

Using Large Language Models to Generate, Validate, and Apply User Intent Taxonomies

Chirag Shah^{1†‡}, Ryen W. White^{2†}, Reid Andersen², Georg Buscher², Scott Counts², Sarkar Snigdha Sarathi Das^{3‡}, Ali Montazer^{4‡}, Sathish Manivannan², Jennifer Neville², Xiaochuan Ni², Nagu Rangan², Tara Safavi², Siddharth Suri², Mengting Wan², Leijie Wang^{1‡}, Longqi Yang²

¹University of Washington, ²Microsoft, ³Pennsylvania State University, ⁴University of Massachusetts Amherst

†Corresponding authors: chirags@uw.edu, ryenw@microsoft.com

‡Work done while working at Microsoft, USA

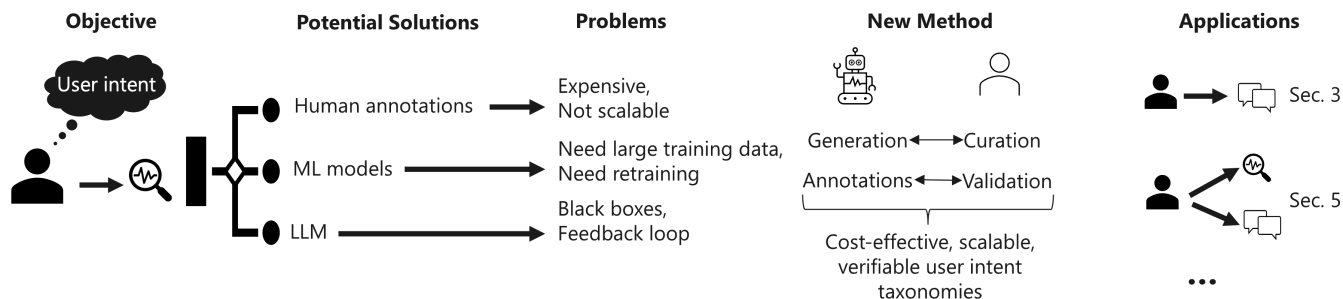


Figure 1: Detecting user intent in search situations is very important, but the potential solutions have different sets of problems. We propose a new method involving a unique collaboration between LLMs and humans to leverage best of both worlds.

ABSTRACT

Log data can reveal valuable information about how users interact with Web search services, what they want, and how satisfied they are. However, analyzing user intents in log data is not easy, especially for emerging forms of Web search such as AI-driven chat. To understand user intents from log data, we need a way to label them with meaningful categories that capture their diversity and dynamics. Existing methods rely on manual or machine-learned labeling, which are either expensive or inflexible for large and dynamic datasets. We propose a novel solution using large language models (LLMs), which can generate rich and relevant concepts, descriptions, and examples for user intents. However, using LLMs to generate a user intent taxonomy and apply it for log analysis can be problematic for two main reasons: (1) such a taxonomy is not externally validated; and (2) there may be an undesirable feedback loop. To address this, we propose a new methodology with human experts and assessors to verify the quality of the LLM-generated taxonomy. We also present an end-to-end pipeline that uses an LLM with human-in-the-loop to produce, refine, and apply labels for user intent analysis in log data. We demonstrate its effectiveness by uncovering new insights into user intents from search and chat logs from the Microsoft Bing commercial search engine. The proposed work’s novelty stems from the method for generating purpose-driven user intent taxonomies with strong validation. This method not only helps remove methodological and practical bottlenecks from intent-focused research, but also provides a new framework

for generating, validating, and applying other kinds of taxonomies in a scalable and adaptable way with minimal human effort.

This article is distributed under Creative Commons Attribution-NonCommercial- NoDerivatives License 4.0 (CC BY-NC-ND).

KEYWORDS

User intents, Large language models, Taxonomies, Log data

1 INTRODUCTION

Understanding the purpose or the task behind a user’s request in an information access context is highly desired for a search or a recommender system to be able to provide the most relevant and meaningful results [57]. However, extracting user intents from log data is extremely difficult due to two main reasons: (1) fluidity in what user intents are or can be; and (2) how these intents can be identified using log data that may not include sufficient context. Additionally, in the case of emerging modalities such as AI-driven chat, users’ understanding, usage, and behaviors are rapidly evolving that call for on-demand, task-focused labels and taxonomies. We need new methods to identify, extract, and apply user intents in IR systems, especially those with emerging modalities.

Traditional qualitative methods such as coding and thematic analysis are time-consuming and require human expertise [8]. Conversely, existing quantitative methods may not capture the nuances and diversity of user intents and experiences [36]. Large language models (LLMs) have become quite capable of generating coherent

texts from various inputs [10]. But can they be useful in reliable and verifiable ways to conduct such intent-focused research?

There are several attempts in the recent years to use LLMs in various applications ranging from ranking and recommendations to content generation and evaluation (e.g., [6, 20, 38]). However, many of these works often lack rigor and reliability since LLMs are used as black boxes without a meaningful understanding of their inner workings or there are feedback loops with weak or non-existent validation for the method. Simply focusing on the promising results without sufficient support of scientific rigor in the methodology can lead to misleading and even dangerous outcomes. We believe that while LLMs have shown great promise for aiding us in various informational tasks, they must be used with responsibility and sufficient validation. This leads us to the following research questions (RQs): (1) Can we use LLMs to reliably generate taxonomies for analyzing user intents in log data? (2) Can an LLM correctly apply a user intent taxonomy to annotate logs? (3) If and when can an LLM perform better than human annotators? Or are there uses of LLMs that go beyond just reducing efforts and increasing efficiency?

To address these RQs, we investigated if/how LLMs can help in creating an end-to-end solution for developing user intent taxonomies from AI chat logs. We use GPT-4¹, for most of our experiments since it is a leading LLM and we wanted to test the potential of the most advanced models. In the process, we devised a new methodology for employing any LLM as a collaborator in an iterative qualitative analysis process that leverages its ability to generate summaries, questions, and categories from chat transcripts.

We already know that LLMs can generate taxonomies, and generate effective annotations [24]. Using human-in-the-loop with computational models is also not new. What brings **novelty** in this paper is the methodology that harnesses the potential of an LLM to infer a structure from data and the curation by human assessors in an integrative way. This significantly reduces human effort and provides validity of machine-generated outputs. We demonstrate the value of this by quickly and reliably running the end-to-end pipeline with heterogeneous data from search and chat logs and generating insights. Therein lies the **significance** of this work – it can remove bottlenecks for research that aims to identify and apply user intents. Beyond that, this methodology can also be useful for generating and applying other kinds of taxonomies in IR and beyond. The **technical soundness** of our work comes from rigorous validation of the proposed method and a set of experiments involving multiple datasets, human assessors, and three different LLMs. Supporting material is available online² improving replicability and enabling a wider application of our methods.

2 RELATED WORK

We are focusing on generating user intent taxonomies from logs, but we are not the first ones to study user intents, build taxonomies, or use such taxonomies to generate insights from logs. Therefore, before we dive into our novel contributions, viz., using LLMs for generating and using user intents taxonomies with scientific rigor, it is important to briefly review some of the related work.

