog systems and virtual assistants (Kiseleva et al., 2016; Park et al., 2020; Sekulić et al., 2021; Lotze et al., 2021; Siro et al., 2022). However, these approaches lack a theoretical backing and have not been studied for CTAs.

The work that bears the most similarity with ours is Wizard of Curiosities (WoC) (Vicente et al., 2023), which also adds facts within dialogues between users and CTAs. However, both works differ in the following ways: (i) WoC targets *information-seeking* traits of humans based on *curiosity* theory (Berlyne, 1966; Kidd and Hayden, 2015). Our work is more principled and draws upon the broader theories of *interest* which encompass informativeness and curiosity, and are better suited for studying and improving user engagement (Pekrun, 2019);

(ii) WoC does not consider several features that we consider in this work to add facts into the CTA dialogue flow, namely novelty, recency, frequency, interest and emotion (defined in Section 3). (iii) Unlike WoC, we translate the theoretical aspects into a novel annotation schema for interestingness that is more objective, consistent and precise for a CTA (Sections 3.1 and 3.3). (iv) We study several user experience options on how to effectively present users interesting facts during the conversation flow.

3 Boosting CTA Engagement with Facts

We now describe our framework to offer interesting facts to users conversing with a multi-modal CTA, having both display and voice components, to entertain users and increase their engagement. We focus on cooking-specific conversations, but as explained in Section 5, we believe that our framework can generalize to any CTA-feasible domain.

3.1 Feature Schema for Interestingness

We first define a feature schema inspired by socio-psychological work on interestingness (Berlyne, 1949, 1970), and use it to curate a dataset of facts that may increase user interest and lower monotony.

(i) Conciseness: A fact should not be too long in terms of number of words and listening time.

(ii) Specificity: A fact should be understandable in general to a large number of users, without requiring specific background or domain knowledge.

(iii) Novelty: A fact should offer new, unique information that may surprise an average user not a knowledge expert in the domain.

(iv) Relevance: A fact should be relevant and related to the considered conversation turn.

(v) Informativeness: A fact should deliver useful, helpful knowledge or insights about the task.

3.2 Interesting Facts Dataset

We create a dataset of 1,379 interesting facts associated with cooking, following the above feature schema. We use an off-the-shelf named entity recognition model to extract a list \mathcal{E} of entities of the types *ingredient*, *recipe*, and *tool*, from a large corpus of recipes \mathcal{R} .² We then crawl specific websites³ to extract diverse candidate facts for \mathcal{E} .

To filter out irrelevant facts, we split each fact into individual sentences. We then fine-tune and use a domain classifier, consisting of a linear layer

on top of the 5B Alexa Teacher Model (FitzGerald et al., 2022), with positive training sentences from \mathcal{R} , and negative instances from out-of-domain online sources.⁴ We match each of the domain-relevant candidate facts to our list of identified entities \mathcal{E} . We only retain facts if they contain an entity in \mathcal{E} as either the subject or the object of the fact sentence, based on its dependency parse tree.⁵ This gives us a set of potential interesting facts \mathcal{F}' .

We next sample a set of 750 facts from \mathcal{F}' , to be manually annotated for their level of *interestingness*. For each fact, 2 expert annotators familiar with this domain give it a binary label with respect to each feature of our schema (Section 3.1), followed by an overall label for interestingness with respect to the entity under consideration. We also ask annotators to choose the feature most important to them to decide if a fact selected by a CTA could be interesting. Having such a structured schema can mitigate the subjectivity and/or biases likely to occur while annotating the inherently subjective element of *interestingness*, giving more objective, consistent, and precise annotations.

We find that annotators prefer novel and easily understood facts. Aggregating their preferences yields this ranked order of feature importance for relevant facts: novelty, specificity, conciseness and informativeness, with a per-feature binary inter-annotator agreement of 0.68, 0.84, 0.88 and 0.75 respectively. We assign each feature a normalized weight based on how many times they were chosen as important. We then compute a linearly combined interestingness score for each fact, resulting in a set \mathcal{F} of 606 annotated interesting facts. These facts serve as positive instances, and irrelevant or uninteresting facts from \mathcal{F}' serve as negative instances, to train a multi-label binary classifier on top of a RoBERTa (Liu et al., 2019) language model. The two labels denote relevance and interestingness to the input text (the current step of the ongoing task).

