Understanding Cognitive Intents in Conversational Information Seeking

Shuoqi Sun

(sure-chi swn)

PhD Candidate, RMIT University, Australia

Research Intern, National Institute of Informatics, Japan

Topical Search Intent

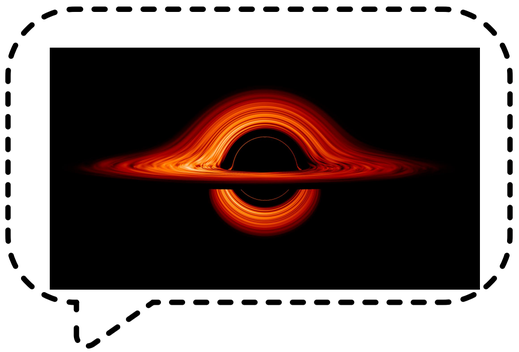



Black hole

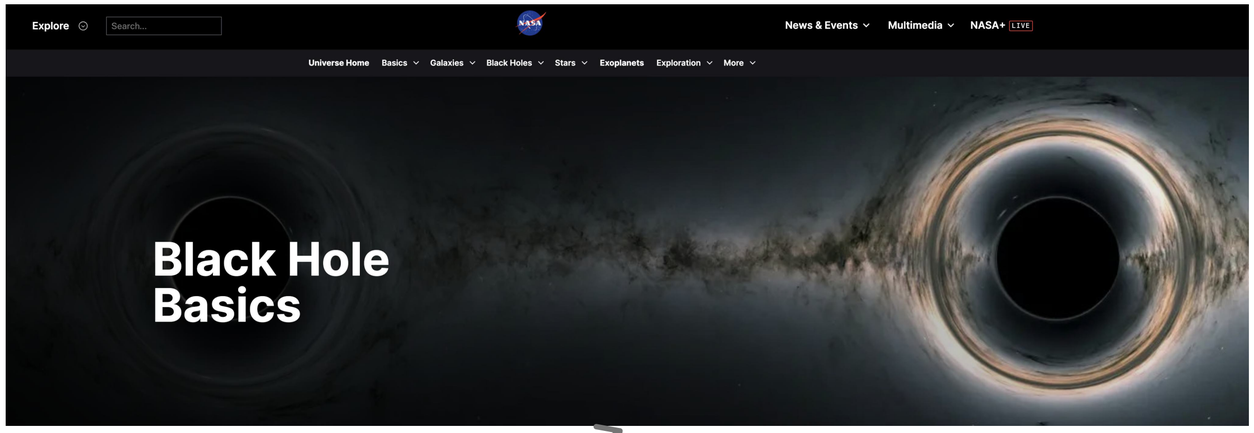

Topical search intent:

Users' needs for topics of the obtained information.
Topics are actually written in the obtained information.

Cognitive Search Intents (CSIs)