2.1 Taxonomy Generation, Validation, and Use

Taxonomies are hierarchical classifications of concepts, terms, or entities. They can facilitate information seeking, retrieval, or behavior by providing structure, organization, and navigation for users and systems [12, 13, 54]. However, generating and validating taxonomies is a challenging task that requires balancing multiple criteria such as coverage, coherence, consistency, granularity, usability, and adaptability, e.g., [29, 35, 41, 46]. Moreover, different domains and contexts may have different requirements and preferences for taxonomy design and evaluation [32]. Taxonomies can be generated manually through an iterative process and the research community has developed tools to generate taxonomies automatically from document collections using methods such as clustering [55]. We are the first to leverage the power of LLMs to automatically generate taxonomies in a search context, focused specifically on conversational search; we also validate that LLM-based methodology with human assessors. More importantly, we provide a method for other researchers to do the same for their specific needs.

2.2 Use of LLMs in Research

The emergence of LLMs has unlocked many opportunities for rapid research advances. LLMs have been used to enable scientific discovery [26], with remarkable progress in areas such as medicine [45] and finance [4]. Early language models, such as BERT, and various natural language processing methods have been used to auto-code qualitative data [1, 23], although not at a near-human level. LLM-driven AI is capable of qualitative analysis and can generate nuanced results comparable to human researchers [11]. Pham et al. [40] used LLMs to generate topic models from large corpora.

This progress has stirred discussions about the role of LLMs in quantitative and qualitative research. Some view these as tools to enrich human understanding, while others perceive them as threats to core scientific values [5]. Watkins [51] proposes considerations that can be applied by researchers to guide the use of LLMs in their workflows. Bano et al. [5] predict that we will see LLMs redefining the interplay between humans and AI in qualitative research. LLMs can derive insights from unstructured data and collect and analyze massive amounts of data from diverse sources. However, in their current form they lack the intuitions, nuances, and sensitivity to culture and context that humans have. A shift toward more human-AI teaming for qualitative analysis is more likely in the short term given current LLM capabilities than a total LLM-centric transformation in how qualitative research is performed [19].

In information retrieval (IR), LLMs have been shown to be effective in supporting humans in judging document relevance, an activity that is central to search engine design and evaluation [20]. They have also been recently used for synthetic dataset generation to support IR research [28] and richer user modeling to support IR experimentation [33]. Other applications of LLMs in IR have also been discussed at length by others in the IR community [3]. In this paper, we will show that working directly with humans, LLMs have the potential to support two other critical activities in search engine research and development: intent understanding and generating intent taxonomies.

¹<https://openai.com/gpt-4/>

²<https://anonymous.4open.science/r/LLM-for-taxonomy/>

2.3 Log Analysis and Insights

Search log analysis has been used extensively to gain insights about search interactions, including queries, search engine result page (SERP) clicks, and post-SERP interactions [18, 47, 52]. Analyzing log data has historically been a highly interactive process: researchers first write scripts to extract data, then analyze that data manually using data science tools and methods, and (optionally) human annotators label data to better understand patterns and trends and generate training data for machine learned (ML) models, e.g., [2].

More recently, there has been an increasing focus on user engagement with conversational search systems [22]. Researchers and practitioners have sought to understand user intents and behaviors in the context of chat-based systems [42, 50]. Chat is becoming an increasingly popular modality for information seeking, especially in domains where users have complex or exploratory queries, need guidance or clarification, or prefer a conversational style of interaction. Datasets of chat logs have also been created and released to the community to promote research in this emerging area [43].

ML models can help support researchers in performing data analysis [25]. Recently, from applications in other domains, we have seen that LLMs may have the potential to play a supportive role in the analysis of text data, providing insights and annotations that expedite experiments and reduce human effort [48]. In this paper, for the first time, we introduce methods for using LLMs collaborating with researchers to derive insights and taxonomies from retrospective log data.

2.4 Understanding User Intents in Search

Intent has been a well explored area in IR and various approaches have been proposed for intent understanding and intent representation [31, 56]. Taxonomies of search intent can help systems better understand user intent. Several search intent taxonomies have been proposed, e.g., [9, 27, 30, 44]. These have been generated iteratively via manual inspection of log data. Intent taxonomies developed from search logs include: navigation, information, transactions, browsing, and resource finding. We expect chat interactions to exhibit new intents (e.g., creation) compared to traditional search. Automatic generation of query taxonomies has also been attempted [15, 16] using query clustering to derive a taxonomy from existing query data and categorization to assign new queries to the taxonomy. Taxonomies have also been used to represent intent in question-answering [7, 14]. Xie [53] derived an empirically based classification of search intents that motivate different search behaviors. Mitsui et al. [37] developed a set of information-seeking intentions based on that classification and studied differences in intentions as a function of the search task. Building on this previous work, we leverage LLMs to generate intent taxonomies with humans-in-the-loop and evaluate the performance of LLMs in assigning search activities to intent categories.

3 METHODOLOGY FOR GENERATING AND VALIDATING A USER INTENT TAXONOMY

In this section, we describe a new methodology that we developed and tested for employing LLMs to generate a user intent taxonomy that can be used to generate insights and construct hypotheses from log data, focused for this analysis on AI chat logs.

Let us begin with a problem scenario. We have access to log data from user interactions with an AI-driven chat system (Bing Chat). This data primarily includes user requests and AI responses in natural language. We can analyze this data in a number of ways, answering questions about what the users are doing (topics, domains). But if we want to understand their intents, we need a set of labels or a taxonomy of intents. In practice, one looks for relevant literature for an existing taxonomy, but if an appropriate taxonomy does not exist, one needs to create it. This can be done by taking an existing taxonomy and modifying it to fit the data or the task (top-down approach), or by building it fresh using the available data (bottom-up approach). Following this, one needs to validate the taxonomy to ensure it meets several criteria for a good taxonomy. Finally, the newly generated taxonomy could be applied to a specific task to generate the desired insights from the data.

Given that we were interested in analyzing AI chat logs, a seemingly newer type of modality, we found it to be desirable to adopt a bottom-up approach. In this approach, one typically analyzes available data to generate codes or labels, leading to a classification scheme or a taxonomy. As detailed in [9, 27, 30, 44], this process could involve one or more researchers and a considerable effort. We wanted to use an LLM to build such a taxonomy using relevant data and instructions. However, given that we do not have enough knowledge about how an LLM is constructed and how it creates or links various concepts from given data, we needed a way to validate LLM's generation and fine-tune it as needed. For that, we used two researchers with many years of experience in doing qualitative analysis and building taxonomies. These researchers guided the taxonomy generation process with the LLM and two human assessors. Once the taxonomy was built using training data and validated, it was applied to annotating test data. In short, our methodology uses LLMs as the backbone of taxonomy generation and application with humans in the loop for curation and validation.

The outline of our methodology is shown in Figure 2. Here, we used GPT-4 as the LLM for generating the taxonomy (Phase 1), engaged human assessors to validate that taxonomy (Phase 2), and then employed both GPT-4 and human assessors to apply it (Phase 3). Through the phases of validation and application, we evaluated not only the generated taxonomy (RQ1), but also GPT-4's ability and potential to perform such research-based tasks reasonably and reliably (RQ2). Following are the details.

3.1 Data

We took a random sample of 1,149 conversations from May-June 2023, available through Bing Chat. Each conversation contained one or more turns of user request and AI response. We ensured that these conversations were in English, however, some of the user requests were interspersed with non-English words. We do not believe this impacted any text processing by GPT-4 for our purposes. We used 1,000 conversations for training (building a user intent taxonomy) and set aside the rest for validation and testing.

