

Novelty Assessment Report

Paper: Shoot First, Ask Questions Later? Building Rational Agents that Explore and Act Like People

PDF URL: <https://openreview.net/pdf?id=EQhUvWH78U>

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2025-12-29

Abstract

Many high-stakes applications of AI require forming data-driven hypotheses and making targeted guesses; e.g., in scientific and diagnostic settings. Given limited resources, to what extent do agents based on language models (LMs) act rationally? We develop methods to benchmark and enhance agentic information-seeking, drawing on insights from human behavior. First, we introduce a strategic decision-oriented dialogue task called Collaborative Battleship, in which a partially-informed Captain must balance exploration (asking questions) and action (taking shots), while a fully-informed Spotter must provide accurate answers under an information bottleneck. Compared to human players (N=42), we find that LM agents struggle to ground answers in context, generate informative questions, and select high-value actions. Next, to address these gaps, we develop novel Monte Carlo inference strategies for LMs based on principles from Bayesian Experimental Design (BED). For Spotter agents, our approach boosts accuracy by up to 14.7% absolute over LM-only baselines; for Captain agents, it raises expected information gain (EIG) by up to 0.227 bits (94.2% of the achievable noise ceiling). Combined, these components yield sharper targeting (+0.303-0.374 F1), and enable weaker LMs, such as Llama-4-Scout, to outperform both humans (8% → 82% win rate) and frontier models (0% → 67% win rate vs. GPT-5) at ≈1% of GPT-5's cost. We replicate these findings on Guess Who? where our methods significantly boost accuracy (+28.3-42.4 p.p.), demonstrating their general applicability for building rational information-seeking agents.

Disclaimer

This report is **AI-GENERATED** using Large Language Models and WisPaper (a scholar search engine). It analyzes academic papers' tasks and contributions against retrieved prior work. While this system identifies **POTENTIAL** overlaps and novel directions, **ITS COVERAGE IS NOT EXHAUSTIVE AND JUDGMENTS ARE APPROXIMATE**. These results are intended to assist human reviewers and **SHOULD NOT** be relied upon as a definitive verdict on novelty.

Note that some papers exist in multiple, slightly different versions (e.g., with different titles or URLs). The system may retrieve several versions of the same underlying work. The current automated pipeline does not reliably align or distinguish these cases, so human reviewers will need to disambiguate them manually.

If you have any questions, please contact: mingzhang23@m.fudan.edu.cn

Core Task Landscape

This paper addresses: **rational information-seeking through question-asking and strategic action**

A total of **50 papers** were analyzed and organized into a taxonomy with **27 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Theoretical Foundations and Cognitive Mechanisms**
- **Developmental and Individual Differences**
- **Strategic Communication and Linguistic Approaches**
- **Human Decision-Making Under Uncertainty**
- **Artificial Intelligence and Computational Agents**
- **Robotic and Autonomous Systems**
- **Domain-Specific Applications and Methodologies**
- **Metacognitive and Reflective Processes**

Complete Taxonomy Tree

- rational information-seeking through question-asking and strategic action Survey Taxonomy
- Theoretical Foundations and Cognitive Mechanisms
 - Normative Models and Computational Frameworks (4 papers)
 - [1] What makes a good query? Prospects for a comprehensive theory of human information acquisition (Björn Meder, 2022) [View paper](#)
 - [8] A unified decision-theoretic model for information gathering and communication planning (Jennifer Renoux, 2020) [View paper](#)
 - [13] Computational mechanisms of curiosity and goal-directed exploration (P. Schwartenbeck, 2019) [View paper](#)
 - [48] Optimal sequential decision making with probabilistic digital twins (C. Agrell, 2023) [View paper](#)
 - Goal-Directed Exploration and Cognitive Maps (2 papers)
 - [21] Aligning Learning and Endogenous Decision-Making (Harsha, 2025) [View paper](#)
 - [24] Cognitive maps for goal-directed decision making (Boorman, 2025) [View paper](#)
- Developmental and Individual Differences
 - Developmental Trajectories in Children and Adolescents (2 papers)
 - [3] Causal information-seeking strategies change across childhood and adolescence (Kate Nussenbaum, 2020) [View paper](#)
 - [23] Toddlers strategically adapt their information search. (F. Poli, 2024) [View paper](#)
 - Experience-Based Differences in Strategy Selection (2 papers)
 - [14] Goal-oriented training affects decision-making processes in virtual and simulated fire and rescue environments. (Sabrina R. Cohen-Hatton, 2015) [View paper](#)
 - [35] Influence of experience on information-gathering strategies. (Susan S. Kirschenbaum, 1992) [View paper](#)
- Strategic Communication and Linguistic Approaches
 - Rhetorical and Argumentative Question Functions (2 papers)
 - [10] Strategic Questioning (Ilie, 2022) [View paper](#)
 - [33] The argumentative functions of strategic questions in Portuguese parliamentary speeches (Fabrizio Macagno, 2025) [View paper](#)