I want to have easy-to-understand information



Simple explanation black hole

Potential Unsatisfactory

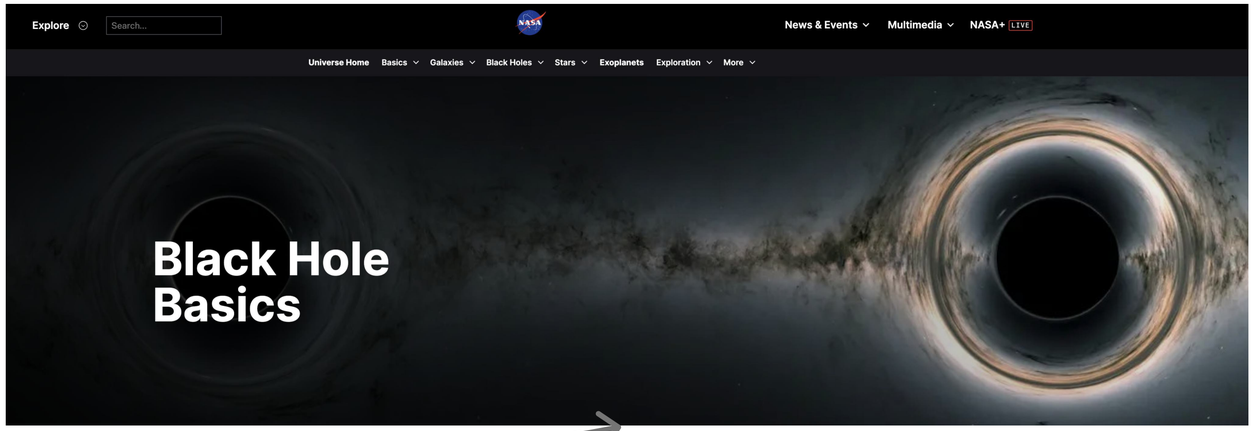

Cognitive Search Intent



Cognitive Search Intents (CSIs)



Cognitive search intent:

Users' needs for cognitive characteristics of the obtained information.

Cognitive characteristics of documents are perceived by information seekers,

.... which can be affected by the writing style, diction, structure, appearance, and meta data such as the author, domain, creation date, etc.

Kato, Makoto P., et al. "Investigating users' query formulations for cognitive search intents." SIGIR 2014

Cognitive Search Intents (CSIs)



needs for
charact

Exhaustiveness. The exhaustiveness (or coverage) of a document for a query is the fraction of query topics covered by the document, where the query topics are information generally expected from the query. Exhaustiveness has been considered to be an important factor by information seekers, and has been used as a document ranking criterion [47] as well as a summarization criterion [23]. We selected ten topics from the queries of NTCIR-10 INTENT-2 Japanese Subtask to investigate CSIs regarding exhaustiveness [35], which is an IR evaluation campaign for search result diversification and its queries were designed to have multiple topics, *e.g.* "Yahoo!" and "South Africa".

Comprehensibility. The comprehensibility of a document means how easily users can read and understand the document. Even though a retrieved document is perfectly relevant in terms of the topic, it is useless unless the user can comprehend it. Some researchers have tried to estimate comprehensibility by using a Wikipedia link structure [28], and a PageRank-based method that propagates comprehensibility scores in a semi-supervised manner [3]. We selected ten topics that were used in the research by Nakatani, Jatowt, and Tanaka, which came from difficult domains such as medicine and economics [28].

Subjectivity and Objectivity. The subjectivity of a document is represented by the fraction of personal opinions in the document, while objectivity is the opposite concept of subjectivity. There has been some work on estimating the subjectivity of documents, mainly within the context of opinion mining. For example, Yu and Hatzivassiloglou proposed a method of separating opinions from fact at the document and sentence levels [46]. The ten topics we used were extracted from queries used in TREC 2008 Blog Track, where opinion-finding and polarity opinion-finding retrieval tasks were carried out, and they used topics that could be described subjectively such as "Mozart" and "Blackberry" [29].

Concreteness and Abstractness. Tanaka *et al.* defined concreteness by combining perceivability and imaginability in the first trial to computationally estimate the concreteness of documents [38]. Perceivability is a measure of the extent to which people can sense an object, while imaginability is a measure of how quickly and easily people can imagine a given referent. Since it is difficult to convey to subjects notions of their own concreteness, we defined the concreteness of a document in a simpler way: to what extent the content of the document describes a topic with a specific scope and actual examples. In contrast, abstractness is defined as to what extent the content of the document describes a topic in a general and definitional way. We selected ten topics from the ones that Tanaka *et al.* used [38] for concreteness and abstractness, which were either queries used in TREC Million Query Tracks [12], or the titles of Wikipedia articles. Examples of these topics are "arithmetic" and "monopoly".

Kato, Makoto P., et al. "Investigating users' query formulations for cognitive search intents." SIGIR 2014