We use this classifier to automatically label the remaining domain-relevant candidate facts in \mathcal{F}' . Each fact is linked to a step in the corpus \mathcal{R} , based on the entities present in the task step. To minimize redundancy and ensure that the facts associated with each step are sufficiently diverse, we cluster and filter facts with a high (≥ 0.85) pairwise cosine similarity based on Sentence-BERT (Reimers and Gurevych, 2019) embeddings. Each fact is then

²The recipe provider in this work is Whole Foods Recipes.

³Wikipedia, and fun facts websites like www.facts.net

⁴The model achieves a 0.91 F1-score on a 13.8K test set.

⁵We use www.spacy.io to compute the parse tree.

manually cleaned for grammar, style and appropriateness. We finally obtain a novel dataset of 1, 379 interesting cooking facts (see dataset details in Table 1), to present to real-time customers conversing and performing a task with a CTA. Each fact has accompanying evidence in the form of the website URL it is sourced from.

Algorithm 1: Presenting interesting facts to a user interacting with a CTA during search.

Data: Facts associated with each task

```

1  $max\_facts \leftarrow initialize;$ 
   $min\_turns\_btw\_facts \leftarrow initialize;$ 
2 Function showFactAtStep():
  // Show user fact at this step?
3 if fact exists for step and # facts already shown <  $max\_facts$  and # turns betw. facts  $\geq$   $min\_turns\_btw\_facts$  and final turn is voice-friendly then
4   | return True;
5 else
6   | return False;
7 User searches for a task
8 if showFactAtStep() then
9   | Present current search result info;
10  | Show interesting fact & record count;
11 Repeat steps 8-10 until the user chooses a task to execute or exits conversation;
```

3.3 Incorporating Interesting Facts into Conversations

The interestingness feature schema described in Section 3.1 only includes *fact-specific* features, to recognize a given fact as interesting. Recall that a CTA aids users explore or search for tasks, provides information about tasks, and guides users through step-by-step execution instructions. For a satisfying user experience (UX), several *conversational* dimensions need to be considered to effectively inject interesting facts into user-CTA interactions. We consider alternative dialogue policies:

(i) **Present Facts Before & During a Task:** We offer interesting facts to users both when they are exploring or searching for a task to perform, and during the step-by-step execution of their chosen task. Algorithms 1 and 2 show both these UXs.

(ii) **Select Turns to Present Facts:** Engaging users is crucial, but the goal of a user-CTA interaction is to complete a task (e.g., cook a recipe). We

use the below criteria to avoid creating distractions and a poor experience by showing too many facts.

Algorithm 2: Presenting interesting facts while a user executes a task with a CTA.

Data: Facts associated with chosen task

```

1  $max\_facts \leftarrow initialize;$ 
   $min\_turns\_btw\_facts \leftarrow initialize;$ 
2 Function seekUserFeedback():
  // Seek user feedback about fact
3 if fact shown and feedback not sought before then
4   | Ask user if fact was interesting;
5   | Record user response;
6 Step-by-step task execution
7 if showFactAtStep() then
8   | Present current step text;
9   | Ask user if they want to see a fact;
10  | If user agrees, show fact & record count;
11  | If user rejects, set  $max\_facts = 1$ ;
12 seekUserFeedback() about fact shown
13 Repeat steps 6-12 until user completes the task execution or exits the conversation;
```

- **Recency:** After presenting a fact, the next one is shown only after a set number of turns.
- **Frequency:** A CTA should not present too many facts per conversation. The recency and frequency parameters are implemented in the showFactAtStep() function in Algorithms 1 and 2, and their values are chosen empirically.
- **Diversity:** Our dataset creation in Section 3.2 ensures that users are shown diverse facts.
- **Voice-friendly:** The word count of all conversation turns with interesting facts is bounded.