- Task-Oriented Dialogue and Clarification Strategies (2 papers)
- [29] : How to learn from procedural How-to questions (G Dagan, 2025) [View paper](#)
- [50] Towards Asking Clarification Questions for Information Seeking on Task-Oriented Dialogues (Feng Yue, 2023) [View paper](#)
- Strategic Questioning in Professional and Institutional Settings (2 papers)
- [16] Optimal Questionnaires for Screening of Strategic Agents (Anuj S. Vora, 2020) [View paper](#)
- [26] Asking the Right Questions or Asking Questions Right: The Effect of Strategic Questioning on Deception in Earnings-Conference-Calls (Julia Haag, 2019) [View paper](#)
- Human Decision-Making Under Uncertainty
 - Consumer and Financial Decision-Making (2 papers)
 - [5] Information search and decision-making in credit acquisition (Felipe Alves, 2024) [View paper](#)
 - [9] How Behavioural Biases Influence Rational Decision-Making in Life Insurance Purchases Decision: A Mediating Role of Information Gathering (N. Ramanjaneyalu, 2025) [View paper](#)
 - Environmental Scanning and Organizational Information Gathering (2 papers)
 - [17] Scanning actions and environmental dynamism: Gathering information for strategic decision making (JR Hough, 2004) [View paper](#)
 - [45] Toward a goal-oriented, business intelligence decision-making framework (A. Pourshahid, 2011) [View paper](#)
 - High-Stakes and Crisis Decision-Making (2 papers)
 - [34] Information gathering and decisionmaking under stress (Elliot E. Entin, 1990) [View paper](#)
 - [49] Crisis decision making: Lessons from a SWAT team (Kevin Cyr, 2022) [View paper](#)
 - Behavioral Biases and Bounded Rationality (3 papers)
 - [27] Information seeking in initial stages of relational development (WA Afifi, 2008) [View paper](#)
 - [36] The addition of anticipated regret to attitudinally based, goal-directed models of information search behaviours under conditions of uncertainty and risk (Steven A. Taylor, 2007) [View paper](#)
 - [40] Information Regulation in Work-life: Applying the Comprehensive Model of Information Seeking to Organizational Networks (Johnson, 2015) [View paper](#)
 - Internal Versus External Information Access (1 papers)
 - [47] External and internal information gathering during decision making (Yoni, 2023) [View paper](#)
 - Search and Information Gathering During Decision Making (2 papers)
 - [32] Search during decision making (Nutt, 2005) [View paper](#)
 - [38] Information gathering and motivated decision making in health and disease (Attaallah, 2022) [View paper](#)
- Artificial Intelligence and Computational Agents
 - Language Model Agents for Strategic Dialogue ★ (2 papers)
 - [0] Shoot First, Ask Questions Later? Building Rational Agents that Explore and Act Like People (Anon et al., 2026) [View paper](#)
 - [2] Doctor-R1: Mastering Clinical Inquiry with Experiential Agentic Reinforcement Learning (Liu Kai-ming, 2025) [View paper](#)
 - Multi-Agent and Distributed Information Systems (2 papers)
 - [31] Distributed intelligent agents (Katia Sycara, 1996) [View paper](#)
 - [37] BIG: An agent for resource-bounded information gathering and decision making (Victor Lesser, 2000) [View paper](#)
 - Question-Answering and Strategy Selection Systems (2 papers)
 - [6] Strategy selection in question answering (Lynne M. Reder, 1987) [View paper](#)
 - [7] Jury-trial story construction and analysis using goal-directed answer set programming (Zesheng Xu, 2023) [View paper](#)
 - Human-AI Collaboration and Decision Support (1 papers)
 - [22] Technology-supported self-triage decision making (Marvin Kopka, 2025) [View paper](#)
- Robotic and Autonomous Systems
 - Gaussian Process-Based Information Gathering (2 papers)
 - [11] POAM: Probabilistic Online Attentive Mapping for Efficient Robotic Information Gathering (Chen, 2024) [View paper](#)
 - [12] MexGen: An Effective and Efficient Information Gain Approximation for Information Gathering Path Planning (Joshua Chesser, 2024) [View paper](#)
 - Communication-Aware Robotic Information Gathering (1 papers)
 - [44] CARING: Communication-Aware Robotic Information Gathering (Moon, 2021) [View paper](#)
- Domain-Specific Applications and Methodologies
 - Educational and Learning Contexts (3 papers)
 - [15] Learning strategies for acquiring useful knowledge (Lorna Idol, 2013) [View paper](#)
 - [18] Implementation of Information Gathering Method in Academic Information System (SIKAD) (Fifit Alfiah, 2025) [View paper](#)
 - [41] The Role of Information-Gathering on Students' Satisfaction in Engineering Majors (Haleh Brotherton, 2024) [View paper](#)
 - Healthcare and Medical Decision-Making (1 papers)
 - [30] Societal perceptions of COVID-19 vaccine hesitancy: strategic communicative action in problem solving (Linda Dam, 2025) [View paper](#)
 - Engineering and Technical Systems (2 papers)
 - [25] Value of Information Analysis for rationalising information gathering in building energy analysis (Langtry, 2023) [View paper](#)
 - [42] Decisions for information or information for decisions? Optimizing information gathering in decision-intensive processes (Simon Voorberg, 2021) [View paper](#)
 - Legal and Investigative Contexts (2 papers)
 - [39] On Plaintiffs' Strategic Information Acquisition and Disclosure during Discovery (Mehdi Ayouni, 2024) [View paper](#)
 - [46] Appellate Courts And The Adversary System: Information-gathering And Appellate Decision-making. (Marvell, 2023) [View paper](#)
 - Design and User Experience Research (1 papers)
 - [20] User-centered design, activity-centered design, and goal-directed design: a review of three methods for designing web applications (Williams, 2009) [View paper](#)
 - Qualitative and Narrative Research Methods (1 papers)
 - [19] Decision-making in narrative inquiry in education: ethical gathering, transcription, validation and analysis of lived experiences (André Freitas, 2025) [View paper](#)
 - Biological and Ecological Systems (1 papers)
 - [4] Information gathering and decision making about resource value in animal contests (Gareth Arnott, 2008) [View paper](#)

