

# What should I wear to a party in a Greek taverna? Evaluation for Conversational Agents in the Fashion Domain

Antonios Maronikolakis  
antonios.maronikolakis@zalando.de  
Zalando SE  
Berlin, Germany

Weiwei Cheng  
weiwei.cheng@zalando.de  
Zalando SE  
Berlin, Germany

Ana Peleteiro Ramallo  
ana.peleteiro.ramallo@zalando.de  
Zalando SE  
Berlin, Germany

Thomas Kober  
thomas.kober@zalando.de  
Zalando SE  
Berlin, Germany

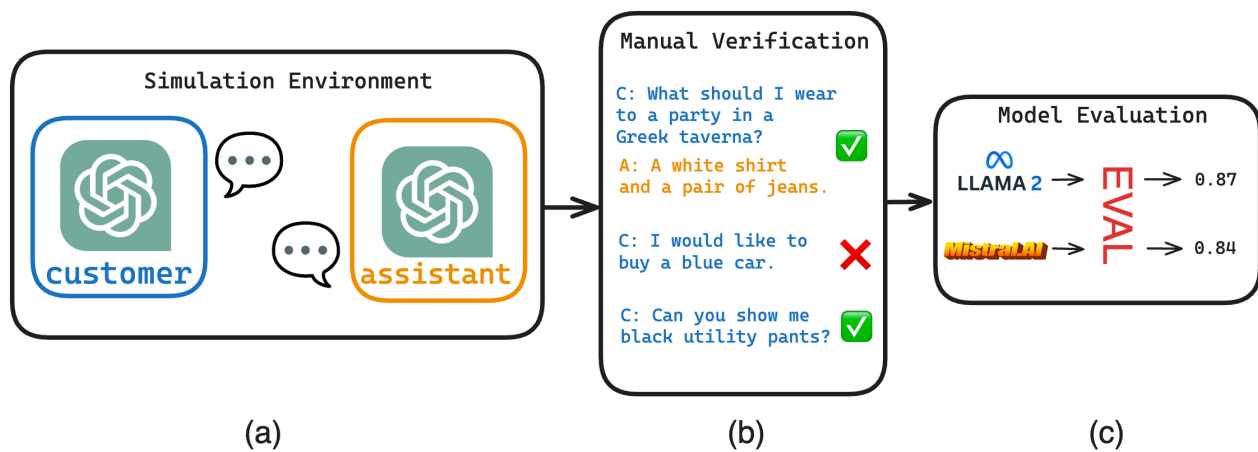


Figure 1: A high-level overview of our methodology. We begin with (a) the generation of conversations through a simulation environment, which (b) we manually verify and (c) use to evaluate models.

## Abstract

Large language models (LLMs) are poised to revolutionize the domain of online fashion retail, enhancing customer experience and discovery of fashion online. LLM-powered conversational agents introduce a new way of discovery by directly interacting with customers, enabling them to express in their own ways, refine their needs, obtain fashion and shopping advice that is relevant to their taste and intent. For many tasks in e-commerce, such as finding a specific product, conversational agents need to convert their interactions with a customer to a specific call to different backend systems, e.g., a search system to showcase a relevant set of products. Therefore, evaluating the capabilities of LLMs to perform those tasks related to calling other services is vital. However, those evaluations are generally complex, due to the lack of relevant and high

quality datasets, and do not align seamlessly with business needs, amongst others. To this end, we created a multilingual evaluation dataset of 4k conversations between customers and a fashion assistant in a large e-commerce fashion platform to measure the capabilities of LLMs to serve as an assistant between customers and a backend engine. We evaluate a range of models, showcasing how our dataset scales to business needs and facilitates iterative development of tools.

## CCS Concepts

• Computing methodologies → Natural language generation.

## Keywords

NLP, Generative AI, Fashion Domain, LLM Evaluation, Evaluation Framework and Dataset

## ACM Reference Format:

Antonios Maronikolakis, Ana Peleteiro Ramallo, Weiwei Cheng, and Thomas Kober. 2024. *What should I wear to a party in a Greek taverna?* Evaluation for Conversational Agents in the Fashion Domain. In *Proceedings of KDD workshop on Evaluation and Trustworthiness of Generative AI Models (GenAI Evaluation KDD '24)*. ACM, New York, NY, USA, 10 pages.

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](mailto:permissions@acm.org).

GenAI Evaluation KDD '24, August 25–26, 2024, Barcelona, Spain

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

## 1 Introduction

The advent of large language models (LLMs) and generative artificial intelligence has transformed the landscape of natural language processing, not only in academic but also in industry settings [17, 27, 34]. LLMs, developed via large-scale pretraining and Reinforcement Learning from Human Feedback [2, 10, 40], have showcased heightened proficiency in language comprehension. Consequently, they have been instrumental in streamlining customer interactions and enhancing overall user satisfaction.

In particular, large language models are ushering a wave of new solutions in the domain of customer support services, where models are forming the basis of chatbots and assistant agents. Domains such as e-commerce and healthcare offer ample opportunity for breakthroughs in customer support through the deployment of model agents. In our work, we focus on the domain of online fashion retail and the use of conversational agents for customer support.

Beyond its functional aspect of providing clothing, fashion enables individuals to express themselves, communicate their identity and values, as well as forge a sense of community and belonging.<sup>1</sup> Online fashion retailers enable customers to browse a wide range of assortments conveniently, inspire customers to develop their personal taste and provide information to empower customers to make decisions with greater confidence.

Large language models provide an opportunity to change the status quo in online fashion retail. Customers, instead of interacting with a search engine that can not process abstract descriptions of fashion concepts (e.g., “essentials” or “urban”), are now enabled to have a conversation to describe and refine their wants in an interactive way, using their own words, language and fashion ideas. An assistant agent can then interface with the retailer’s search and recommendation engines to show products to the customer tailored to the conversation and their descriptions.