(iii) **Proactive Inquiry:** A CTA explicitly asking a user if they want to see an interesting fact or not adds two extra turns in the conversation, but reduces the possibility of distracting or frustrating users with undesirable information. When we present facts to users during their search (step 10 of Algorithm 1, first two turns of Figure 1), we do not seek prior permission from them, to avoid additional dialogue overhead before the task even begins. However, during task execution, our CTA seeks user approval each time before showing interesting facts (steps 9-10 of Algorithm 2). Users

Number of interesting facts	1,379
Number of unique fact providers	5
Number of unique entities	420
Mean fact length in words	13

Table 1: Details of our curated interesting facts dataset.

Model	F1-relevance	F1-interestingness	F1-rel-int
roberta-base-rel	0.76	N/A	0.76
roberta-base-int	N/A	0.67	0.67
roberta-base-rel-int (Ours)	0.83	0.80	0.82
gpt-3.5-turbo-fs-rel	0.84	N/A	N/A
gpt-3.5-turbo-fs-int	N/A	0.72	N/A
gpt-3.5-turbo-fs-rel-int	0.78	0.66	0.72

Table 2: Detecting relevance and interestingness of facts. ‘rel’, ‘int’ and ‘fs’ denote relevance, interestingness, and few shot respectively. N/A means that a model does not predict the corresponding label.

can choose to not answer or bypass this request, in which case the next step of the task continues.

(iv) **Seek User Feedback on Facts:** We explicitly ask the user *once* during the conversation if they liked the fact that was shown to them (seekUserFeedback() function in Algorithm 2). Users can choose to not answer or bypass this question, and continue to execute their task normally.

4 Experiments and Results

4.1 Interesting Fact Quality Evaluation

Relevance and Interestingness Detection: In Section 3.2, we proposed a multi-label classifier to detect relevant and interesting facts. We now assess multiple approaches in detecting relevance and interestingness in Table 2, on a test set of 400 facts with a 50% split of positive and negative candidates. Our proposed classifier (roberta-base-rel-int) improves upon the single-label classifier baselines (roberta-base-rel, roberta-base-int) by at least 5% F1 score points, in detecting if an input fact is relevant and interesting. We also prompt the gpt-3.5-turbo (OpenAI, 2023a) LLM with 8 randomly chosen few-shot examples from our training set to recognize relevance and interestingness individually (gpt-3.5-turbo-fs-rel and gpt-3.5-turbo-fs-int), as well as together in the same prompt (last row of the table). Our multi-label classifier is comparable to GPT-3.5 in detecting relevant facts despite being a much smaller model. Moreover, it significantly outperforms GPT-3.5 by 10% in detecting the interestingness of facts.

Evaluating LLMs for Annotation: We prompt GPT-4 (OpenAI, 2023b) with our well-defined interestingness feature schema from Section 3.1.

Similar to the human annotation in Section 3.2, we ask GPT-4 to label (i) each input fact for specificity, novelty, relevance and informativeness; (ii) overall interestingness based only on the schema features; and (iii) the most important feature to decide a fact’s interestingness. On 75 randomly selected facts from Section 3.2, we observe 44% agreement with human labels for overall interestingness, showing that it is non-trivial to replace humans with LLMs for this task. The highest agreement (> 58%) is observed for ‘specificity’ and ‘relevance’. The lowest agreement of 31% is seen for the ‘novelty’ feature, likely due to its inherent subjectivity. GPT-4 chose ‘informativeness’ as the key interestingness indicator, in contrast to the ‘novelty’ feature chosen by human annotators in Section 3.2.

Interesting Facts and User Engagement: We first evaluate the quality of the interesting facts in the dataset we created in Section 3.2, and their effect on the user engagement with a CTA. We provide crowd workers with 200 sampled facts,⁶ and ask them if the fact is (i) interesting to them; and (ii) likely to engage users interacting with a CTA. In this experiment, we do not provide our feature schema from Section 3.1 to the annotators, letting the definition of ‘interestingness’ remain subjective and open to interpretation. 89% of the input facts were chosen as both interesting and engaging to CTA users, reinforcing the effectiveness of our feature schema in defining interestingness, and the utility of our interesting fact dataset.

4.2 Conversation Quality Evaluation

We randomly select 50 user-CTA cooking conversations, having 63 turns with interesting facts. We ask human judges if the facts within the conversations might bore, frustrate or distract users from their original goal. The judges negatively labeled 22% of the turns with facts, possibly because those facts were not appealing enough to justify the additional fact-related turns (examples in Appendix A).