- Metacognitive and Reflective Processes (2 papers)
 - [28] If You Ask Yourself a Question, Seek the Answer! (Meylan, 2025) [View paper](#)
 - [43] Question storming: The power of Questions-A method to increase options for action in uncertainty. (S. Mennin, 2023) [View paper](#)

Narrative

Core task: rational information-seeking through question-asking and strategic action. This field examines how agents—whether human, computational, or robotic—decide what questions to ask and which actions to take in order to reduce uncertainty and achieve goals efficiently. The taxonomy reflects a broad interdisciplinary landscape. Theoretical Foundations and Cognitive Mechanisms explore normative models and psychological underpinnings of curiosity and inquiry, often drawing on work like Good Query Theory[4] and Causal Information Seeking[6]. Developmental and Individual Differences investigate how question-asking emerges across the lifespan, including studies on toddlers' information search strategies. Strategic Communication and Linguistic Approaches focus on the pragmatics of dialogue, examining how speakers frame questions to elicit useful responses or even deceive, as seen in Strategic Questioning Deception[29]. Human Decision-Making Under Uncertainty addresses how people gather and weigh information before committing to choices, with contributions like Anticipated Regret Search[38] and Motivated Decision Making[40]. Meanwhile, Artificial Intelligence and Computational Agents and Robotic and Autonomous Systems branches cover algorithmic methods for active learning, dialogue management, and sensor planning, exemplified by works such as BIG Agent[39] and POAM Robotic Mapping[14]. Domain-Specific Applications span areas from medical triage to legal discovery, while Metacognitive and Reflective Processes consider self-monitoring and adaptive inquiry.