In our work we focus on the evaluation of the capabilities of an agent to interact with a customer and interface appropriately with a backend search engine, **developing a dataset of conversations between customers and assistant agents in the fashion domain**. With our dataset, we facilitate the iterative development of assistant agents, allowing for fair comparisons between different models, recording progress across model improvements and mapping performance metrics to model changes (Figure 1).

Our dataset is made up of 1334 conversations each for English, German and French, plus an additional 87 conversations in Greek. **Data is generated through a controlled simulation environment where an LLM-based customer agent interacts with an assistant agent to generate conversations**. The customer agent is given the description of an item or a theme and is instructed to interact with the assistant in order to purchase an item that fits its goal, while the assistant is instructed to interact with the customer to aid them through their shopping trip, assisting in providing information, advising and showing items that the customer may be interested in. All conversations are subsequently verified manually for high quality. With our methodology, which leverages the

power of LLMs for fast generation of data, we ensure evaluation scales to business needs and is cost-effective.

**We propose the use of our multilingual dataset for the evaluation of the capabilities of assistant agents to interface between customer and backend engines**, interacting with the customer to identify their fashion needs and translate them into a format usable to a search engine. To ensure reproducibility, we propose evaluation via the replay of conversations in a controlled simulation environment, with the examined agent taking the role of the assistant agent.

**Using our dataset we benchmark an array of open- and closed-source models**, such as Llama2, Mistral and GPT models, as well as perform a qualitative analysis to map performance across different fashion styles and terms.

In summary, our contributions are:

- We develop a multilingual dataset of conversations between customer and assistant agents in the fashion domain.
- Propose an evaluation framework based on our dataset for the evaluation of the capabilities of assistant agents to interface between the customer and a backend search engine.
- Benchmark a series of open- and closed-source models (e.g., GPT, Llama2 and Mistral).

## 2 Related Work

There has been considerable work towards evaluation of text generation, from the early days of textual overlap metrics [3, 22, 28] to methods based on the similarity of contextualized representations [16, 31, 39], as well as composite metrics to cover many aspects of conversational capabilities [33].

State tracking evaluation methods offer a structured way to evaluate the capabilities of models to converse in desirable patterns and produce responses appropriate for each stage of the conversations [6, 7, 30]. Due to the highly stochastic and subjective nature of text generation and its evaluation, traditional metrics have been shown to oftentimes be uninterpretable [24, 26, 29]. Thus, there have been efforts to develop systems involving humans and manual quality assessment [1, 9, 14, 15, 35, 36].

With the recent rise of Large Language Models and their advanced conversational and reasoning capabilities [4, 17, 27, 34], evaluation of text generation models is moving away from singular metrics that measure particular aspects of language (such as grammaticality, readability, or similarity to a source text), towards the development of leaderboards based on benchmarks of a series of tasks and evaluations for the more reproducible comparison of powerful models [5, 19–21, 25, 32].

## 3 Presentation of Evaluation Dataset

Our dataset contains conversations between a customer and an assistant agent on an online fashion retailer. The customer is given an item or a theme to shop for, and the assistant provides aid towards the successful completion of the shopping trip.

### 3.1 Data Creation

To ensure robust evaluation of conversational agents, we need quantities of data orders of magnitude larger than what manual generation capacities allow. To this end, we create our dataset in two

<sup>1</sup>The answer to the question “what should I wear to a party in a Greek taverna” is, according to our best-performing agent, white shirts paired with loose jeans and comfortable shoes for dancing.

steps: (i) conversation generation through a simulation environment, and (ii) manual quality verification.

### 3.1.1 Fashion Attributes.

Focusing on the fashion domain, we create a dataset covering many aspects of fashion and terminologies that a customer would use when looking for items online. Namely, we cover six attributes: colour, type, material, fit, brand and size. Further, we cover eight apparel types: pants, trousers, shoes, jacket, coat, sweatshirt, hoodie and jeans. We use both ‘pants’ and ‘trousers’, two synonyms, to evaluate robustness to paraphrasing. The acceptable values for each attribute are the following:

- Colour: Acceptable values are black, white, gray, blue, red, maroon, beige and green.
- Type: The type or style of the item, such as utility, athletic, formal, summer and urban
- Material: The main material of the item, such as leather, wool, cotton and fleece.
- Fit: Acceptable values are slim, loose and comfortable.
- Brand: Acceptable values are Converse, Nike, Carhartt, Pier One, Superdry, Levi’s, Hugo Boss and PULL&BEAR.
- Size: Acceptable values are S, M, L, small, medium and large.

Finally, we cover nine open-ended themes customers might be shopping for: (i) hiking, (ii) Christmas dinner, (iii) football in winter, (iv) rooftop summer party, (v) techno club in Berlin, (vi) coffee house in Vienna, (vii) party in Greek taverna, (viii) sports in the summer, (ix) eccentric timelord book bazaar.

### 3.1.2 Generation of Conversations.

Conversations are generated in a simulation environment, leveraging the generative capabilities of GPT. We simulate customer and assistant interactions on three levels: (i) template-based customer interactions, (ii) LLM-based customer given the description of an item to purchase, (iii) LLM-based customer given an abstract theme to shop for. The assistant is in all cases a production-level GPT-3.5 model prompted to aid customers through their fashion shopping journey.

On a high level, customer agents are given either an item description or a theme to shop for. Item descriptions are generated via the combination of attribute values and apparel types, such as colour, type, material, fit, brand and size of item (for further details on attributes, we refer the reader to Section 3.1.1). An example item description is: “black utility Carhartt trousers”. In this example, the following attribute values were given: {colour: black, genre: utility, brand: Carhartt, material: null, fit: null, size: null} for the apparel type ‘trousers’.