CSIs In a Decade Ago



needs for cognitive characteristics

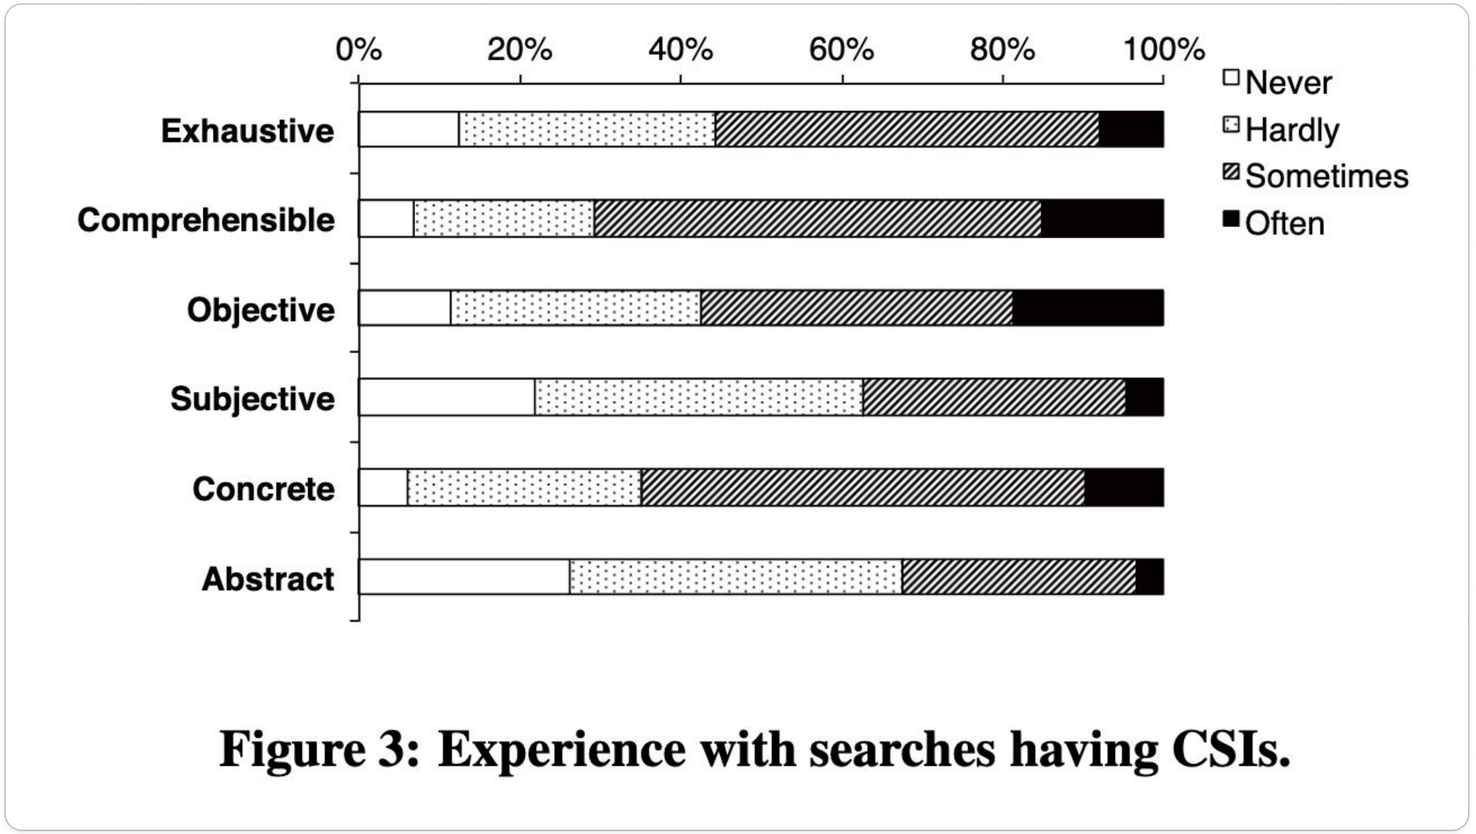

Figure 3: Experience with searches having CSIs.

Kato, Makoto P., et al. "Investigating users' query formulations for cognitive search intents." SIGIR 2014

CSIs In a Decade Ago

needs for cognitive characteristics



- **Easy-to-understand** black hole explanations
- black hole **examples**
- black hole **all I need to know**



Table 3: Query types for CSIs.