Within the AI and computational agents branch, a particularly active line of work centers on language model agents that engage in strategic dialogue to gather information or negotiate outcomes. These systems must balance exploration—asking clarifying questions—with exploitation of known facts, often under constraints like limited interaction turns or noisy environments. The original paper `json[0]` sits squarely in this cluster, focusing on Language Model Agents for Strategic Dialogue. It shares thematic ground with `aliases[2]`, which also addresses dialogue-based information-seeking, and contrasts with earlier heuristic approaches like `Shoot First Ask Later[3]`, which prioritized rapid action over careful inquiry. Compared to domain-specific agents such as `Doctor R1[5]`, which targets medical reasoning, `json[0]` appears to emphasize more general-purpose strategic communication, exploring how agents can dynamically adapt their questioning strategies across varied conversational contexts. This positioning highlights ongoing tensions between task-specific fine-tuning and broadly transferable dialogue policies.

Related Works in Same Category

The following **1 sibling papers** share the same taxonomy leaf node with the original paper:

1. Doctor-R1: Mastering Clinical Inquiry with Experiential Agentic Reinforcement Learning

Authors: Liu Kai-ming, Yunghwei Lai, Wang, Ziyue, Kai Liu, et al. (11 authors total) | **Year/Venue:** 2025 | **URL:** [View paper](#)

Abstract

The professionalism of a human doctor in outpatient service depends on two core abilities: the ability to make accurate medical decisions and the medical consultation skill to conduct strategic, empathetic patient inquiry. Existing Large Language Models (LLMs) have achieved remarkable accuracy on medical decision-making benchmarks. However, they often lack the ability to conduct the strategic and empathetic consultation, which is essential for real-world clinical scenarios. To address this gap, ...

Relationship Analysis

Both papers belong to the Language Model Agents for Strategic Dialogue category, focusing on LM-based agents that balance exploration and action through question-asking. While the original paper (Battleship) develops general Bayesian inference strategies for rational information-seeking in game-based environments with explicit hypothesis spaces, Doctor-R1 specializes in clinical inquiry, using experiential reinforcement learning to train medical consultation agents that ask strategic diagnostic questions. The key difference is that the original paper emphasizes theoretical Bayesian experimental design principles applicable across domains, whereas Doctor-R1 applies domain-specific RL with multi-agent simulation and two-tiered rewards tailored to medical consultation scenarios.

Contributions Analysis

Overall novelty summary. The paper introduces a strategic dialogue task (Collaborative Battleship) and Monte Carlo inference methods for language model agents that balance question-asking and action-taking under uncertainty. It sits within the 'Language Model Agents for Strategic Dialogue' leaf, which contains only two papers total (including this one). This is a relatively sparse research direction within the broader taxonomy of 50 papers across 27 leaf nodes, suggesting the work addresses an emerging rather than saturated area of inquiry.

The taxonomy reveals that strategic dialogue agents occupy a distinct niche within the AI and Computational Agents branch, separated from multi-agent systems, question-answering strategy selection, and human-AI collaboration. Neighboring leaves focus on distributed architectures or strategy selection without the emphasis on balancing exploration-exploitation trade-offs in dialogue. The paper's use of Bayesian Experimental Design principles connects it conceptually to the Theoretical Foundations branch (normative models), though it remains firmly an applied AI contribution rather than a cognitive modeling effort.

Among 24 candidates examined across three contributions, the Collaborative Battleship task and evaluation framework showed no clear refutation (10 candidates each, zero refutable). However, the Monte Carlo inference strategies based on Bayesian Experimental Design encountered one refutable candidate among four examined, indicating some prior work in this methodological space. The limited search scope (top-K semantic matches plus citation expansion) means these statistics reflect a targeted rather than exhaustive literature review, particularly for the BED-based methods.