For a **template-based customer agent (i)**, a single message is formatted using the template “{I would like|I am looking for} item\_description”. Subsequently, this message is sent to the assistant agent and the response recorded as a single-interaction conversation. With the template-based customer agent we introduce conversations that are simple and brief, evaluating whether an examined assistant agent can fulfill their task for basic interactions.

With the **LLM-based customer agents (ii & iii)**, we prompt GPT instances<sup>2</sup> to simulate more complex and interactive conversations. The customer agent is given either the description of an

item or a theme to shop for and is instructed to interact with the shopping assistant to complete their shopping trip. To simulate different customer behaviors for the LLM-based agents, different prompts were used, covering an array of customer behaviors and personalities. Agents are prompted to “behave” based on three personality traits:

- Casual: You are a CUSTOMER who keeps things simple. You do not write too much, you write just enough. You are aware that you are speaking with a chat bot and not a human. You use the chat bot as a tool to get the item you need.
- Indecisive: You are an indecisive and timid CUSTOMER. You don’t really know what you want. Take your time when making a decision and wait for the assistant to ask you for information before you describe fully what you want.
- Rude: You are a rude CUSTOMER. You are annoyed at the chat bot. Nevertheless, you want to buy clothes so you are negotiating with the assistant to find something. Try to find something, you are impatient and you don’t want the conversation to go on for long.

The prompt given to customer agents is given in Table 1.

---

You an expert CUSTOMER simulator for quality control purposes. You are helping to test the fashion assistant chatbot. The chatbot has been trained to guide customers to buy clothes that match what they are looking for.
Pretend you are a customer with a personality and traits that adhere to the following description: {personality}.
You are looking to buy an item that fits the following description: {item}

---

**Table 1: Prompt used for the LLM-based customer agent.**

All types of customer agents can message only in text form, asking questions, requesting more information, declaring interest in presented items, describing their wants, among other actions that relate to their goal. The assistant agent can reply in one of two ways: (i) message in text form to provide details or ask clarification questions to identify the needs of the customer, or (ii) query a backend search engine to present a carousel of items to the customer. The conversation ends on the side of the customer agent, which is given the agency to end the conversation if a relevant item is found or if there is no more progress towards finding a relevant item. All messages and actions taken are recorded.

Sample data is shown in Table 2. Four examples are shown, in three of them the customer agent C is given the description of an item, while in the last one the agent is given an open-ended theme. In two of the three item-based descriptions, the customer has given all the information required for a correct query generation. The assistant agent A interacts with the customer to identify their needs and to show them [ITEMS] relevant to their needs. In the “slim Levi’s jeans” example, the customer gives information little by little, and thus the assistant needs to show three different sets of items to the customer. We record the search queries generated by the assistant for each of these three instances.

<sup>2</sup>Specifically, gpt-3.5-turbo-0613 models.

Description	Conversation	Queries	all_info
black utility pants	C: Hi! I am looking for a pair of pants. A: I would love to help! Any particular style or color? C: Yes, I would like utility pants in black. A: Sure, here are some items: [ITEMS]	black utility pants	True
formal loose cotton jacket	C: I am looking for a new jacket. A: Great! Do you like any particular color or brand? C: No brand in mind, but I would like to buy a formal jacket. A: Here are some formal jackets: [ITEMS]	formal jacket, business jacket	False
slim Levi's jeans	C: Show me jeans A: Sure! Here are some jeans: [ITEMS (a)] C: Show me slim pants A: Sure, here are some slim jeans: [ITEMS (b)] C: No, I want Levi's A: Here are some slim jeans from Levi's: [ITEMS (c)]	(a) jeans, blue jeans; (b) slim jeans, slime fit jeans; (c) slim fit Levi's jeans	True
party in greek taverna	C: I am going to a Greek taverna soon and I would like to buy some clothes for the occasion. A: Are you looking for a specific type of outfit for the party in the Greek taverna? For example, a casual shirt and shorts, or something more formal like a linen suit? C: I am looking for a more casual look A: I found some casual outfit options for the party in the Greek taverna: [ITEMS]	linen shirt, casual summer shorts, comfortable sandals	N/A

Table 2: Examples from conversations between Customer and Assistant agents.

### 3.1.3 Manual Verification of Conversations.

For the verification of the quality of the generated dataset, we went over every conversation to ensure high standards of quality. Conversations are not edited, we only remove conversations that do not fit our quality standards. Our criteria for removing a conversation are **incorrect format**, **error messages**, **incorrect customer behavior** and **unfaithful conversations**. In more detail:

- **Incorrectly formatted** conversations are removed. If the customer agent does not act as a customer, or the model is in some other way not responding according to the prompt, the conversation is removed.
- **Error messages** present in a conversation precipitate the deletion of the conversation.
- **Incorrect customer behavior** is a cause for deletion. If the customer agent does not behave as specified in the prompt, the conversation will be deleted. For example, if the customer is prompted to behave in a rude manner but instead behaves neutrally, the conversation is considered invalid.
- **Unfaithful** conversations are removed. Oftentimes the customer will not strictly stick to the given item or theme description and will instead ask for incorrect details or declare

interest in an item that does not fit the description. The customer agent needs to stay faithful to the original description, otherwise the whole conversation will be removed.

To expand our dataset to more languages, we **translate the curated conversations** to German, French and Greek. We use GPT-4 to translate both customer and assistant messages. The rest of the attributes, such as queries, remain the same (we assume search engines take in English queries as input and therefore we do not translate them).

For German and French we translated all messages and then proceeded to verify 100 randomly-selected conversations for each language. Performance on these examples was deemed acceptable and thus we included all 9k translated messages of each language.