3.2 Evaluating the Taxonomy

We first start by describing how a taxonomy should be evaluated. This will inform how we generate the taxonomy using GPT-4 (how we provide prompts), how we validate and revise it, as well as how

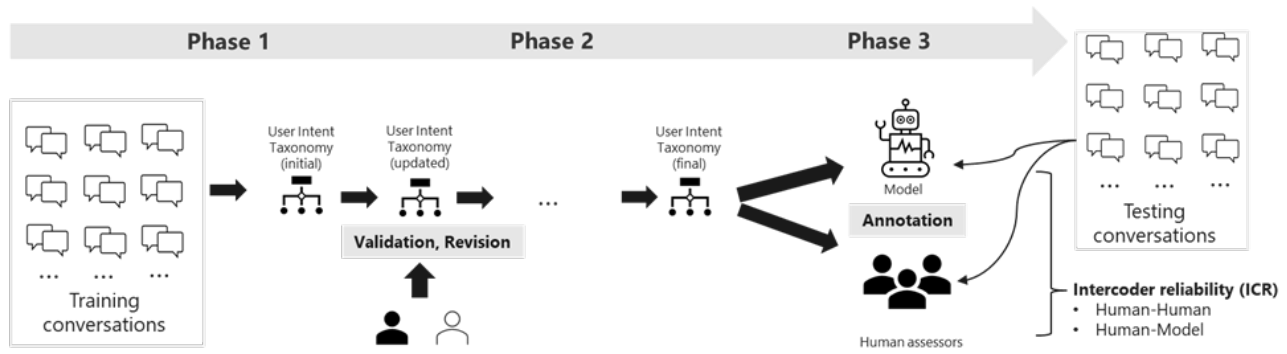


Figure 2: Three phases of user intent taxonomy generation, validation, and application.

we measure the effectiveness of the taxonomy for developing insights from logs. Using the relevant literature (some of which is summarized in the previous section) concerning taxonomy generation and validation, we consolidated the following criteria, taken from Raad and Cruz [41], with appropriate modifications.

- **Comprehensiveness:** All the data should be reliably classified using this taxonomy.
- **Consistency:** The taxonomy does not include or allow for any contradictions.
- **Clarity:** The taxonomy should communicate the intended meaning of the defined terms. Definitions should be objective and independent of the context.
- **Accuracy:** The definitions, descriptions of classes, properties, and individuals in a taxonomy should be correct.
- **Conciseness:** The taxonomy should not include any irrelevant elements with regards to the user intents in AI chat.

Table 1 reports how we planned to evaluate or ensure that the generated taxonomy meets these criteria.

3.3 Phase 1: Taxonomy Generation

Considering how the taxonomy should be evaluated (see Table 1), we constructed a detailed prompt for GPT-4 for generating the first version of taxonomy. We made a few design choices here, including the depth of the taxonomy (single level) and the number of categories (4-6). We asked the LLM to generate labels, descriptions, and examples for these categories. The full prompt can be seen in the supporting material.

There were variations in how different versions described the same category. For instance, ‘Learning’ had slightly different meaning and definition in each version of the taxonomy, but generally included concepts and examples of understanding and explanation. Two researchers (two of the co-authors) discussed these three versions and decided to create a consolidated version of the taxonomy, which is shown in Table 2.

3.4 Phase 2: Taxonomy Validation

Next, we provided the taxonomy in Table 2 to two human coders along with 10 segments of conversation. They coded them independently, after which we compared their labels. They only agreed three out of 10 times. We repeated the whole process with a new

sample of 10 segments. It improved, but still had a high level of disagreement (60%). We, therefore, had another round of discussions and deliberations.

More than trying to reach a higher level of agreement, the goal here was to revise the current version of the taxonomy and develop a better understanding of how a reliable taxonomy could be generated that meets the criteria reported earlier and leads to a common and robust comprehension among the annotators. We learned that the annotators were often extrapolating why a user might have tried to do something. That led to most divergence among them. For instance, even when all we could interpret from the data that the user asked for factual information (e.g., “Does the state of Washington have income tax?”), one of the annotators often extended that to ‘Learning’ intent. It is possible that the user was collecting such information as a part a learning task, but without additional context, it may be impossible to determine that. In such cases, it is advisable to not overextend our understanding and mark the intent based on evidence. Thus, we found it was useful to include in the taxonomy not only positive examples, but also negative examples per category to improve overall clarity. The taxonomy was further modified using negative examples for each category, and the prompt for generating a taxonomy was edited to explicitly ask for negative examples (negative examples are not listed here due to space constraints).

Once we had GPT-4 provide such examples and clarify definitions of ‘Information Retrieval’ and ‘Learning’ categories, we achieved a good match with only 20% disagreements between the two annotators. In addition, the human assessors did not find a need for any intention not covered here. Thus, the validation state of taxonomy generation was completed and we had the final version of user intent taxonomy (see the supporting material referenced earlier).

3.5 Phase 3: Taxonomy Application and Testing

We then took a different set of 124 conversations and asked GPT-4 to code them using the modified taxonomy generated from the above process. We also gave the same instructions to two human assessors for annotation. These instructions to humans and the prompt to GPT-4 can be found in the supporting material.

For the human coders, not a single datapoint was labeled ‘Other’. GPT-4, on the other hand, marked one out of 124 conversations as

Table 1: Evaluating GPT-4 generated taxonomy using human assessors.

Criteria	How to evaluate/ensure?
Comprehensiveness	→ At the time of generation (Phase 1) by making sure the taxonomy considers as broad a scope as possible in the prompt to GPT-4. → At the time of validation (Phase 2) by looking at what proportion of instances by human assessors end up in the ‘Other’ category.
Consistency	→ At the time of generation (Phase 1) by making sure GPT-4 generates a taxonomy where different categories do not overlap in their meanings. → At the time of validation (Phase 2) by assessing how often the human assessors have difficulty distinguishing between two labels.
Clarity	→ At the time of generation (Phase 1) by ensuring GPT-4 provides detailed description or definition along with examples with each category. → At the time of validation (Phase 2) by eliciting from the human assessors how clear the definitions and examples are for them.
Accuracy	→ At the time of validation (Phase 2) by comparing annotations by at least two human assessors and computing inter-coder reliability (ICR).
Conciseness	→ At the time of application (Phase 3) by analyzing how well the generated and validated taxonomy serve the purpose of understanding user intents in chat logs after GPT-4 and/or human annotators use the taxonomy for annotation of test data.

Table 2: A consolidated version of the user intent taxonomy generated by GPT-4. The examples are collected from the three versions that GPT-4 generated. Slight modifications are made in the user intent title and description using those versions. Short phrases in parentheses next to the intent name are added by researchers for clarity.

User intent	Description	Examples
Information retrieval (looking for factual information that already exists)	The user wants to search, query, or find some information, data, or resources about a topic.	Find out the airing dates and channels of women’s world cup; Search for information about a phone number; Search for corruption and unemployment statistics for a country.
Problem solving (extracting facts or answers by computing something)	The user wants to perform a mathematical or logical operation, such as a conversion, a percentage, a formula, or a function.	Compare the size of a human to a hydrogen atom and the observable universe; Compare interest rates for savings accounts; Calculate the distance between a point and a line; Convert a message from Chinese to English.
Learning (satisfying curiosity, helping learn a concept or a phenomenon)	The user wants to learn, study, or acquire new skills, concepts, or understanding about a subject. This often involves operations of calculations, comparison, and conversion.	Learn about different structural systems; Compare GPT-3 and GPT-4 versions; Explain the difference between Newtonian and non-Newtonian flow.
Content creation	The user wants to write or edit a text for a specific purpose or audience.	Write an introduction about geothermal energy; Modify a poem into different formats; Improve a report and find adverbs and connectors.
Leisure	The user wants to chat or interact with the AI or another agent about various topics or play a game with the AI or another agent.	Ask about the AI’s sexual orientation and name; Listen to a romantic story; Play tennis and flirt with the user.