Given the sparse taxonomy leaf and the modest search scale, the work appears to occupy relatively novel ground in applying BED principles to LM-based strategic dialogue. The task design and human-agent comparison framework show stronger novelty signals than the inference methods, where at least one overlapping prior approach was identified. The analysis covers semantic neighbors and citations but does not claim comprehensive coverage of all related work in active learning or dialogue systems.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: Collaborative Battleship task and BATTLESHIP QA dataset

Description: The authors develop a two-player dialogue and decision-making task extending the classic Battleship game, where players ask natural language questions to gain information about hidden ships. They collect 126 full human-human game trajectories (N=42 participants) including dialogue and actions, creating the BATTLESHIP QA dataset with 931 gold yes/no questions for evaluating grounded answering and strategic gameplay.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Human-Agent Coordination in Games under Incomplete Information via Multi-Step Intent

URL: [View paper](#)

Brief Assessment

Human-Agent Coordination in Games under Incomplete Information via Multi-Step Intent[63] focuses on turn-based cooperative games with multi-step intent and online planning algorithms (intentMCTS), not on dialogue-based question-asking tasks or dataset collection for evaluating grounded answering in strategic gameplay.

2. Manipulative underspecification

URL: [View paper](#)

Brief Assessment

Manipulative underspecification (Bayesian 21st Century",[55]) focuses on strategic dialogue and underspecification in political/conversational contexts, not on developing game-based datasets for evaluating AI agents' information-seeking abilities in grid-based environments.

3. Autonomous agents for collaborative task under information asymmetry

URL: [View paper](#)

Brief Assessment

Autonomous agents for collaborative task under information asymmetry (Reverse Annealed Monte Carlo,[56]) focuses on multi-agent collaboration under information asymmetry in social networks, not on strategic dialogue tasks with exploration-action decisions like Battleship. The candidate addresses different collaboration paradigms and does not present prior work on Battleship-style games with natural language questions.

4. Communication and Verification in LLM Agents towards Collaboration under Information Asymmetry

URL: [View paper](#)

Brief Assessment

Game Theoretic Steering[64] focuses on communication and verification in LLM agents under information asymmetry using Einstein puzzles adapted to a tabletop game, not the Battleship game framework with natural language questions and dialogue trajectories.

5. Conversation as action under uncertainty

URL: [View paper](#)

Brief Assessment

Conversation as action under uncertainty (Multi Turn Puzzles,[61]) focuses on general spoken dialog systems for tasks like PowerPoint navigation and receptionist duties, not strategic board game dialogue with information asymmetry and exploration-action tradeoffs.

6. Prompt, information, and game theory: A strategic guide to existence

URL: [View paper](#)

Brief Assessment

Autonomous Agents Asymmetry[58] discusses information asymmetry in abstract theoretical terms (ethical phase divergence, curvature of action). It does not present empirical dialogue tasks, game-based datasets, or human behavioral studies comparable to the BATTLESHIP QA contribution.

7. Multi-Turn Puzzles: Evaluating Interactive Reasoning and Strategic Dialogue in LLMs

URL: [View paper](#)

Brief Assessment

Multi-Turn Puzzles[59] focuses on multi-turn dialogue and reasoning tasks with deterministic scoring, not specifically on strategic dialogue with exploration-action decisions under information asymmetry like Collaborative Battleship.

8. Steering language models with game-theoretic solvers

URL: [View paper](#)

Brief Assessment

Human Agent Cooperation,[62] focuses on steering language models using game-theoretic solvers in negotiation domains (scheduling, trading, debate), not on developing question-asking tasks with information asymmetry like Collaborative Battleship.

9. Human-agent cooperation in games under incomplete information through natural language communication

URL: [View paper](#)

Brief Assessment

Human-agent cooperation in games under incomplete information through natural language communication (Prompt Game Theory,[60]) focuses on a shared-control game where players alternate turns controlling a single token, not a two-player dialogue task where players ask natural language questions about hidden ships. The candidate's gnomes at night testbed involves cooperative maze navigation with incomplete information about transition functions, which is fundamentally different from the ORIGINAL paper's Battleship variant with question-asking and grounded answering.