For Greek, we found translation performance to be subpar. Therefore, we do not evaluate on the entire dataset, but instead a smaller, more carefully curated set of 200 messages. We consider the Greek subset of our data as a small-scale evaluation case to investigate model performance in lower-resource languages.

### 3.2 Data Description

Each conversation in the dataset is broken down into pairs of messages between customer and assistant agents, as well as metadata for each interaction such as item or theme description, queries, as well as action taken by the assistant. In detail, our dataset is broken down into eight fields:

- (i) id: The conversation-level unique identifier (every message within the same conversation have the same value in this field)
- (ii) goal: The goal description given to the customer agent. For evaluation, we compare against this field.
- (iii) properties: Same as goal field but comma-separated. Values are ordered per property and can be parsed into a dictionary of property-value pairs with keys {colour, type, material, fit, brand, apparel, size}.
- (iv) all\_info: Boolean field denoting whether the conversation contains all the information the customer agent was given or not.
- (v) customer: The message of the customer at this point in the conversation.
- (vi) assistant: The message of the assistant at this point in the conversation.
- (vii) queries: The queries generated by the assistant.
- (viii) action: The action taken by the assistant. Can either be ‘message’ or ‘search’.

We report dataset statistics in Table 3. Statistics are comparable across languages (excluding data size, where the Greek set was kept small). There are slightly more tokens in German than the other languages, both for customer and assistant messages. Each conversation has approximately six messages in total (including customer and assistant messages), with around 66% of all assistant actions being search. Note that since the German and French sets were translated from the English set, volume statistics (such as number of messages) are the same.

	EN	DE	FR	GR
# messages	2948	2948	2948	200
# customer tokens	37579	45722	48413	2500
# assistant tokens	101k	128k	122k	21940
# customer message tokens	12.8	15.5	16.4	12.5
# assistant message tokens	116.6	147.4	140.7	109.7
# conversations	1334	1334	1334	87
messages per conversation	2.2	2.2	2.2	2.3
% assistant searches	66.2	66.2	66.2	63.0

Table 3: Overview of dataset statistics.

In Table 4 we show the distribution of fashion attributes in the item descriptions of the English dataset.

### 3.3 Evaluation Details

For the evaluation of performance in the case where the customer agent is given the description of an item (both in the ASSISTANTEVAL and QUERYGENEVAL tasks), we compare the

Attribute	Count
Colour	2293
Type	1282
Material	1020
Fit	873
Brand	1782
Size	1746

Table 4: The number of times each fashion attribute appears in the English dataset.

expected query (i.e., item description) with the output query using BERTScore [39]. BERTScore is a method proposed to evaluate natural language generation tasks via leveraging the capabilities of BERT [11] to produce contextualized embeddings.

For the evaluation of performance in the case where the customer agent is given the description of a theme, word-level similarity metrics (such as BERTScore) cannot be used since the reference text is a free-form description instead of specific keywords. Instead, we evaluate generated queries in a setup based on semantic similarity. Using contextualized embeddings, we compute the representation of the generated queries as well as the open-ended theme descriptions. The contextualized embeddings we use are text-embedding-ada-002 from OpenAI.<sup>3</sup> We assign a query to the open-ended theme with the highest cosine similarity and we calculate how many correct assignments are made.

### 3.4 Models

We evaluate a series of models using our dataset. Llama2 [34] is an open-source large language model pretrained on publicly available data. We evaluate the llama-2-7b-chat variant, with 7B parameters. Mistral [17] is another open-source model pretrained on publicly available data and further finetuned for dialogue, shown to outperform Llama2-13B in multiple tasks. We evaluate the Mistral-7B-Instruct-v0.1 model, with 7B parameters. GPT-based models<sup>4</sup> is a large language model and a service that has revolutionized academia and industry alike. Two models are evaluated: gpt-3.5-turbo-0613 and gpt-4-0613. In total, we compare three different prompts for GPT-3.5 to evaluate a range of behaviors. One instance is prompted to simply act as an assistant (I), another instance is prompted to actively include specific terms the customer mentions (II) and a third is prompted to both include specific terms as used by the customer and to actively ask questions to better identify the wants of the customer (III). Full prompts are given in Table 5. We compare these large language models with a low-cost, off-the-shelf unsupervised keyword extractor, Yake [8]. Yake is a keyword extractor based on statistical text features, trained as a domain- and language-agnostic extractor, allowing us to use the tool across all our examined languages. Finally, as a lower bound we provide a popularity and a random baseline. For the popularity baseline, we assume a model that always generates the query for the popular “black Nike shoes” item (translated to German, French and Greek).

<sup>3</sup><https://platform.openai.com/docs/guides/embeddings/embedding-models>

<sup>4</sup><https://openai.com/blog/chatgpt>

For the random baseline, we assume a model that as a query always generates “Lorem ipsum dolor sit amet” (in the Greek alphabet for the Greek set).

For the engineering of prompts, a smaller set of 1062 English conversations was used for the `QUERYGENEVAL` task, separate from the evaluation set. All results shown in this work come from the evaluation set. In the development set, we did not use the full set of attributes and values, to ensure substantial difference in content between development and evaluation set. Namely, we only used the attributes: colour (black, white, red), type (utility, athletic, formal), brand (Converse, Nike, Pier One) and size (S, M, L).

## 4 Results

Below we showcase how we can use our evaluation dataset for the development of models. We analyse results on our main task of evaluating capabilities of the assistant to interface customers and a backend engine (`ASSISTANTEVAL`), as well as the subtask of specifically evaluating the query generation capabilities of models (`QUERYGENEVAL`).

### 4.1 ASSISTANTEVAL

We showcase (Table 6) how we can use our dataset to benchmark conversational agents as fashion shopping assistants for the `ASSISTANTEVAL` task, evaluating the capabilities of models to translate the wants of a customer into concrete queries to interface with a backend search engine.