Interesting Fact Placement in Conversations:

We compare different strategies of inserting facts during a conversation between users and CTAs. We select 20 seed conversations, and present the following variants to crowd workers:⁷ (i) showing interesting facts only during search, before the user selects a task; (ii) showing interesting facts only

⁶Average binary inter-annotator agreement = 0.68

⁷Average binary inter-annotator agreement = 0.60

during task execution; (iii) always showing facts after seeking user approval; (iv) always showing facts without seeking user approval; and (v) our proposed hybrid fact placement approach in Section 3.3. In 65% of cases, human judges preferred hybrid conversations where facts were shown both before task selection and during execution. Asking user permission before showing a fact was preferred 60% of the time. This further validates the UX design choices we made in Section 3.3.

4.3 Live Testing

Setup: We test our proposed interesting facts feature experience within a CTA of a leading commercial voice assistant, for live interaction and use by thousands of users in the US. Our CTA was built on top of the base Alexa Prize infrastructure (Agichtein et al., 2023). The design, implementation and evaluation details of our final set up are described in Sections 3.2, 3.3, 4.1 and 4.2. We study two versions of the CTA, i.e., with and without the interesting facts feature, via A/B testing over a period of 6 weeks. We obtain explicit real-user feedback on our system in the following ways: (i) we ask users if they liked the interesting facts shown to them; (ii) users rate their experience conversing with the CTA with a score from 1 to 5; and (iii) users provide additional feedback via a free-form text field at the end of the conversation.

User Satisfaction and Engagement: We analyzed 450 user-CTA conversations, of which 300 belonged to the control group, and 150 were exposed to our interesting facts UX. Users stated that they liked about 66% of the interesting facts shown to them during the conversation. The average satisfaction rating provided by users to their conversations grew by 40% (p -value < 0.05 with a t -test) over the set up without the interesting facts.

The average conversation length between users and the CTA grew by 37% (p -value < 0.05 with a t -test). This is a favorable outcome overall, because while longer conversations can increase the amount of time spent by users to complete their task, shorter conversations are also likely to indicate task abandonment or user disinterest in continuing the conversation. Recall that the goal of our work is to maximize users' engagement with the CTA, since this can make our CTA more appealing for further adoption and usage (Sections 1 and 2). In fact, we noted a rise in the number of users both starting a task, and continuing their guided task interaction,

after seeing an interesting fact. We also observed a 5% increase in the task completion rate (i.e, the users completing all execution steps of their chosen task), and a 5% increase in the repeat users interacting with the CTA after testing this feature.

Finally, we examined users' next turn responses to the provided interesting facts, excluding those who explicitly mentioned liking the facts. We did not find any users abruptly ending their conversations after being presented with an interesting fact, or interrupting the CTA via voice as it spoke the fact. Manual inspection showed a positive or neutral sentiment in $> 70\%$ of these user responses.

5 Discussion

Scope and Generalizability: The goals of our work are two-fold. The first is the generation of factually accurate interesting facts that are novel, specific, concise, informative and relevant to the task at hand. The second, more complex goal of our work is strategically injecting interesting facts into a user's conversation with a CTA, and involves multiple steps (Section 3.3). In this paper, we empirically evaluate our proposed approach only on the domain of cooking. However, we believe that the various steps and considerations involved in our framework, namely, our theoretically backed interestingness feature schema, fact curation and cleaning, and conversational presentation techniques can serve as an inspiration or a useful starting point for other applicable conversational task domains to consider, evaluate and expand upon.

Scalability: Though potentially applicable across domains where interesting facts are available, our interesting fact dataset creation approach presented in Section 3.2 may not scale due to the manual annotation effort involved. We investigated the use of LLMs to annotate interesting facts in Section 4.1. We next explore the automatic generation of interesting facts using an LLM (gpt-3.5-turbo). We sample 50 entities from set \mathcal{E} (Section 3.2), and prompt GPT-3.5 to generate 50 relevant and interesting facts as per our schema in Section 3.1. The facts generated by GPT-3.5 are suitable for a user conversing with a CTA. However, we found that 16% of them are either factually inaccurate, or of questionable accuracy with no supporting evidence available. For instance, the fact: *The Can Opener Museum located in Cincinnati, Ohio, celebrates the history and*

evolution of can openers is completely fabricated – such a museum does not exist.