‘Other’. This further demonstrates comprehensiveness of the taxonomy. We computed inter-coder reliability (ICR) between two human coders and found Cohen’s kappa to be 0.7620, which indicates a substantial level of agreement [17].

Next, we asked a third coder to code these 124 conversations. This allowed us to triage among the three coders. When the three annotators disagreed, we took the majority vote. If all three picked a different label, we labeled that case as ‘Other’. Finally, we computed ICR between GPT-4 labels and those generated by the majority of human coders. We computed Cohen’s kappa to be 0.7212. This also indicates a substantial level of agreement.

Overall, what we learned is that when a taxonomy is generated by GPT-4 and verified by humans, it leads to a very high amount of agreement for annotation. That speaks to the **validity** of the generated taxonomy. Also, given that GPT-4’s own coding achieves a high level of ICR with human coders shows that GPT-4 can be used with high **reliability** for the annotation task.

3.6 Insights About and From Annotations

Now that we have demonstrated the end-to-end methodology for generating, validating, and using a taxonomy for understanding user intents in chat logs, let us consider what insights we could derive from the 124 conversation segments analyzed by annotators and GPT-4.

Table 3: Confusion matrix for user intent annotations between two human annotators. IR=Information Retrieval, PS=Problem Solving, LR=Learning, CR=Content Creation, LS=Leisure.

		Annotator-2				
		IR	PS	LR	CR	LS
Annotator-1	IR	42	2	10	0	0
	PS	0	8	0	4	0
	LR	3	0	36	0	0
	CR	0	0	0	8	0
	LS	1	0	0	0	9

Table 3 presents the confusion matrix between the two human annotators. We can see that Information Retrieval (IR) is the largest category, followed by Learning (LR). The greatest number of times the two annotators disagree is for IR and LR categories. This is understandable since LR always contains IR in a search setting, but it may not always be easy to evaluate if an IR process extends enough to qualify as LR. As noted earlier, this was the biggest factor leading to disagreements among the annotators.

Table 4 presents the confusion matrix between human annotations (after triaging of three annotators’ annotations) and those of GPT-4. Once again, we find that IR and LR are the largest categories and also where we see the most disagreements. Specifically, several

Table 4: Confusion matrix for user intent annotations between human and GPT-4 assessments. IR=Information Retrieval, PS=Problem Solving, LR=Learning, CR=Content Creation, LS=Leisure, OT=Other.

		GPT-4					
		IR	PS	LR	CR	LS	OT
Human	IR	46	1	5	0	0	1
	PS	0	8	2	0	0	0
	LR	12	3	26	1	0	0
	CR	0	0	0	10	0	0
	LS	0	0	3	0	5	0
	OT	1	0	0	0	0	0

(12 out of 124) conversation segments that are marked as LR by humans are labeled as IR by GPT-4. To understand who may be better or more appropriate in picking the labels here, we examined these conversations closely. They all include IR components, but the question is: do they go far enough to indicate an LR task?

Unfortunately, we do not have the ground truth here since we do not have access to the original user who conducted the conversation. We interviewed the assessors and found that they extended their understanding of what the users were doing in those segments of conversations to what they might want to do with that information beyond the logged interactions. This often led to a conversation segment being marked as an LR instead of an IR. GPT-4 here is strictly labeling the data without making further assumptions, which is desirable. But how consistent this LLM is while making such subjective decisions? To test this, we ran the same test data through GPT-4 four more times and measured ICR among the five sets of annotations by the LLM. We found Fleiss’ kappa [21] to be 0.8516, indicating a very high level of agreement and consistency. Therefore, we believe that the labels generated by GPT-4 are better than those generated by humans in this case as they are more objectively and consistently assigned without undesirable extrapolation or assumptions that may not be well-founded, addressing RQ3.

4 ADDITIONAL VALIDATIONS USING OPEN-SOURCE LLMs

While we selected GPT-4 as the LLM of choice for developing our method due to its state-of-the-art performance, we wanted to make sure that the results are not typical of GPT-4 and can be replicable by other LLMs, especially ones available as open-source. Therefore, we used Mistral³ and Hermes⁴, both available from Huggingface as open-source and free LLMs, to do similar experiments for taxonomy generation and application.

4.1 Single Level Taxonomy Generation

Similar to the process we instituted for GPT-4 during Phase 1 (Figure 2), we fed a prompt and the set of training conversations to Mistral and Hermes. Since this prompt was already validated and optimized for taxonomy generation, we skipped Phase 2. But we still needed a way to get some assurance that these two LLMs could also reliably generate user intent taxonomies. For this, we used

bootstrapping, where about 80% of the data was randomly sampled from the available data and provided to a given LLM with a prompt to generate a taxonomy. The prompt remained the same – the one that resulted from Phase 2 as described before. We ran this process 10 times with each of the three LLMs, each time resulting in a slightly different taxonomy. We performed a few minor manual adjustments to each version as needed. For instance, we substituted ‘Finding’ with ‘Information Retrieval’ and ‘Enjoy’ with ‘Leisure’ for category labels. Table 5 shows the union of all generated categories through 30 total runs, with each run generating five categories of user intents.

Table 5: Bootstrapping experiments showing frequency of different intent categories over 30 total runs, 10 runs for each of the three LLMs. Top 5 categories in each are bolded.

Category	GPT-4	Mistral	Hermes
Information retrieval/seeking/finding	10	9	10
Problem solving	9	8	8
Learning	8	10	9
Content creation	9	8	8
Leisure/Entertainment	8	10	7
Ask for advice/opinion	3	2	4
Chat	3	1	2
Verify	0	2	2

As Table 5 shows, while a few instances of taxonomies had intents that were different than what is reported in Table 2, most coalesced around the same set of five categories. More importantly, those categories were relatively common across each of the three LLMs. This demonstrated that (1) the prompt we validated and optimized for intent generation is robust and generalizable; and (2) the categories of user intents emerging from the available data are quite stable.

4.2 Multilevel Taxonomy Generation

While a taxonomy could have any number of levels, we have thus far focused on generating and applying single level or flat taxonomy only. This is not a limitation of the proposed methodology, but a manifestation of why we are generating such taxonomies and what we intend to do from their applications. The next section presents a specific application where our interest in understanding user intents is to explore how high level tasks are directed and done by two different modalities of information access (search and chat). A small, flat taxonomy is sufficient for this. However, other applications may need more levels or granularity.

To understand how well LLMs can do for generating multilevel taxonomies, we ran Phases 1 and 2 in Figure 2 with a slightly modified prompt. We added instructions to generate a taxonomy with two levels, with each node not having more than five children. This means we could have up to 30 intent categories for a taxonomy with two levels. This is not unprecedented. There are several works in the literature (e.g., [34, 36, 37]) that have defined and used 20 or more use intents through a multilevel structure.

³https://huggingface.co/docs/transformers/main/model_doc/mistral

⁴<https://huggingface.co/NousResearch/Nous-Hermes-Llama2-13b>

We ran this modified prompt through all three LLMs – GPT-4, Mistral, and Hermes – multiple times with bootstrapping. We found that level-1 of the taxonomies maintained the consistency close enough to what is shown in Table 5. Examining the subcategories generated for level-2, we found there to be more variance among the three LLMs. A similar behavior can be expected if different human coders were asked to generate up to 30 categories/subcategories. These initial disagreements among human coders or LLMs could be resolved with more data and more iterations. A full description of this process is beyond the scope for this paper, but a few important findings are reported below.