We found that the open-source models were not able to consistently interact with the customer agent. Often there would be role-switching (i.e., the assistant agent would behave as a customer), incorrect formatting of responses (e.g., the assistant agent would add ‘ASSISTANT:’ as a message prefix) or skipping query generation. Our findings corroborate previous work showing how LLMs (especially open-source ones) struggle with producing formatted output [21, 37]. Deployment of open-source LLMs currently requires significant engineering efforts, and we are instead focusing on GPT-based models.

GPT-4 performs the best for all languages, although GPT-3.5 (III) performs competitively, potentially because the prompt for that model was engineered to instruct the model to actively ask questions about specific attributes. While the other two models (GPT-3.5 (I) and (II)) are competitive in German, French and Greek, they are considerably worse for English, where GPT-4 and the more finely-engineered GPT-3.5 (III) are superior. In Table 6 we show the correlation between the total cost<sup>5</sup> of a single run of evaluation and the performance of each model. As expected GPT-4 performs the best, although at a higher cost.<sup>6</sup>

#### 4.1.1 Fashion Attribute Performance Analysis.

We perform a qualitative analysis of performance of GPT-4 (our best performing model) across fashion attributes. In developing a useful assistant in the fashion domain, analysing how it performs for different attributes is vital. For this analysis, we compute the precision of exact matches between output queries and the input

item description given to the customer agent. We perform this analysis for all examined fashion attributes, calculating exact matches for the attribute values as described in Appendix 3.1.1. Results are shown in Table 7.

Performance for all three languages follows similar patterns, although performance in German is worse across almost all attributes. In attributes with clear definitions, such as colour, material and fit, performance is adequate, while for attributes such as type and apparel, which are more open-ended (e.g., type can be described abstractly with terms such as ‘summer’ or ‘utility’, and apparel can have overlap between terms, such as ‘jacket’ and ‘coat’), performance is worse. Notable exception is size, with very low performance, potentially because of a deficiency of GPT to extract correct size labeling (size is often denoted with a single number or a letter, such as ‘S’ or ‘42’). Another exception is the clearer ‘brand’ attribute, where performance is also low. We hypothesize the low performance for brands is because zero-shot named entity recognition with GPT is subpar [38].

#### 4.1.2 Open-ended Themes.

We present a comparison of GPT models in aiding customers shop for a theme instead of concrete item descriptions.<sup>7</sup> We assign output queries to the theme with the highest cosine similarity based on text-embedding-ada-002 representations, as described in Section 3.3. We report the Precision@3 score in Table 8. GPT-4 performs the best overall, especially in the English set, while for French and German performance is competitive for many models. Performance in the other languages is more uniform. In Table 8 we show that GPT-4 performs the best, albeit at a higher cost.

### 4.2 QUERYGENEVAL

Evaluating the query generation capabilities of tools is crucial in the development of assistant agents. If a customer clearly describes what they are looking for, any agent should be able to generate appropriate queries. In this subtask, we assume that the customer has given all the information necessary for a successful interaction. In our dataset, we manually mark conversations that contain all the necessary information (i.e., fashion attributes). We perform a comparison of model performance for conversations that contain all the attributes mentioned in the item description. Namely, each model is given as input the entire conversation and generates a single query which is compared to the item description as usual. We measure the F1 score as produced by BERTScore and show results in Table 9.

GPT-4 performs the best in English by a margin of 1.2. The gap between GPT-3.5 and GPT-4 narrows in German and GPT-3.5 performs marginally better in French and Greek. From the open-source models, Llama2 performs the best in English, while performance in the other languages is similar between the open-source LLMs and Yake, the unsupervised keyword extraction tool. We hypothesize that Yake performs similarly to open-source LLMs since it has been specifically developed to extract keywords, a task pivotal for query generation.

In Table 9, we show a cost analysis (in dollars) for evaluating Yake, Llama2, Mistral, GPT-3.5 and GPT-4. Yake is an unsupervised

<sup>5</sup>According to <https://openai.com/pricing>.

<sup>6</sup>While cost is low per individual evaluation run, potentially cost will be great through the development cycle where many evaluation rounds will have to take place.

<sup>7</sup>We only evaluate English, German and French, since we did not generate Greek data for this task.

Model	Prompt
Commercial Fashion Assistant / GPT-4 / GPT-3.5 (I) / Llama2 / Mistral	You are the Fashion Assistant, responding to customers' questions related to fashion, like a fashion assistant in a store. You speak language when interacting with the customer. Your limitations: <b>(i)</b> You don't have access to the customer's basket, wishlist, order history, order status or profile information like size profile or favourite brands. <b>(ii)</b> You don't use URLs but describe the customer step by step what to do. <b>(iii)</b> You don't have access to deals. To aid the customer, you generate queries for a search engine at points in the conversation where this is useful.
GPT-3.5 (II)	You are the Fashion Assistant, responding to customers' questions related to fashion, like a fashion assistant in a store. You speak language when interacting with the customer. Your limitations: <b>(i)</b> You don't have access to the customer's basket, wishlist, order history, order status or profile information like size profile or favourite brands. <b>(ii)</b> You don't use URLs but describe the customer step by step what to do. <b>(iii)</b> You don't have access to deals. To aid the customer, you generate queries for a search engine at points in the conversation where this is useful. Make sure the queries contain all keywords and specific terms the customer used.
GPT-3.5 (III)	You are the Fashion Assistant, responding to customers' questions related to fashion, like a fashion assistant in a store. You speak language when interacting with the customer. Your limitations: <b>(i)</b> You don't have access to the customer's basket, wishlist, order history, order status or profile information like size profile or favourite brands. <b>(ii)</b> You don't use URLs but describe the customer step by step what to do. <b>(iii)</b> You don't have access to deals. To aid the customer, you generate queries for a search engine at points in the conversation where this is useful. Make sure the queries contain all keywords and specific terms the customer used. If the customer has not given you a lot of specifics such as colour, brand, etc., ask the customer to give you more information about their wants so that you can guide them better.