10. The Traitors: Deception and Trust in Multi-Agent Language Model Simulations

URL: [View paper](#)

Brief Assessment

The Traitors: Deception and Trust in Multi-Agent Language Model Simulations[57] focuses on social deception games with asymmetric information where agents deceive or detect deception through dialogue, not on strategic question-asking tasks with exploration-action tradeoffs like Battleship.

Contribution 2: Monte Carlo inference strategies based on Bayesian Experimental Design

Description: The authors formalize three Bayesian-inspired inference-time strategies that leverage sequential Monte Carlo approximation: QBayes for asking questions that maximize expected information gain, MBayes for selecting moves that maximize hit probability, and DBayes for deciding between asking questions or taking actions via one-step lookahead. These strategies enable weaker language models to achieve superhuman performance while maintaining significant cost savings.

This contribution was assessed against **4 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Doing experiments and revising rules with natural language and probabilistic reasoning

URL: [View paper](#)

Brief Assessment

Doing experiments and revising rules with natural language and probabilistic reasoning[52] focuses on Zendo-style concept learning tasks using Sequential Monte Carlo Samplers (SMC-S) for hypothesis revision. The original paper applies Monte Carlo methods to interactive dialogue and decision-making in Battleship, a fundamentally different domain requiring real-time question-answering and strategic action selection rather than batch concept induction.

2. Bayesian Computations in the 21st Century

URL: [View paper](#)

Brief Assessment

BED LLM²[53] appears to be a general survey/review paper on Bayesian computations mentioning sequential Monte Carlo and Bayesian Experimental Design as established techniques. The candidate does not present these as novel contributions for language model inference, whereas the original paper specifically develops three novel strategies (QBayes, MBayes, DBayes) that apply BED principles to LM-based agents in interactive decision-making tasks.

3. BED-LLM: Intelligent Information Gathering with LLMs and Bayesian Experimental Design

URL: [View paper](#)

Prior Art Analysis

BED-LLM[51] demonstrates prior work that uses Bayesian Experimental Design principles with sequential inference for information-seeking tasks with LLMs. The candidate paper explicitly describes using BED to iteratively choose questions that maximize expected information gain (EIG), which is the same core principle claimed as novel in the original paper. Both papers formalize EIG-based question selection strategies and apply them to similar information-gathering scenarios (20 questions game vs. Battleship). The candidate's approach of 'iteratively choosing questions or queries that maximize the expected information gain (eig)' directly parallels the original's QBayes strategy for 'asking questions that maximize expected information gain.'

Evidence

Evidence 1 - **Rationale:** Both papers explicitly use Bayesian Experimental Design with EIG maximization as their core methodology for improving LLM information-seeking behavior. The candidate paper's BED-LLM framework predates the original's claimed novelty. - **Original:** we develop novel monte carlo inference strategies for lms based on principles from bayesian experimental design (bed). for spotter agents, our approach boosts accuracy by up to 14.7% absolute over lm-only baselines; for captain agents, it raises expected information gain (eig) by up to 0.227 bits (9... - **Candidate:** we propose a general-purpose approach for improving the ability of large language models (llms) to intelligently and adaptively gather information from a user or other external source using the framework of sequential bayesian experimental design (bed). this enables llms to act as effective multi-tu...

Evidence 2 - **Rationale:** The QBayes strategy of selecting questions that maximize EIG is directly paralleled by BED-LLM's approach of 'iteratively choosing questions or queries that maximize the expected information gain.' Both formalize EIG estimation from probabilistic models. - **Original:** asking questions to maximize information gain (qbayes). to maximize expected information gain, a simple strategy is to sample a set of candidate questions q (e.g., from an lm) and select the one with highest eig: $q^* t \in \arg \max_{q \in \mathcal{Q}} \text{eig}_\epsilon(q | x, h_{1:t})$ - **Candidate:** our approach, which we call bed-llm (bayesian experimental design with large language models), is based on iteratively choosing questions or queries that maximize the expected information gain (eig) about the task of interest given the responses gathered previously. we show how this eig can be formu...