- (1) Consistencies in subcategories can be improved substantially by holding the level-1 categories constant.
- (2) Often subcategories are generated by the model to address very specific instances found in the data. For instance, one subcategory that often emerged was ‘Look for review’ under ‘Information Retrieval’. We examined the data and discovered several requests in which the user was asking for restaurant or movie reviews. This level of granularity may be essential for some applications, whereas it may add noise for others. If it is the latter, one could prune the taxonomy and remove such subcategories.
- (3) To generate meaningful subcategories, we found it to be useful to instruct the LLM that a given subcategory must have at least some minimum number of potential examples from the data, otherwise it should be removed or merged with another subcategory. This is similar to the notion of *support* in data mining [39].
- (4) In general, it seems essential that each new level of a taxonomy will need its own round of prompt optimization and human intervention for doing appropriate edits and pruning.

4.3 Taxonomy Application

We now move to Phase 3 of Figure 2 for the two open-source models. We executed this phase with Mistral and Hermes by giving them the same prompt and test data fed to GPT-4. We then measured ICR between each pair of GPT-4, Mistral, Hermes, and human annotations as a way to understand how humans and different LLMs agree on their understanding and application of the taxonomy. The results are shown in Table 6.

Table 6: ICR using Cohen’s Kappa.

	Human	GPT-4	Mistral	Hermes
Human	0.7620	–	–	–
GPT-4	0.7212	–	–	–
Mistral	0.6943	0.6343	–	–
Hermes	0.6521	0.5732	0.6772	–

As we can see, the ICR values as measured by Cohen’s Kappa were moderate (0.41-0.60) to substantial (0.61-0.80). This is an indirect validation of the taxonomy that was generated as it indicates that the categories and their descriptions were comprehensive, consistent, clear, accurate, and concise enough that different entities – humans and three LLMs – could apply them to new data in a very similar way. There are still other elements of a good taxonomy that cannot be measured by such ICR scores. For example, the usefulness

and objectivity of the taxonomy may depend on the application. In the next section, we consider a specific application of user intent taxonomy as a way to further validate our proposed methodology.

5 APPLICATION OF USER INTENT PIPELINE

Can we apply the proposed methodology to other, perhaps more challenging tasks that call for identifying user intents? We were driven by a hypothesis that there is a shift in users’ behaviors and intents happening with the emergence of AI chat modality. To test this, we built a pipeline using the method proposed here. Depicted in Figure 3, the pipeline shows that there is still a human in the loop but given that we know how to construct and apply a reliable taxonomy using a process that is already validated, we can now use human intervention in shaping the process and doing light touch validations. Notably, the intent taxonomy here is created with heterogeneous log data containing only user requests.

5.1 Stepwise Process for the Full Pipeline

In the steps below, we describe this process as a full pipeline for how one could leverage LLM for analyzing log data. Through the process, we will also focus on evaluating various aspects of the generated taxonomy. These aspects include comprehensiveness, consistency, clarity, accuracy, and conciseness.

But why generate a new taxonomy if there are several existing taxonomies, including the one generated in the previous section? While one of these taxonomies could be fitting, it is desirable that we have a taxonomy that is rooted in specific data and application we have under consideration. Given the small cost of generating a taxonomy may also justify at least attempting to construct a new taxonomy and deciding if it is more fitting than anything available.

Step-1: Identify Application and Data

The first step is to identify what kind of data we want to extract user intents from and why. Here, we are interested in understanding how users have different or overlapping intents between two modalities – search and chat. Using log data available to use from Bing Search and Bing Chat, we needed to first build a new user intent taxonomy and then apply that taxonomy to annotate log data. We extracted a random sample of users who had used both Bing Search and Bing Chat from May-June 2023. From those users, we extracted 2,456 queries and 15,531 chat requests they had sent to the respective services. We used 500 search queries and 500 chat requests (a total of 1,000 user inquiries to Bing) for training and set aside the rest for testing. We randomized their order, forming our training set with 1,000 data points. There are two ways this data is different from the data used earlier, making this a different and a more challenging task: (1) it contains log data from different modalities; and (2) it does not contain system response.

Step-2: Build/Fine-tune Taxonomy with Human-in-the-loop

To get started with the LLM (here, GPT-4), we built the initial prompt that explained what we are trying to do, what the data contains, and what are some of the criteria or constraints. For example, we indicated that we are looking for a taxonomy of user intents with no more than five categories and the criteria for a good taxonomy are comprehensiveness, consistency, clarity, accuracy, and conciseness as defined earlier. The full prompt is given in the

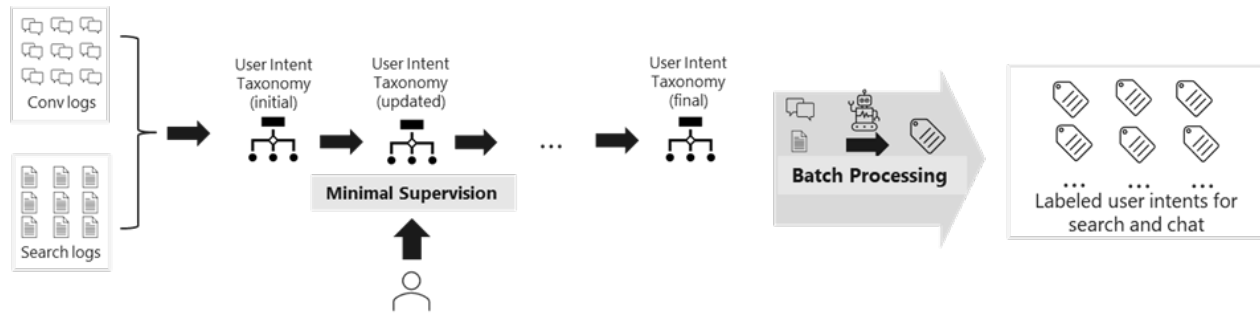


Figure 3: Using LLM in an end-to-end pipeline for generating, validating, and applying a taxonomy for user intents.

supporting material. This resulted in the zero-shot version of the taxonomy with the following five categories:

- *Ask for Advice or Recommendation*: The intent to seek suggestions, opinions, or guidance from others on a specific topic or situation.
- *Create*: The intent to use AI tools or platforms to generate, edit, or manipulate information objects.
- *Information Retrieval*: The intent to find existing information or answers on the internet.
- *Learn*: The intent to acquire new knowledge or skills on a subject of interest.
- *Leisure*: The intent to enjoy oneself by engaging in amusing activities such as games, jokes, stories, etc.

This is not very different from what we saw in the previous section. However, we noticed that the descriptions and examples associated with these labels were different and more suitable for our purpose. Even if there was evidence or intuition that an existing taxonomy would be sufficient for our purpose, given the reasonable cost for generating a new taxonomy, it may be desirable to go through these two steps to validate and revise that taxonomy with a goal to fare better along the criteria for a quality taxonomy described before.

Step-3: Measure Taxonomy Comprehensiveness/Consistency

We now need to test how complete and consistent this taxonomy is. For that, we fed it as a prompt to the LLM and have it label each of the samples we used before separately. This time, we asked it explicitly to label anything that does not fit the provided labels as 'Other'. We found that no sample fell under this category. This indicated that the taxonomy was comprehensive and consistent.