CSI	#	Query type (%)			
		direct	trans.	none	other
Exhaustive	134	0.7	36.6	**62.7**	0.0
Comprehensible	166	13.3	32.5	**53.0**	1.2
Objective	114	2.6	31.6	**65.8**	0.0
Subjective	128	0.8	**62.5**	36.7	0.0
Concrete	124	5.6	44.4	**49.2**	0.8
Abstract	84	0.0	45.2	**54.8**	0.0
Total	750	4.5	41.6	**53.5**	0.4

Kato, Makoto P., et al. "Investigating users' query formulations for cognitive search intents." SIGIR 2014

CSIs In a Decade Ago



Hypothesis:

People do not verbalize intents could because the over-adoption that they know Web search engine do not process such words effectively.

- **Easy-to-understand** black hole explanations
- black hole **examples**
- black hole **all I need to know**

Table 4: Most frequent terms in *transformed* queries.

CSI	Topic independent (%)	Topic dependent (%)
Exh.	details (13.4%) wiki (6.7%) summary (5.2%)	type (1.5%)
		cine (0.6%) logy (0.6%)
		(0.9%)
Sub.	blog (10.2%) opinion (2.3%) about (1.6%)	impression (20.3%) usability (3.9%) word-of-mouth (3.1%)
Con.	concrete example (15.3%) example (8.9%) details (3.2%)	real experience (0.8%) real world (0.8%)
Abs.	summary (8.3%) wiki (7.1%) Wikipedia (4.8%)	

Kato, Makoto P., et al. "Investigating users' query formulations for cognitive search intents." SIGIR 2014

Modern Information Seeking

What's on the

Ask ChatGPT

Help me organize my notes so I can find what

Google

easy-to-understand black hole

AI Mode | All | Images | Videos | Shopping | Short videos | Forums | More | Tools

AI Overview

A black hole is a region in space where a huge amount of matter is packed into a tiny area, creating gravity so strong that nothing—not even light—can escape. NASA Science (.gov) +1



Key Parts of a Black Hole

- **The Event Horizon:** This is the outer boundary or the "point of no return". Once anything crosses this line, the pull is too strong to ever get back out. (Think of a boat drifting past the point of no return near a giant waterfall).
- **The Singularity:** This is the very center of the black hole. All the crushed-up matter is

Show more

NASA Science (.gov)
Black Holes - NASA Science
Black Holes. Black holes are among the most mysterious cosmic objects, much studied b…



UChicago News
What is a black hole? | University of Chicago News
Definition: Regions of massive density and extreme gravity preventing light escape,…



NASA Science (.gov)
https://science.nasa.gov › universe › black-holes

Black Hole Basics

12 Aug 2026 — A **black hole** is so dense that gravity just beneath its surface, the event horizon, is strong enough that nothing – not even light – can escape. Read more

Modern Information Seeking

Can you give me a simple general study plan for daily use?

REASON Content Expectation *I asked this to get a simple study plan without personalization.*

Can you make it even simpler and easier to follow daily?

REASON Style Expectation *I asked this to make the study plan easier to follow consistently.*

Should I include bitcoin in my long term FIRE portfolio? answer in very short sentences. be brief.

REASON Task Motivation & Goal *I want to learn about why or why not bitcoin is good for investment.*

Jin, Chuanyang, et al. "Thoughttrace: Understanding user thoughts in real-world llm interactions." arXiv preprint arXiv:2605.20087 (2026).

Research Questions

Understanding
Cognitive Search Intents

RQ 1: What **types of cognitive intents** (CIs) people have in conversational information seeking with generative AI?

RQ 2: When people have cognitive intents, **do they verbalize** them in their prompts?

RQ 3: **How people verbalize** their cognitive intents?

Research Questions

Understanding
Cognitive Search Intents



RQ 1: What **types of cognitive intents** (CIs) people have in conversational information seeking with generative AI?

RQ 2: When people have cognitive intents, **do they verbalize** them in their prompts?

RQ 3: **How people verbalize** their cognitive intents?