Therefore, directly showing unverified LLM-generated facts to users in an online manner can severely impair users’ experience and trust in a commercial CTA that aims to be a reliably and knowledgeably guide users through costly and lengthy tasks. Online fact generation by an LLM during a conversation also increases the CTA’s response latency. One solution, which we use in this work, involves offline correction and retrieval of facts before presenting to users. Another solution could be to use feedback dataset as demonstration data to tune an LLM or an explicit dialogue policy.

Using LLMs in commercial settings also poses a practical challenge, as legal requirements mandate appropriate crediting of original information sources. When prompted for the original sources it uses to generate facts, GPT-3.5 responds: *I’m unable to directly access external websites or provide specific sources for facts as my knowledge.* We will thus need to provide the LLM a retrieval or API-based search component to get these sources automatically, which adds a risk of hallucination. We leave the end-to-end automation of our approach with one or more LLMs for future work.

6 Conclusions

Drawing on psychological theories on human interest, this work defines what users might find interesting during guided task-based conversations with CTAs, and empirically demonstrates the value of engaging and satisfying users with appealing external knowledge. Engaged users are, in fact, more likely to complete tasks and leverage the power of CTAs for complex task assistance. To achieve this goal, we defined a semi-automatic framework to build a dataset of 1,379 contextual and interesting facts for the cooking domain, and investigated different options to effectively inject the facts into user-CTA conversations. Offline human evaluations and live testing on a multi-modal voice assistant showed significant gains in user satisfaction and engagement in a CTA enriched with our interesting facts experience.

Limitations and Future Work

Our work did not consider the following aspects, which we leave for future work.

Domain Generalization: Based on the theoretical backing of our framework, we fully expect it to

generalize across multiple task domains feasible for a CTA. However, we would like to verify this empirically, by generating interesting facts for other task assistant domains such as home improvement, inserting those facts into user-CTA conversations, and addressing any potential domain-specific challenges that may occur, e.g. recognizing and linking complex named entities (Malmasi et al., 2022).

Fact Presentation: Our current approach only considers the presentation of interesting facts in the form of simple text. Our multi-modal CTA currently shows the entire fact text on the screen, and reads out the entire fact by voice. We seek to explore varied modes of presentation combining both voice and the display screen, such as diverse text formats, catchy summary phrases, or visual aids to accompany the facts.

Utilization of LLMs: We plan to reduce the number of steps and manual effort involved in our approach to increase its scalability and generalizability, by employing one or more LLMs to automate the process in an end-to-end fashion. This involves using a retrieval augmented LLM with explicit hallucination reduction techniques to generate factually accurate and interesting facts based on our feature schema. We will then use the insights we discovered in this work as well as any existing UX specifications or constraints, to either prompt or fine-tune an LLM to decide when, where and how best to naturally incorporate interesting facts into task-specific conversations.

Furthermore, though we explored the ability of LLMs to annotate as well as generate interesting facts, we did not directly present the factually accurate interesting facts generated by LLMs to real-world users interacting with our CTA. Given the disagreement on interestingness between human and GPT-4’s annotations in Section 4.2, we would like to investigate the user engagement and user satisfaction obtained when users are presented facts generated and selected by GPT-4, and how this differs compared to the interesting facts preferred by humans annotators.

Acknowledgements

We would like to express our sincere gratitude to Zhiyu Chen, Jason Choi, Saar Kuzi as well as the anonymous reviewers for their insightful feedback and comments on our work.