Step-4: Improve Taxonomy Clarity

Next, we asked the LLM to expand each category label with more description and examples to improve its clarity. Taking the lesson from before, we also asked GPT-4 for negative examples per category, improving on the taxonomy's clarity.

Step-5: Measure Validity and Accuracy

As the final step of validation and refinement, we asked the LLM to use the constructed taxonomy to label the same data that was used to generate the taxonomy. Normally, this is not a practice for testing, but here we are looking for internal validity and accuracy of the taxonomy. Recall that we had 1,000 data points (500 search queries and 500 chat requests) for training. Once the LLM labeled each of these, we took a random sample of 100 and manually checked if the assigned label follows the definition for that label as generated before. We found that the answer was 'yes' for 95 of these samples and that there was no sample assigned 'Other' category.

This analysis provided us with the assurance the taxonomy was valid and accurate.

Step-6: Perform Annotations and Measure Conciseness

Finally, we ran our test data – 1,956 search queries and 15,031 chat requests – through GPT-4 with the final version of the taxonomy as a part of the prompt. This prompt is given in the supporting material. We found that no sample was marked with 'Other' label, ensuring that all the important concepts were covered. Also, no category had too few (subjective, but in our case < 2%) samples, indicating that the taxonomy was concise.

5.2 Insights about Intents in Search vs. Chat

The steps above demonstrated that we could create a user intent taxonomy fulfilling all the criteria for a high-quality, reliable, and robust taxonomy. We did not have ground truth for the test data to be able to compute accuracy, but given the rigor taken in building and validating the taxonomy, we feel confident in the labels generated by the LLM. If one needs additional assurance at this point, a small sample of this test data can be taken for human assessment and ICR can be computed between that assessment and the one from the LLM.

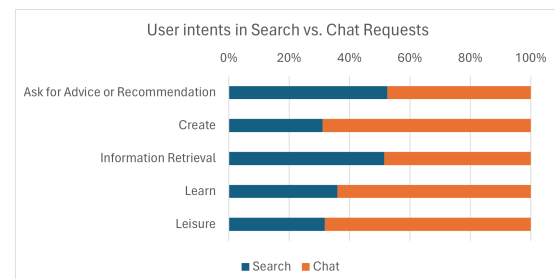


Figure 4: Comparing user intents between search and chat.

For our purposes, we decided to move on to deriving insights from this test data. Given that we had an uneven number of queries and chat requests, we normalized them around each intent category by counting the number of instances for a given modality w.r.t. an intent and dividing it by the total instances of search and chat for that intent. Figure 4 shows the distribution of user intents for search and chat. As shown, 'Ask for Advice or Recommendation' and 'Information Retrieval' are almost evenly distributed between search and chat, with a little bit of skew toward search. The other

three categories ('Create', 'Learn', and 'Leisure') are heavily leaning toward chat. This requires a close examination.

First off, it is important to understand that this figure shows a view from the user intent perspective. If a user had an intent or a task related to one of the five intents considered here, where would they go – search or chat? We found that while they could use either for their 'Information Retrieval' or 'Ask for Advice or Recommendation' needs, they are favoring (with a significant tilt toward) chat for their create, learn, and leisure intents. 'Create' makes sense because it is more suitable to use a generative AI tool like Bing Chat for a creation task than a search engine. Of course, users are still sending search engines with create-related requests, but we hypothesize that as generative IR tools such as AI chat become more capable and known, that intent will shift more dramatically from search to chat. Similarly, 'Leisure' makes sense because here the user intends to have a social or open-ended conversation, which is by design supported through chat. Perhaps more interesting finding here is with respect to 'Learn'. Learning is considered to be a higher-level goal or task in information seeking [49]. While people have used keyword-based search system for such a task, with chat-based generative IR systems, the intent fits the modality more appropriately. Through a manual inspection of some of the logs available to us, we could see that users are indeed issuing higher-level and complex requests, often associated with learning, through the chat interface. We should note this with a caveat that our unit of analysis here is a single request from the user. It is possible that the user issued multiple queries in a given search session to accomplish their learning task. Even then, it is interesting to learn that users are preferring to issue their single-request learning requests through chat. As the information access systems with emerging technology such as generative AI and conversation-based modalities chart their course to support users in new meaningful ways, they should consider their designs from a user intent perspective.

5.3 Steps for Generating Intent Taxonomies

Now that we have described the methodology and demonstrated how it could be executed using an application, we summarize the lessons from these experiments and provide a guidance to anyone who wants to use LLMs for generating, validating, or applying user intent taxonomies.

(1) *Identify Application and Data.* A taxonomy must fulfill the purpose for which it is built. That also means an existing taxonomy may not be right for your application. Assuming you want to build a purpose-driven taxonomy, prepare a detailed description of what user intent means for your application and how it should be used. For instance, in our case, it was important for us to stay focused on users' actions in a task rather than the objects when recognizing intents. This means we would not want intents that are tied to an object (e.g., 'finding information about tax') and stay close to a general action or objective (e.g., 'information retrieval'). It is also important to have as clean data as possible for an LLM to process it appropriately. Depending on which LLM you use, you may need to check for input requirements such as the size and language of input tokens.

(2) *Build and Fine-tune Taxonomy.* Pay attention to the first prompt you prepare for building the taxonomy. Add details of your application/task, your criteria for a good taxonomy (see Table 1), and relevant constraints (e.g., number of levels, number of categories, length of a label). We recommend bootstrapping to build different versions of the taxonomy to see how sensitive it is to the data used, like what we showed in Section 3.3. If a zero-shot taxonomy built with your LLM with this prompt is good enough for you, you can skip the following sub steps.

(a) *Check for Comprehensiveness.* Construct a prompt to annotate input data using the taxonomy built. Feed the training data to the LLM with this prompt to have it label the data. If what falls under 'Other' category is more than, say, 5% of your data, you may need to create additional categories or levels to make your taxonomy more complete.

(b) *Check for Consistency.* Assuming your training data is of a reasonable size, it may not be feasible to manually check labels for each of the samples, but you can take an appropriate random sample and see if the LLM consistently applied the definitions of various categories. You could also perform multiple runs of Step-2 and see if a sample gets labeled the same way every time. Since this is an iterative and exploratory process, you can decide how far and deep you want to go before seeing good enough convergence and consistency.

(c) *Improve Taxonomy's Clarity.* Once the above steps are done reasonably well or skipped as appropriate, your taxonomy is now fixed. At this point, you may ask the LLM to revise and expand the definition or description for each of your labels to improve its clarity. Often, feeding appropriate examples (positive and negative) can be useful – similar to how a human annotator is trained.

(3) *Measure Accuracy and Conciseness.* As a final and another optional step, you can give the LLM test data, ensuring this data was not used before for any training purposes, for doing annotations using the final version of the taxonomy. Is there any category that does not get enough samples? If so, you may decide to remove that category to improve your taxonomy's conciseness. Note that if you do this, you may have to repeat some of the steps from before because now those samples will fall under other categories, which may affect some of the criteria evaluated before. Now take a random sample of labeled data and have a human annotator label it using the same instructions given to the LLM as prompt. Measure the ICR between human annotations and those from the LLM for the same data. This measurement will give you a sense of how accurate or valid your taxonomy is as well as your LLM's annotation capabilities. If at this point you have taken all the steps before (or skipped them as appropriate) and found a high enough ICR, your taxonomy and your LLM have been thoroughly tested.