Step 1: Updating CSIs Taxonomy

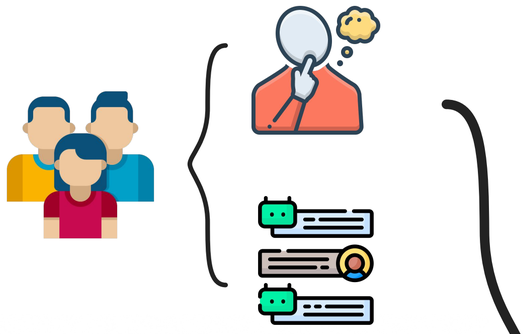

- Exhaustiveness
- Comprehensibility
- Subjectivity
- Objectivity
- Concreteness
- Abstractness

ThoughtTrace dataset

- 1,058 users
- 2,155 conversations
- 17,058 turns
- 10,174 thought annotations

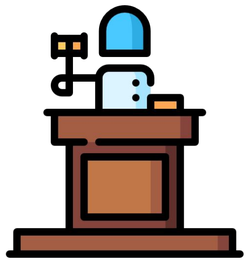

LLM Judge



Validation with Human in the Loop (HITL)

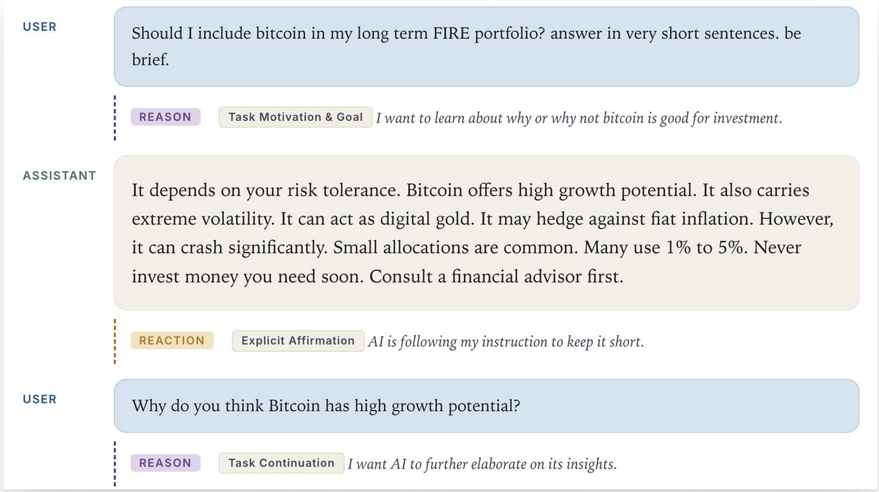

Shah, Chirag, et al. "Using large language models to generate, validate, and apply user intent taxonomies." ACM Transactions on the Web 19.3 (2025): 1-29.
Jin, Chuanyang, et al. "Thoughttrace: Understanding user thoughts in real-world llm interactions." arXiv preprint arXiv:2605.20087 (2026).

Pilot Test

ThoughtTrace

- 174 users
- **200 conversations**
- 602 turns (prompt + response)
- 1,184 thought annotations
  - 607 reasons
  - 577 reactions

- GPT-6.2-Sol
- GPT-6-Luna
- GPT-5.4

- Task ir
- Conce
- The ha
  or spe
- The re
- Tips th
  (e.g., a

s
nts
y topics
ks
xonomy
y, etc.)