References

- Eugene Agichtein, Michael Johnston, Anna Gottardi, Lavina Vaz, Cris Flagg, Yao Lu, Shaohua Liu, Sattvik Sahai, Giuseppe Castellucci, Jason Ingyu Choi, et al. 2023. Advancing conversational task assistance: the second alexa prize taskbot challenge.
- Daniel E Berlyne. 1949. Interest as a psychological concept. *British Journal of Psychology*, 39(4):184.
- Daniel E Berlyne. 1966. Curiosity and exploration: Animals spend much of their time seeking stimuli whose significance raises problems for psychology. *Science*, 153(3731):25–33.
- Daniel E Berlyne. 1970. Novelty, complexity, and hedonic value. *Perception & psychophysics*, 8(5):279–286.
- Deborah R Carvalho, Alex A Freitas, and Nelson Ebecken. 2005. Evaluating the correlation between objective rule interestingness measures and real human interest. In *Knowledge Discovery in Databases: PKDD 2005: 9th European Conference on Principles and Practice of Knowledge Discovery in Databases, Porto, Portugal, October 3-7, 2005. Proceedings 9*, pages 453–461. Springer.
- Jessy Ceha, Ken Jen Lee, Elizabeth Nilsen, Joslin Goh, and Edith Law. 2021. Can a humorous conversational agent enhance learning experience and outcomes? In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–14.
- Jason Ingyu Choi, Saar Kuzi, Nikhita Vedula, Jie Zhao, Giuseppe Castellucci, Marcus Collins, Shervin Malmasi, Oleg Rokhlenko, and Eugene Agichtein. 2022. Wizard of tasks: A novel conversational dataset for solving real-world tasks in conversational settings. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 3514–3529.
- Leigh Clark, Nadia Pantidi, Orla Cooney, Philip Doyle, Diego Garaialde, Justin Edwards, Brendan Spillane, Emer Gilmartin, Christine Murad, Cosmin Munteanu, et al. 2019. What makes a good conversation? challenges in designing truly conversational agents. In *Proceedings of the 2019 CHI conference on human factors in computing systems*, pages 1–12.
- Mihai Gabriel Constantin, Miriam Redi, Gloria Zen, and Bogdan Ionescu. 2019. Computational understanding of visual interestingness beyond semantics: literature survey and analysis of covariates. *ACM Computing Surveys (CSUR)*, 52(2):1–37.
- Jack FitzGerald, Shankar Ananthkrishnan, Konstantine Arkoudas, Davide Bernardi, Abhishek Bhagia, Claudio Delli Bovi, Jin Cao, Rakesh Chada, Amit Chauhan, Luoxin Chen, et al. 2022. Alexa teacher model: Pretraining and distilling multi-billion-parameter encoders for natural language understanding systems. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 2893–2902.
- Asbjørn Følstad and Petter Bae Brandtzaeg. 2020. Users’ experiences with chatbots: findings from a questionnaire study. *Quality and User Experience*, 5(1):3.
- Liqiang Geng and Howard J Hamilton. 2006. Interestingness measures for data mining: A survey. *ACM Computing Surveys (CSUR)*, 38(3):9–es.
- Jessica He, David Piorkowski, Michael Muller, Kristina Brimijoin, Stephanie Houde, and Justin D Weisz. 2023. Understanding how task dimensions impact automation preferences with a conversational task assistant. In *ACM CHI Conference on Human Factors in Computing Systems*.
- Celeste Kidd and Benjamin Y Hayden. 2015. The psychology and neuroscience of curiosity. *Neuron*, 88(3):449–460.
- Julia Kiseleva, Kyle Williams, Ahmed Hassan Awadallah, Aidan C Crook, Imed Zitouni, and Tasos Anatasakos. 2016. Predicting user satisfaction with intelligent assistants. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, pages 45–54.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Tom Lotze, Stefan Klut, Mohammad Aliannejadi, and Evangelos Kanoulas. 2021. Ranking clarifying questions based on predicted user engagement. *arXiv preprint arXiv:2103.06192*.
- Shervin Malmasi, Anjie Fang, Besnik Fetahu, Sudipta Kar, and Oleg Rokhlenko. 2022. [SemEval-2022 task 11: Multilingual complex named entity recognition \(MultiCoNER\)](#). In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, pages 1412–1437, Seattle, United States. Association for Computational Linguistics.
- OpenAI. 2023a. Chatgpt. <https://openai.com/chatgpt>.
- OpenAI. 2023b. Gpt-4 technical report. *arXiv preprint arXiv: 2303.08774*. <https://openai.com/gpt4>.
- Dookun Park, Hao Yuan, Dongmin Kim, Yinglei Zhang, Matsoukas Spyros, Young-Bum Kim, Ruhi Sarikaya, Edward Guo, Yuan Ling, Kevin Quinn, et al. 2020. Large-scale hybrid approach for predicting user satisfaction with conversational agents. *arXiv preprint arXiv:2006.07113*.
- Reinhard Pekrun. 2019. The murky distinction between curiosity and interest: State of the art and future prospects. *Educational Psychology Review*, 31(4):905–914.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*.