Evidence 3 - **Rationale:** Both papers use sequential inference methods to estimate EIG in intractable spaces. The original's SMC approach and the candidate's 'updating procedure' for probabilistic models represent the same methodological contribution of using sequential inference for BED with LLMs. - **Original:** sequential monte carlo (smc) approximation. exact sums over s^x are typically intractable; we therefore maintain a weighted particle approximation $\{(s_j, w(t, j))\}_{j=1}^n$ which is updated via sequential monte carlo (smc; doucet et al., 2001) with per-turn resampling - **Candidate:** we show how this eig can be formulated (and then estimated) in a principled way using a probabilistic model derived from the llm's predictive distributions and provide detailed insights into key decisions in its construction and updating procedure.

4. Reverse-Annealed Sequential Monte Carlo for Efficient Bayesian Optimal Experiment Design

URL: [View paper](#)

Brief Assessment

Reverse-Annealed Sequential Monte Carlo for Efficient Bayesian Optimal Experiment Design[54] focuses on improving computational efficiency of evidence estimation in BOED through reverse-annealed SMC, not on developing inference-time strategies for language model agents in question-asking and decision-making tasks.

Contribution 3: Evaluation framework comparing human and agent information-seeking behavior

Description: The authors create a reusable evaluation harness that systematically compares language model agents against human behavior and idealized resource rational strategies in information-seeking tasks. The framework tests distinct agent capabilities including asking informative questions, providing grounded answers, taking strategic actions, and navigating explore/exploit tradeoffs, with demonstrated generalizability to other information-seeking games like Guess Who.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. The turn: Integration of information seeking and retrieval in context

URL: [View paper](#)

Brief Assessment

The candidate paper 'The turn: Integration of information seeking and retrieval in context' appears to focus on information retrieval and search strategies in document collections, not on systematic evaluation frameworks comparing language model agents against human behavior in interactive information-seeking tasks like the original paper's Battleship game.

2. ChatGPT's Information Seeking Strategy: Insights from the 20-Questions Game

URL: [View paper](#)

Brief Assessment

ChatGPT's Information Seeking Strategy[69] focuses on evaluating ChatGPT's question-asking strategy in the 20-questions game using feature norms and hierarchical hypothesis spaces, without systematic comparison against human behavior in interactive tasks or generalization to other information-seeking games like the original paper's Battleship and Guess Who environments.

3. Mind2Web 2: Evaluating Agentic Search with Agent-as-a-Judge

URL: [View paper](#)

Brief Assessment

Mind2Web 2: Evaluating Agentic Search with Agent-as-a-Judge (BEARCUBS Web Agents,[68]) focuses on evaluating agentic search systems on long-horizon web browsing tasks with time-varying answers, not on comparing agents against human behavior in strategic information-seeking games like the original paper's Battleship task.

4. ClarQ-LLM: A Benchmark for Models Clarifying and Requesting Information in Task-Oriented Dialog

URL: [View paper](#)

Brief Assessment

ClarQ-LLM focuses on clarification question-asking in task-oriented dialogues with provider-seeker interactions, not on comparing agents against human behavior or resource-rational strategies in information-seeking games like the original paper's Battleship framework.

5. Student translators' web-based vs. GenAI-based information-seeking behavior in translation process: A comparative study

URL: [View paper](#)

Brief Assessment

The candidate paper focuses on student translators' web-based vs. GenAI-based information-seeking behavior in translation processes, not on evaluating language model agents against human strategies in interactive information-seeking tasks like the original paper's Battleship framework.

6. BEARCUBS: A benchmark for computer-using web agents

URL: [View paper](#)

Brief Assessment

BEARCUBS[66] focuses on evaluating web agents' computer-use abilities for information-seeking on live web content, not on comparing agent behavior against human strategies or resource-rational baselines in interactive dialogue tasks like the original paper's Battleship framework.

7. DiscipLink: Unfolding Interdisciplinary Information Seeking Process via Human-AI Co-Exploration

URL: [View paper](#)

Brief Assessment

DiscipLink focuses on human-AI collaborative information seeking in interdisciplinary research contexts, not on benchmarking agent behavior against human strategies or idealized rational models in interactive tasks.