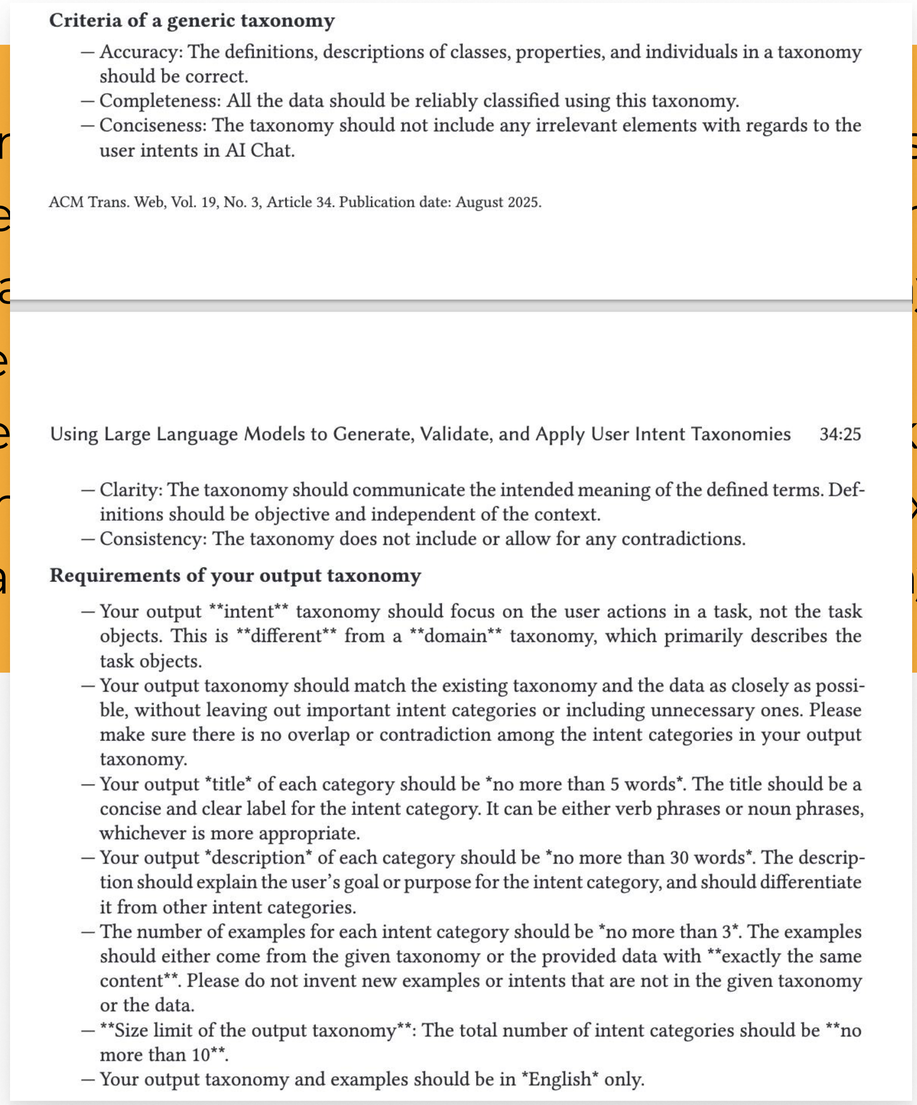

Criteria of a generic taxonomy

— Accuracy: The definitions, descriptions of classes, properties, and individuals in a taxonomy should be correct.
— Completeness: All the data should be reliably classified using this taxonomy.
— Conciseness: The taxonomy should not include any irrelevant elements with regards to the user intents in AI Chat.

ACM Trans. Web, Vol. 19, No. 3, Article 34. Publication date: August 2025.

Using Large Language Models to Generate, Validate, and Apply User Intent Taxonomies 34:25

— Clarity: The taxonomy should communicate the intended meaning of the defined terms. Definitions should be objective and independent of the context.
— Consistency: The taxonomy does not include or allow for any contradictions.

Requirements of your output taxonomy

— Your output **intent** taxonomy should focus on the user actions in a task, not the task objects. This is **different** from a **domain** taxonomy, which primarily describes the task objects.
— Your output taxonomy should match the existing taxonomy and the data as closely as possible, without leaving out important intent categories or including unnecessary ones. Please make sure there is no overlap or contradiction among the intent categories in your output taxonomy.
— Your output *title* of each category should be *no more than 5 words*. The title should be a concise and clear label for the intent category. It can be either verb phrases or noun phrases, whichever is more appropriate.
— Your output *description* of each category should be *no more than 30 words*. The description should explain the user's goal or purpose for the intent category, and should differentiate it from other intent categories.
— The number of examples for each intent category should be *no more than 3*. The examples should either come from the given taxonomy or the provided data with **exactly the same content**. Please do not invent new examples or intents that are not in the given taxonomy or the data.
— **Size limit of the output taxonomy**: The total number of intent categories should be **no more than 10**.
— Your output taxonomy and examples should be in *English* only.

Shah, Chirag, et al. "Using large language models to generate, validate, and apply user intent taxonomies." ACM Transactions on the Web 19.3 (2025): 1-29.