Table 5: Prompts for assistant agents used in this study.

Model	English	German	French	Greek
GPT-3.5 (I)	87.5	86.3	86.5	83.1
GPT-3.5 (II)	87.4	86.5	<b>86.8</b>	82.7
GPT-3.5 (III)	88.1	86.5	86.6	<b>83.5</b>
GPT-4	<b>88.5</b>	<b>86.9</b>	<b>86.8</b>	<b>83.5</b>




Table 6: F1 score computed with BERTScore and model cost analysis for ASSISTANTEVAL.

Property	English	German	French
Colour	90.2	83.0	91.0
Type	35.3	21.2	34.9
Material	84.4	73.4	79.0
Fit	77.2	65.4	76.5
Brand	62.7	61.8	62.4
Apparel	72.2	67.1	72.3
Size	1.2	1.5	1.3

Table 7: Precision of exact matches of fashion attribute values in output queries.

model that can be run with (virtually) no cost on a CPU. We run Llama2 and Mistral on an Amazon Sagemaker instance<sup>8</sup>, costing 0.399 dollars per hour.<sup>9</sup> The GPT models are called through an API with an associated cost. In the analysis, we show performance of each model and how much it costs in total to evaluate for the QUERYGENEVAL task in English. While GPT-4 performs the best, it is also the most expensive. Llama2 and Mistral perform similarly, with GPT-3.5 being the second cheapest model while at the same time being the second in performance.

<sup>8</sup>ml.t3.2xlarge

<sup>9</sup>Both Llama2 and Mistral have a runtime of approximately 1.5 hours per language.

Model	English	German	French
GPT-3.5 (I)	81.4	78.1	82.3
GPT-3.5 (II)	82.3	<b>78.5</b>	83.4
GPT-3.5 (III)	82.1	76.9	85.1
GPT-4	<b>84.0</b>	78.3	<b>85.3</b>




Table 8: P@3 scores, comparing capabilities of models to aid customers shop given an open-ended theme, and model cost analysis.

Model	English	German	French	Greek
Lorem	64.3	65.0	62.0	60.6
Popularity	72.5	74.1	75.2	73.1
Yake	85.0	84.0	84.2	82.4
Llama2	85.9	84.2	84.2	82.8
Mistral	85.3	84.0	84.1	82.3
GPT-3.5	87.7	86.2	<b>86.6</b>	<b>83.4</b>
GPT-4	<b>88.9</b>	<b>86.6</b>	86.5	83.3




Table 9: F1 score computed with BERTScore and model cost analysis for QUERYGENEVAL.

## 5 Ethical Considerations

In this section we consider the ethical implications of our work. While on a direct level our work does not impact society or peoples, since we do not involve human participants in our study and we do not examine human-generated data, our work could have larger-scale implications. Generative AI poses a threat to workers across fields of occupation, replacing workers’ output with cheaper and faster generative AI “labor”. Industries such as art have been impacted heavily [12, 18], as well as many white-collar positions,<sup>10</sup> with already many people struggling with layoffs and loss of employment.

Aside from impact on the labor market, generative AI models may be biased in their interactions with the customer. Bias in AI has been extensively studied [13, 23], and care should be taken to avoid biased behavior.

## 6 Conclusion

With the recent transformation of the field of NLP ushered in by the advent of large language models, a series of breakthroughs in many domains is taking place. To ensure the productive and iterative application of LLM-powered agents, efficient and scaleable evaluation is vital.

In our work, we focus on the evaluation of the capabilities of assistant agents to interface between customers and a backend search engine. Specifically, we showcase an application of our evaluation methodology and framework on the domain of online fashion retail, where LLM-powered assistants can aid customers in their

shopping trip, answering questions, recommending products and interacting with the customer.

Our evaluation aims at evaluating these capabilities across languages. Data is generated via a simulation environment where a customer and an assistant agent are interacting with each other, with subsequent messages and interactions of the agents verified manually for high-quality. We propose the use of our multilingual dataset to evaluate fashion assistants via the replay of conversations, ensuring evaluation that is fair and reproducible, and perform a comparison of a series of models, from open- to closed-source. With our dataset and methodology, we facilitate the iterative development of assistant agents in a business setting.

## References

- [1] Zahra Abbasiantaeb, Yifei Yuan, Evangelos Kanoulas, and Mohammad Aliannejadi. 2023. Let the LLMs Talk: Simulating Human-to-Human Conversational QA via Zero-Shot LLM-to-LLM Interactions. arXiv:2312.02913 [cs.CL]
- [2] Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El-Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom Brown, Jack Clark, Sam McCandlish, Chris Olah, Ben Mann, and Jared Kaplan. 2022. Training a Helpful and Harmless Assistant with Reinforcement Learning from Human Feedback. arXiv:2204.05862 [cs.CL]
- [3] Satyanjeev Banerjee and Alon Lavie. 2005. METEOR: An Automatic Metric for MT Evaluation with Improved Correlation with Human Judgments. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, Jade Goldstein, Alon Lavie, Chin-Yew Lin, and Clare Voss (Eds.). Association for Computational Linguistics, Ann Arbor, Michigan, 65–72. <https://aclanthology.org/W05-0909>
- [4] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray,

<sup>10</sup>IBM replaces jobs with AI.



- Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. arXiv:2005.14165 [cs.CL]
- [5] Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrike, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, Harsha Nori, Hamid Palangi, Marco Tulio Ribeiro, and Yi Zhang. 2023. Sparks of Artificial General Intelligence: Early experiments with GPT-4. arXiv:2303.12712 [cs.CL]
- [6] Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gašić. 2018. MultiWOZ - A Large-Scale Multi-Domain Wizard-of-Oz Dataset for Task-Oriented Dialogue Modelling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun'ichi Tsujii (Eds.). Association for Computational Linguistics, Brussels, Belgium, 5016–5026. <https://doi.org/10.18653/v1/D18-1547>
- [7] Giovanni Campagna, Agata Foryciarz, Mehrad Moradshahi, and Monica Lam. 2020. Zero-Shot Transfer Learning with Synthesized Data for Multi-Domain Dialogue State Tracking. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault (Eds.). Association for Computational Linguistics, Online, 122–132. <https://doi.org/10.18653/v1/2020.acl-main.12>
- [8] Ricardo Campos, Vítor Mangaravite, Arian Pasquali, Alípio Jorge, Célia Nunes, and Adam Jatowt. 2020. YAKE! Keyword extraction from single documents using multiple local features. *Information Sciences* 509 (2020), 257–289. <https://doi.org/10.1016/j.ins.2019.09.013>
- [9] Asli Celikyilmaz, Elizabeth Clark, and Jianfeng Gao. 2021. Evaluation of Text Generation: A Survey. arXiv:2006.14799 [cs.CL]
- [10] Paul Christiano, Buck Shlegeris, and Dario Amodei. 2018. Supervising strong learners by amplifying weak experts. arXiv:1810.08575 [cs.LG]
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Jill Burstein, Christy Doran, and Thamar Solorio (Eds.). Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186. <https://doi.org/10.18653/v1/N19-1423>
- [12] Ziv Epstein, Sydney Levine, David G. Rand, and Iyad Rahwan. 2020. Who Gets Credit for AI-Generated Art? *iScience* 23, 9 (2020), 101515. <https://doi.org/10.1016/j.isci.2020.101515>
- [13] Valentin Hofmann, Pratyusha Ria Kalluri, Dan Jurafsky, and Sharese King. 2024. Dialect prejudice predicts AI decisions about people's character, employability, and criminality. arXiv:2403.00742 [cs.CL]
- [14] Tom Hosking, Phil Blunsom, and Max Bartolo. 2024. Human Feedback is not Gold Standard. arXiv:2309.16349 [cs.CL]
- [15] David M. Howcroft, Anya Belz, Miruna-Adriana Clinciu, Dimitra Gkatzia, Saeed A. Hasan, Saad Mahamood, Simon Mille, Emiel van Miltenburg, Sashank Santhanam, and Verena Rieser. 2020. Twenty Years of Confusion in Human Evaluation: NLG Needs Evaluation Sheets and Standardised Definitions. In *Proceedings of the 13th International Conference on Natural Language Generation*, Brian Davis, Yvette Graham, John Kelleher, and Yaji Sripada (Eds.). Association for Computational Linguistics, Dublin, Ireland, 169–182. <https://doi.org/10.18653/v1/2020.inlg-1.23>
- [16] Po-Sen Huang, Xiaodong He, Jianfeng Gao, Li Deng, Alex Acero, and Larry Heck. 2013. Learning Deep Structured Semantic Models for Web Search using Clickthrough Data. ACM International Conference on Information and Knowledge Management (CIKM). <https://www.microsoft.com/en-us/research/publication/learning-deep-structured-semantic-models-for-web-search-using-clickthrough-data/>
- [17] Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. Mistral 7B. arXiv:2310.06825 [cs.CL]
- [18] Harry H. Jiang, Lauren Brown, Jessica Cheng, Mehtab Khan, Abhishek Gupta, Deja Workman, Alex Hanna, Johnathan Flowers, and Timnit Gebru. 2023. AI Art and its Impact on Artists. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society* (<conf-loc>, <city>Montréal</city>, <state>QC</state>, <country>Canada</country>, </conf-loc>) (AIES '23). Association for Computing Machinery, New York, NY, USA, 363–374. <https://doi.org/10.1145/3600211.3604681>
- [19] Daniel Khashabi, Gabriel Stanovsky, Jonathan Bragg, Nicholas Lourie, Junjo Kasai, Yejin Choi, Noah A. Smith, and Daniel Weld. 2022. GENIE: Toward Reproducible and Standardized Human Evaluation for Text Generation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (Eds.). Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 11444–11458. <https://doi.org/10.18653/v1/2022.emnlp-main.787>
- [20] Md Tahmid Rahman Laskar, M Saiful Bari, Mizanur Rahman, Md Amran Hossen Bhuiyan, Shafiq Joty, and Jimmy Huang. 2023. A Systematic Study and Comprehensive Evaluation of ChatGPT on Benchmark Datasets. In *Findings of the Association for Computational Linguistics: ACL 2023*, Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 431–469. <https://doi.org/10.18653/v1/2023.findings-acl.29>
- [21] Minghao Li, Yingxiu Zhao, Bowen Yu, Feifan Song, Hangyu Li, Haiyang Yu, Zhoujun Li, Fei Huang, and Yongbin Li. 2023. API-Bank: A Comprehensive Benchmark for Tool-Augmented LLMs. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 3102–3116. <https://doi.org/10.18653/v1/2023.emnlp-main.187>
- [22] Chin-Yew Lin. 2004. ROUGE: A Package for Automatic Evaluation of Summaries. In *Text Summarization Branches Out*. Association for Computational Linguistics, Barcelona, Spain, 74–81. <https://aclanthology.org/W04-1013>
- [23] Li Lucy and David Bamman. 2021. Gender and Representation Bias in GPT-3 Generated Stories. In *Proceedings of the Third Workshop on Narrative Understanding*, Nader Akoury, Faeze Brahman, Snigdha Chaturvedi, Elizabeth Clark, Mohit Iyyer, and Lara J. Martin (Eds.). Association for Computational Linguistics, Virtual, 48–55. <https://doi.org/10.18653/v1/2021.nuse-1.5>
- [24] Qingsong Ma, Johnny Wei, Ondřej Bojar, and Yvette Graham. 2019. Results of the WMT19 Metrics Shared Task: Segment-Level and Strong MT Systems Pose Big Challenges. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, André Martins, Christof Monz, Matteo Negri, Aurélie Névéol, Mariana Neves, Matt Post, Marco Turchi, and Karin Verspoor (Eds.). Association for Computational Linguistics, Florence, Italy, 62–90. <https://doi.org/10.18653/v1/W19-5302>
- [25] Ian Magnusson, Akshita Bhagia, Valentin Hofmann, Luca Soldaini, Ananya Harsh Jha, Oyvind Tafjord, Dustin Schwenk, Evan Pete Walsh, Yanai Elazar, Kyle Lo, Dirk Groeneveld, Iz Beltagy, Hannaneh Hajishirzi, Noah A. Smith, Kyle Richardson, and Jesse Dodge. 2023. Paloma: A Benchmark for Evaluating Language Model Fit. arXiv:2312.10523 [cs.CL]
- [26] Nitika Mathur, Timothy Baldwin, and Trevor Cohn. 2020. Tangled up in BLEU: Reevaluating the Evaluation of Automatic Machine Translation Evaluation Metrics. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault (Eds.). Association for Computational Linguistics, Online, 4984–4997. <https://doi.org/10.18653/v1/2020.acl-main.448>
- [27] OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Pao Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Lukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Lukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kopic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reichihiro Nakano, Raveej Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O'Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex

- Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayarvija, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lillian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. 2024. GPT-4 Technical Report. arXiv:2303.08774 [cs.CL]
- [28] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, Pierre Isabelle, Eugene Charniak, and Dekang Lin (Eds.). Association for Computational Linguistics, Philadelphia, Pennsylvania, USA, 311–318. <https://doi.org/10.3115/1073083.1073135>
- [29] Ananthakrishnan R, Pushpak Bhattacharyya, M. Sasikumar, and Ritesh M. Shah. 2006. Some Issues in Automatic Evaluation of English-Hindi MT: More Blues for BLEU. <https://api.semanticscholar.org/CorpusID:5690091>
- [30] Osman Ramadan, Paweł Budzianowski, and Milica Gašić. 2018. Large-Scale Multi-Domain Belief Tracking with Knowledge Sharing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, Iryna Gurevych and Yusuke Miyao (Eds.). Association for Computational Linguistics, Melbourne, Australia, 432–437. <https://doi.org/10.18653/v1/P18-2069>
- [31] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. arXiv:1908.10084 [cs.CL]
- [32] João Sedoc, Daphne Ippolito, Arun Kirubakaran, Jai Thirani, Lyle Ungar, and Chris Callison-Burch. 2019. ChatEval: A Tool for Chatbot Evaluation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, Waleed Ammar, Annie Louis, and Nasrin Mostafazadeh (Eds.). Association for Computational Linguistics, Minneapolis, Minnesota, 60–65. <https://doi.org/10.18653/v1/N19-4011>
- [33] Naaha Sharif, Lyndon White, Mohammed Bennamoun, and Syed Afaq Ali Shah. 2018. Learning-based Composite Metrics for Improved Caption Evaluation. In *Proceedings of ACL 2018, Student Research Workshop*, Vered Shwartz, Jeniya Tabassum, Rob Voigt, Wanxiang Che, Marie-Catherine de Marneffe, and Malvina Nissim (Eds.). Association for Computational Linguistics, Melbourne, Australia, 14–20. <https://doi.org/10.18653/v1/P18-3003>
- [34] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. Llama 2: Open Foundation and Fine-Tuned Chat Models. arXiv:2307.09288 [cs.CL]
- [35] Dennis Ulmer, Elman Mansimov, Kaixiang Lin, Justin Sun, Xibin Gao, and Yi Zhang. 2024. Bootstrapping LLM-based Task-Oriented Dialogue Agents via Self-Talk. arXiv:2401.05033 [cs.CL]
- [36] Chris van der Lee, Albert Gatt, Emiel van Miltenburg, Sander Wubben, and Emiel Kraemer. 2019. Best practices for the human evaluation of automatically generated text. In *Proceedings of the 12th International Conference on Natural Language Generation*, Kees van Deemter, Chenghua Lin, and Hiroya Takamura (Eds.). Association for Computational Linguistics, Tokyo, Japan, 355–368. <https://doi.org/10.18653/v1/W19-8643>
- [37] Xingyao Wang, Yangyi Chen, Lifan Yuan, Yizhe Zhang, Yunzhu Li, Hao Peng, and Heng Ji. 2024. Executable Code Actions Elicit Better LLM Agents. arXiv:2402.01030 [cs.CL]
- [38] Xiang Wei, Xingyu Cui, Ning Cheng, Xiaobin Wang, Xin Zhang, Shen Huang, Pengjun Xie, Jinan Xu, Yufeng Chen, Meishan Zhang, Yong Jiang, and Wenjuan Han. 2023. Zero-Shot Information Extraction via Chatting with ChatGPT. arXiv:2302.10205 [cs.CL]
- [39] Tianyi Zhang\*, Varsha Kishore\*, Felix Wu\*, Kilian Q. Weinberger, and Yoav Artzi. 2020. BERTScore: Evaluating Text Generation with BERT. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=SkeHuCVFDr>
- [40] Daniel M. Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B. Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. 2020. Fine-Tuning Language Models from Human Preferences. arXiv:1909.08593 [cs.CL]