6 CONCLUSION

Identifying user intents in online information access is highly crucial for most search and recommender systems. But doing so is often very challenging. Even if one has a pre-defined taxonomy of user intents, training an ML model or using such a model to

annotate rapidly changing behavioral traits in new modalities such as AI chat can be expensive or infeasible. LLMs are shown to be effective at extracting concepts, descriptions or summaries, and examples from given set of text. This could be used for building and using taxonomies containing user intents, but there is a danger of creating a feedback loop without a clear evaluation.

In this paper we presented a novel methodology for using LLMs in generating, validating, and using taxonomies for identifying user intents in various applications. The methodology was demonstrated using an application of understanding user intents in AI chat logs. A case study was then presented with the application of contrasting user intents between search and chat. The results from both the applications are intriguing, presenting a set of new hypotheses and calling for further explorations. However, the primary contribution of this paper is the methodology for deploying LLMs in such research tasks.

As a reference from our own experiments, building the full pipeline in that case study (Section 5) took less than half the time and effort compared to the process executed for developing the method (Section 3). The process described in Section 5.3 will take substantially less even if all the optional steps are executed. Such efficiency is more than simply reducing the effort for one set of experiments. Emerging technologies such as AI-driven chat are being discovered and used by a large set of new users. As they become more accustomed, we can expect to see shifts in the kind of tasks they do and the kinds of intents they have with these modalities. The approach presented here will allow researchers to adapt to these evolving intents quickly and at lower cost and effort.

Through the development of this methodology, we learned that we could use various LLMs for a zero-shot construction of a user intent taxonomy, given some log data with user requests in natural language. While this taxonomy is of reasonably high quality, we found the need to have human verification and fine tuning to ensure that such a taxonomy meets various criteria commonly expected in the literature and in practice, including comprehensiveness, consistency, clarity, accuracy, and conciseness. Through the development of this methodology and its subsequent application in a different case study, we showed how these criteria can be reliably fulfilled using an LLM and human-in-the-loop.

In this regard, we conclude that an LLM can serve as a collaborator or a copilot rather than a replacement for human researchers. This human-LLM collaboration can yield not only faster construction and validation of a new user intent taxonomy, but also higher quality outputs with crisply defined labels, descriptions, and examples. Once the phases of construction and validation are done, the LLM can very effectively and accurately perform the LLM annotation task, turning from copilot to autopilot. This can allow us to analyze large-scale data and generate insights. Finally, we found that often LLMs not only made things go faster, but also better. In cases of disagreements with human annotators, we found that GPT-4 was producing user intent labels truer to the data given rather than extrapolating to situations for which we lacked evidence. In short, the work reported here charts a new territory for using LLMs as collaborators and consignors for user intent analysis in an effective, efficient, and responsible manner.

REFERENCES

- [1] Marissa D Abram, Karen T Mancini, and R David Parker. 2020. Methods to integrate natural language processing into qualitative research. *International Journal of Qualitative Methods* 19 (2020), 1609406920984608.
- [2] Eugene Agichtein, Ryen W White, Susan T Dumais, and Paul N Bnetten. 2012. Search, interrupted: understanding and predicting search task continuation. In *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval*. 315–324.
- [3] Qingyao Ai, Ting Bai, Zhao Cao, Yi Chang, Jiawei Chen, Zhumin Chen, Zhiyong Cheng, Shoubin Dong, Zhicheng Dou, Fuli Feng, et al. 2023. Information Retrieval Meets Large Language Models: A Strategic Report from Chinese IR Community. *arXiv preprint arXiv:2307.09751* (2023).
- [4] Dogu Araci. 2019. Finbert: Financial sentiment analysis with pre-trained language models. *arXiv preprint arXiv:1908.10063* (2019).
- [5] Muneera Bano, Didar Zowghi, and Jon Whittle. 2023. Exploring Qualitative Research Using LLMs. *arXiv preprint arXiv:2306.13298* (2023).
- [6] Garbiel Bénédict, Ruqing Zhang, and Donald Metzler. 2023. Gen-IR@ SIGIR 2023: The First Workshop on Generative Information Retrieval. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 3460–3463.
- [7] Valeriia Bolotova, Vladislav Blinov, Falk Scholer, W Bruce Croft, and Mark Sanderson. 2022. A non-factoid question-answering taxonomy. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1196–1207.
- [8] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3, 2 (2006), 77–101.
- [9] Andrei Broder. 2002. A taxonomy of web search. In *ACM Sigir forum*, Vol. 36. ACM New York, NY, USA, 3–10.
- [10] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [11] Courtnei Byun, Piper Vasicek, and Kevin Seppi. 2023. Dispensing with Humans in Human-Computer Interaction Research. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–26.
- [12] Belen Carrion, Teresa Onorati, Paloma Diaz, and Vasiliki Triga. 2019. A taxonomy generation tool for semantic visual analysis of large corpus of documents. *Multimedia Tools and Applications* 78 (2019), 32919–32937.
- [13] Soumen Chakrabarti, Byron Dom, Rakesh Agrawal, and Prabhakar Raghavan. 1997. Using taxonomy, discriminants, and signatures for navigating in text databases. In *VLDB*, Vol. 97. 446–455.
- [14] Long Chen, Dell Zhang, and Levene Mark. 2012. Understanding user intent in community question answering. In *Proceedings of the 21st international conference on world wide web*. 823–828.
- [15] Pu-Jeng Cheng, Ching-Hsiang Tsai, Chen-Ming Hung, and Lee-Feng Chien. 2006. Query taxonomy generation for web search. In *Proceedings of the 15th ACM international conference on Information and knowledge management*. 862–863.
- [16] Shui-Lung Chuang and Lee-Feng Chien. 2003. Automatic query taxonomy generation for information retrieval applications. *Online Information Review* 27, 4 (2003), 243–255.
- [17] Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and psychological measurement* 20, 1 (1960), 37–46.
- [18] Nick Craswell, Onno Zoeter, Michael Taylor, and Bill Ramsey. 2008. An experimental comparison of click position-bias models. In *Proceedings of the 2008 international conference on web search and data mining*. 87–94.
- [19] Yogesh K Dwivedi, Nir Kshetri, Laurie Hughes, Emma Louise Slade, Anand Jayaraj, Arpan Kumar Kar, Abdullah M Baabdullah, Alex Koohang, Vishnupriya Raghavan, Manju Ahuja, et al. 2023. “So what if ChatGPT wrote it?” Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. *International Journal of Information Management* 71 (2023), 102642.
- [20] Guglielmo Faggioli, Laura Dietz, Charles Clarke, Gianluca Demartini, Matthias Hagen, Claudia Hauff, Noriko Kando, Evangelos Kanoulas, Martin Potthast, Benno Stein, et al. 2023. Perspectives on Large Language Models for Relevance Judgment. *arXiv preprint arXiv:2304.09161* (2023).
- [21] Joseph L Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin* 76, 5 (1971), 378.
- [22] Jianfeng Gao, Chenyan Xiong, and Paul Bennett. 2020. Recent advances in conversational information retrieval. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2421–2424.
- [23] Philipp Grandeit, Carolyn Haberkern, Maximiliane Lang, Jens Albrecht, and Robert Lehmann. 2020. Using BERT for qualitative content analysis in psychosocial online counseling. In *Proceedings of the Fourth Workshop on Natural Language Processing and Computational Social Science*. 11–23.
- [24] Xingwei He, Zhenghao Lin, Yeyun Gong, Alex Jin, Hang Zhang, Chen Lin, Jian Jiao, Siu Ming Yiu, Nan Duan, Weizhu Chen, et al. 2023. Annollm: Making