Taxonomy from Pilot Test

Cognitive Intent	Description	Examples
Coverage, Depth, and Completion	Get missing components, broader coverage, or a fuller explanation, rather than a more precise value or a different presentation.	“In addition, please explain the causes of independence in more detail.”
Language and Comprehensibility	Get an answer in a preferred language or at a level suited to the user's knowledge, without necessarily changing its scope or length.	“can u explain chords to a beginner in an easy way”; “Can we chat in English?”
Tone, Viewpoint, and Support	Get a preferred voice, perspective, or emotional stance, rather than an evaluation of the subject's merits.	“seeking some support”.
Assessment and Transparent Reasoning	Get a judgment based on criteria or trade-offs, or understand the reasoning behind it, rather than identify a value or option on one specified measure.	“I want a personalised ranking rather than a generic list.”
Specificity and Precision	Get a concrete value, estimate, identification, or comparison on a specified measure, without necessarily asking for verification.	“Finally, which option would be the cheapest?”; “Also, could you switch the currency from dollars to Polish złoty (PLN) and give me the approximate prices for the ingredients in Poland?”
Practical Guidance and Fit	Get actionable guidance that is feasible within personal constraints or routines. The focus is whether the guidance can be used, not whether it updates an earlier answer.	“Can you help me put together a efficient three day high protein meal prep plan that’s budget friendly and takes less than an hour to cook, making sure the meals are easy to reheat or eat cold during my work shifts?”
Presentation, Selection, and Disclosure	Control the answer's organization, length, or disclosure of selected details, rather than the depth of its explanation.	“I honestly do not know what to do and your messages are lengthy”; “Sure give me a spoiler free idea”.
Novelty, Interest, and Inspiration	Get unfamiliar, surprising, or thought-provoking ideas because freshness matters, rather than simply getting more detail or guidance that is easier to follow.	“But this tour program is already what I have done before, is there something new to try?”
Reliability, Currency, and Typicality	Get claims checked, sourced, or updated, or learn what a relevant group typically does or recommends, rather than seek the assistant's own judgment.	“Something is wrong with your math, can you double check it?”; “Too many materials, I don't know which ones to choose. Which ones do learners recommend?”
Context Elicitation and Grounding	Get an answer that asks for needed context, uses explicitly referenced information, or updates prior advice in light of a new or corrected premise.	“Can you give me a simple meal plan to go with this gym routine?”; “Could you tweak the plan a bit? I don't eat bread at all, so please swap out the wraps and sandwiches for something else.”

Research Questions

Understanding
Cognitive Search Intents



RQ 1: What **types of cognitive intents** (CIs) people have in conversational information seeking with generative AI?

RQ 2: When people have cognitive intents, **do they verbalize** them in their prompts?

RQ 3: **How people verbalize** their cognitive intents?

Step 2: Annotating Logs

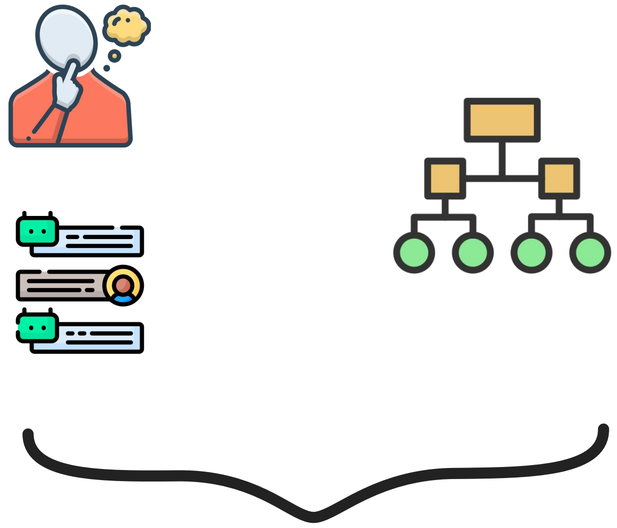

LLM-based Annotation with HITL

- CSIs in prompts only (verbalized)
- CSIs in both prompt and reasons (verbalized)
- CSIs in reasons only (not verbalized)
- CSIs in neither (no CSIs)
- Annotated CSIs as “Others” (part of taxonomy validation)