- Ivan Sekulić, Mohammad Aliannejadi, and Fabio Crestani. 2021. User engagement prediction for clarification in search. In *Advances in Information Retrieval: 43rd European Conference on IR Research, ECIR 2021, Virtual Event, March 28–April 1, 2021, Proceedings, Part I 43*, pages 619–633. Springer.
- Paul J Silvia. 2005. What is interesting? exploring the appraisal structure of interest. *Emotion*, 5(1):89.
- Clemencia Siro, Mohammad Aliannejadi, and Maarten de Rijke. 2022. Understanding user satisfaction with task-oriented dialogue systems. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2018–2023.
- Carl Strathearn and Dimitra Gkatzia. 2022. Task2dial: A novel task and dataset for commonsense enhanced task-based dialogue grounded in documents. *arXiv preprint arXiv:2204.01061*.
- Silvan Tomkins. 1962. *Affect imagery consciousness: Volume I: The positive affects*. Springer publishing company.
- Frederico Vicente, Rafael Ferreira, David Semedo, and João Magalhães. 2023. The wizard of curiosities: Enriching dialogues with fun facts. In *Proceedings of the 24th Meeting of the Special Interest Group on Discourse and Dialogue*, pages 149–155.
- Alexandra Vtyurina and Adam Fourney. 2018. Exploring the role of conversational cues in guided task support with virtual assistants. In *Proceedings of the 2018 CHI conference on human factors in computing systems*, pages 1–7.
- Alexandra Vtyurina, Denis Savenkov, Eugene Agichtein, and Charles LA Clarke. 2017. Exploring conversational search with humans, assistants, and wizards. In *Proceedings of the 2017 chi conference extended abstracts on human factors in computing systems*, pages 2187–2193.
- Qiong Wu and Chunyan Miao. 2013. Curiosity: From psychology to computation. *ACM Computing Surveys (CSUR)*, 46(2):1–26.
- Zheng Zhang, Ryuichi Takanobu, Qi Zhu, MinLie Huang, and XiaoYan Zhu. 2020. Recent advances and challenges in task-oriented dialog systems. *Science China Technological Sciences*, 63(10):2011–2027.

Appendix

A Examples of Conversation Turns with Interesting Facts

Table 3 shows some examples of conversation turns containing interesting facts, and their labels annotated by human annotators.

Entity	Conversation turn	Human annotation
Ingredient: Sweet potato	Did you know that the vibrant colors of sweet potatoes can be extracted and used as a natural dye for fabrics? From facts.net.	Interesting
Recipe: Smoked Salmon Crepes	According to facts.net, crepe-making is an art form in Japan.	Boring, frustrating or distracting
Ingredient: Sausage	Here's something interesting from thefactsite.com: Sausages play a key role in Australian politics!	Interesting
Tool: Chopsticks	According to Wikipedia, the first chopsticks were used not as eating utensils but for cooking, stirring and serving.	Boring, frustrating or distracting
Recipe: Mooncake	According to monoandco.com, half-baked mooncakes must be taken out and cooled for 15 minutes before continuing baking.	Interesting
Ingredient: Candy	From eatthis.com, weirdly enough, cotton candy was invented in 1897 by a dentist.	Interesting
Recipe: Chocolate cake	According to facts.net, the secret to a moist cake is to use ingredients like buttermilk, sour cream, or yogurt, which add moisture.	Interesting
Ingredient: Pumpkin spice	Here's an interesting fact from facts.net: pumpkin spice does not actually contain pumpkins.	Interesting
Tool: Baking tin	From organicfacts.net, aluminium is considered the best material for a baking tin as it allows for a quick transfer of heat.	Boring, frustrating or distracting
Ingredient: Baking soda	According to tasty.co, baking soda must be replaced every month, otherwise a bit more than the recipe calls for can be added.	Interesting

Table 3: Examples of conversation turns annotated for interestingness by human annotators.