- large language models to be better crowdsourced annotators. *arXiv preprint arXiv:2303.16854* (2023).
- [25] Jeffrey Heer. 2019. Agency plus automation: Designing artificial intelligence into interactive systems. *Proceedings of the National Academy of Sciences* 116, 6 (2019), 1844–1850.
- [26] Tom Hope, Doug Downey, Daniel S Weld, Oren Etzioni, and Eric Horvitz. 2023. A computational inflection for scientific discovery. *Commun. ACM* 66, 8 (2023), 62–73.
- [27] Bernard J Jansen, Danielle L Booth, and Amanda Spink. 2007. Determining the user intent of web search engine queries. In *Proceedings of the 16th international conference on World Wide Web*. 1149–1150.
- [28] Vitor Jeronymo, Luiz Bonifacio, Hugo Abonizio, Marzieh Fadaee, Roberto Lotufo, Jakub Zavrel, and Rodrigo Nogueira. 2023. InPars-v2: Large Language Models as Efficient Dataset Generators for Information Retrieval. *arXiv preprint arXiv:2301.01820* (2023).
- [29] Angelika Kaplan, Thomas Kühn, Sebastian Hahner, Niko Benkler, Jan Keim, Dominik Fuchß, Sophie Corallo, and Robert Heinrich. 2022. Introducing an Evaluation Method for Taxonomies. In *Proceedings of the 26th International Conference on Evaluation and Assessment in Software Engineering*. 311–316.
- [30] Melanie Kellar, Carolyn Watters, and Michael Shepherd. 2007. A field study characterizing Web-based information-seeking tasks. *Journal of the American Society for information science and technology* 58, 7 (2007), 999–1018.
- [31] Weize Kong, Rui Li, Jie Luo, Aston Zhang, Yi Chang, and James Allan. 2015. Predicting search intent based on pre-search context. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 503–512.
- [32] Dennis Kundsich, Jan Muntermann, Anna Maria Oberländer, Daniel Rau, Maximilian Röglinger, Thorsten Schoormann, and Daniel Szopinski. 2021. An update for taxonomy designers: methodological guidance from information systems research. *Business & Information Systems Engineering* (2021), 1–19.
- [33] Shuokai Li, Ruobing Xie, Yongchun Zhu, Xiang Ao, Fuzhen Zhuang, and Qing He. 2022. User-centric conversational recommendation with multi-aspect user modeling. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 223–233.
- [34] Yuelin Li and Nicholas J Belkin. 2008. A faceted approach to conceptualizing tasks in information seeking. *Information processing & management* 44, 6 (2008), 1822–1837.
- [35] Helen Lippell. 2022. *Taxonomies: Practical Approaches to Developing and Managing Vocabularies for Digital Information*. Facet Publishing.
- [36] Jiqun Liu, Matthew Mitsui, Nicholas J Belkin, and Chirag Shah. 2019. Task, information seeking intentions, and user behavior: Toward a multi-level understanding of Web search. In *Proceedings of the 2019 conference on human information interaction and retrieval*. 123–132.
- [37] Matthew Mitsui, Chirag Shah, and Nicholas J Belkin. 2016. Extracting information seeking intentions for web search sessions. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*. 841–844.
- [38] Steven Moore, Richard Tong, Anjali Singh, Zitao Liu, Xianguan Hu, Yu Lu, Joleen Liang, Chen Cao, Hassan Khosravi, Paul Denny, et al. 2023. Empowering Education with LLMs-The Next-Gen Interface and Content Generation. In *International Conference on Artificial Intelligence in Education*. Springer, 32–37.
- [39] Tadeusz Morzy, Marek Wojciechowski, and Maciej Zakrzewicz. 2000. Data mining support in database management systems. In *Data Warehousing and Knowledge Discovery: Second International Conference, DaWaK 2000 London, UK, September 4–6, 2000 Proceedings 2*. Springer, 382–392.
- [40] Chau Minh Pham, Alexander Hoyle, Simeng Sun, and Mohit Iyyer. 2023. TopicGPT: A Prompt-based Topic Modeling Framework. *arXiv preprint arXiv:2311.01449* (2023).
- [41] Joe Raad and Christophe Cruz. 2015. A survey on ontology evaluation methods. In *Proceedings of the International Conference on Knowledge Engineering and Ontology Development, part of the 7th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management*.
- [42] Filip Radlinski and Nick Craswell. 2017. A theoretical framework for conversational search. In *Proceedings of the 2017 conference on conference human information interaction and retrieval*. 117–126.
- [43] Pengjie Ren, Zhongkun Liu, Xiaomeng Song, Hongtao Tian, Zhumin Chen, Zhaochun Ren, and Maarten de Rijke. 2021. Wizard of search engine: Access to information through conversations with search engines. In *Proceedings of the 44th International ACM SIGIR Conference on research and development in information retrieval*. 533–543.
- [44] Daniel E Rose and Danny Levinson. 2004. Understanding user goals in web search. In *Proceedings of the 13th international conference on World Wide Web*. 13–19.
- [45] Karan Singhal, Shekoofeh Azizi, Tao Tu, S Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, et al. 2023. Large language models encode clinical knowledge. *Nature* (2023), 1–9.
- [46] Scott Spangler and Jeffrey Kreulen. 2002. Interactive methods for taxonomy editing and validation. In *Proceedings of the eleventh international conference on Information and knowledge management*. 665–668.
- [47] Jaime Teevan, Eytan Adar, Rosie Jones, and Michael AS Potts. 2007. Information re-retrieval: Repeat queries in Yahoo’s logs. In *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*. 151–158.
- [48] Petter Törnberg. 2023. How to use LLMs for Text Analysis. *arXiv preprint arXiv:2307.13106* (2023).
- [49] Pertti Vakkari. 2016. Searching as learning: A systematization based on literature. *Journal of Information Science* 42, 1 (2016), 7–18.
- [50] Alexandra Vtyurina, Charles LA Clarke, Edith Law, Johanne R Trippas, and Horatiu Bota. 2020. A mixed-method analysis of text and audio search interfaces with varying task complexity. In *Proceedings of the 2020 ACM SIGIR on International Conference on Theory of Information Retrieval*. 61–68.
- [51] Ryan Watkins. 2023. Guidance for researchers and peer-reviewers on the ethical use of Large Language Models (LLMs) in scientific research workflows. *AI and Ethics* (2023), 1–6.
- [52] Ryen W White and Steven M Drucker. 2007. Investigating behavioral variability in web search. In *Proceedings of the 16th international conference on World Wide Web*. 21–30.
- [53] Hong Xie. 2002. Patterns between interactive intentions and information-seeking strategies. *Information processing and Management* 38, 1 (2002), 55–77.
- [54] Hui Yang. 2012. Constructing task-specific taxonomies for document collection browsing. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*. 1278–1289.
- [55] Oren Zamir and Oren Etzioni. 1998. Web document clustering: A feasibility demonstration. In *Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval*. 46–54.
- [56] Hongfei Zhang, Xia Song, Chenyan Xiong, Corby Rosset, Paul N Bennett, Nick Craswell, and Saurabh Tiwary. 2019. Generic intent representation in web search. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 65–74.
- [57] Peiyan Zhang, Jiayan Guo, Chaozhuo Li, Yueqi Xie, Jae Boum Kim, Yan Zhang, Xing Xie, Haohan Wang, and Sunghun Kim. 2023. Efficiently leveraging multi-level user intent for session-based recommendation via atten-mixer network. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*. 168–176.