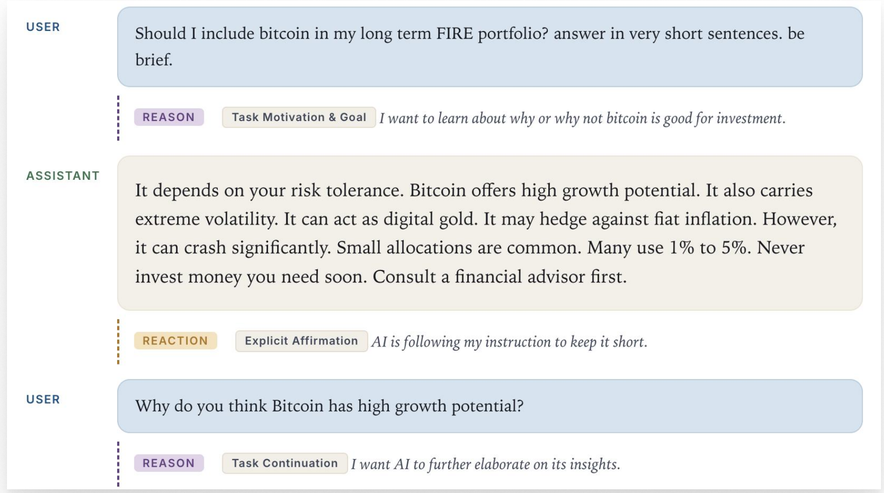

Shah, Chirag, et al. "Using large language models to generate, validate, and apply user intent taxonomies." ACM Transactions on the Web 19.3 (2025): 1-29.
Jin, Chuanyang, et al. "Thoughttrace: Understanding user thoughts in real-world llm interactions." arXiv preprint arXiv:2605.20087 (2026).

Research Questions

Understanding
Cognitive Search Intents



RQ 1: What **types of cognitive intents** (CIs) people have in conversational information seeking with generative AI?

RQ 2: When people have cognitive intents, **do they verbalize** them in their prompts?

RQ 3: **How people verbalize** their cognitive intents?

Step 3: Log Analysis


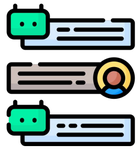
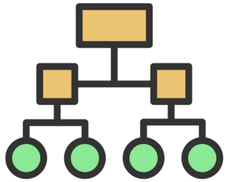

- Intents in prompts only (verbalized)
- Intents in both prompt and reasons (verbalized)

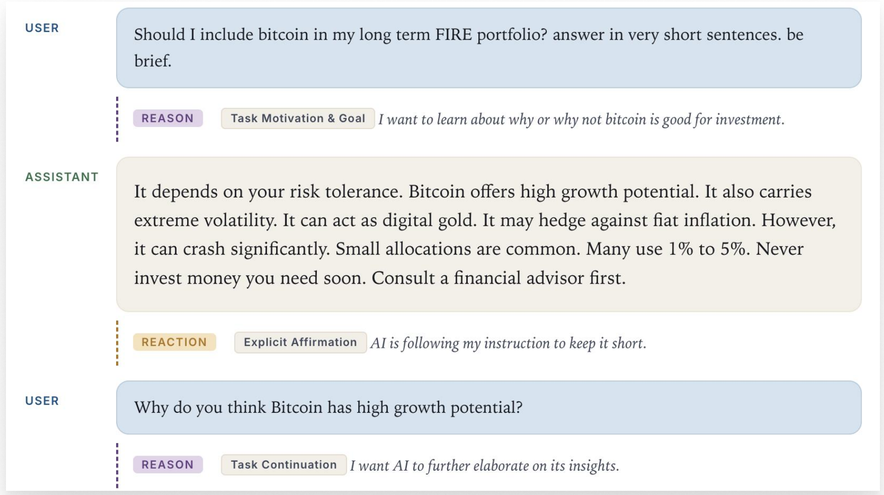

- Qualitative thematic analysis on sampled conversations
- Quantitative analysis

Understanding Cognitive Intents in Conversational Information Seeking

Shuoqi Sun

shuoqi.sun@student.rmit.edu.au

Shuoqi's Personal Website

Shuoqi's Linkdin

RMIT UNIVERSITY

大学共同利用機関法人 情報・システム研究機構 国立情報学研究所 National Institute of Informatics

ADM + S ARC Centre of Excellence for Automated Decision-Making and Society

CHAI