8. On the motivations to seek information from artificial intelligence agents versus humans: A risk information seeking and processing perspective

URL: [View paper](#)

Brief Assessment

The candidate paper examines human preferences for AI versus human information sources in risk communication contexts, not the development of evaluation frameworks for benchmarking agent capabilities against human strategies in interactive tasks.

9. Long-Horizon Prediction for Human-Robot Collaboration

URL: [View paper](#)

Brief Assessment

The candidate paper (ClarQ LLM,[72]) focuses on long-horizon prediction for human-robot collaboration, not on benchmarking agentic information-seeking behavior against human strategies in interactive tasks like the original paper's Battleship framework.

10. Thought-Augmented Planning for LLM-Powered Interactive Recommender Agent

URL: [View paper](#)

Brief Assessment

The candidate paper (Thought-Augmented Planning for LLM-Powered Interactive Recommender Agent[74]) focuses on interactive recommendation systems with thought-augmented planning, not on systematic evaluation frameworks comparing human and agent information-seeking behavior in controlled tasks like the original paper's Battleship environment.

Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

References

- [0] Shoot First, Ask Questions Later? Building Rational Agents that Explore and Act Like People [View paper](#)
- [1] What makes a good query? Prospects for a comprehensive theory of human information acquisition [View paper](#)
- [2] Doctor-R1: Mastering Clinical Inquiry with Experiential Agentic Reinforcement Learning [View paper](#)
- [3] Causal information-seeking strategies change across childhood and adolescence [View paper](#)
- [4] Information gathering and decision making about resource value in animal contests [View paper](#)
- [5] Information search and decision-making in credit acquisition [View paper](#)
- [6] Strategy selection in question answering [View paper](#)
- [7] Jury-trial story construction and analysis using goal-directed answer set programming [View paper](#)
- [8] A unified decision-theoretic model for information gathering and communication planning [View paper](#)
- [9] How Behavioural Biases Influence Rational Decision-Making in Life Insurance Purchases Decision: A Mediating Role of Information Gathering [View paper](#)
- [10] Strategic Questioning [View paper](#)
- [11] POAM: Probabilistic Online Attentive Mapping for Efficient Robotic Information Gathering [View paper](#)
- [12] MexGen: An Effective and Efficient Information Gain Approximation for Information Gathering Path Planning [View paper](#)
- [13] Computational mechanisms of curiosity and goal-directed exploration [View paper](#)
- [14] Goal-oriented training affects decision-making processes in virtual and simulated fire and rescue environments. [View paper](#)
- [15] Learning strategies for acquiring useful knowledge [View paper](#)
- [16] Optimal Questionnaires for Screening of Strategic Agents [View paper](#)
- [17] Scanning actions and environmental dynamism: Gathering information for strategic decision making [View paper](#)
- [18] Implementation of Information Gathering Method in Academic Information System (SIKAD) [View paper](#)

- [19] Decision-making in narrative inquiry in education: ethical gathering, transcription, validation and analysis of lived experiences [View paper](#)
- [20] User-centered design, activity-centered design, and goal-directed design: a review of three methods for designing web applications [View paper](#)
- [21] Aligning Learning and Endogenous Decision-Making [View paper](#)
- [22] Technology-supported self-triage decision making [View paper](#)
- [23] Toddlers strategically adapt their information search. [View paper](#)
- [24] Cognitive maps for goal-directed decision making [View paper](#)
- [25] Value of Information Analysis for rationalising information gathering in building energy analysis [View paper](#)
- [26] Asking the Right Questions or Asking Questions Right: The Effect of Strategic Questioning on Deception in Earnings-Conference-Calls [View paper](#)
- [27] Information seeking in initial stages of relational development [View paper](#)
- [28] If You Ask Yourself a Question, Seek the Answer! [View paper](#)
- [29] : How to learn from procedural How-to questions [View paper](#)
- [30] Societal perceptions of COVID-19 vaccine hesitancy: strategic communicative action in problem solving [View paper](#)
- [31] Distributed intelligent agents [View paper](#)
- [32] Search during decision making [View paper](#)
- [33] The argumentative functions of strategic questions in Portuguese parliamentary speeches [View paper](#)
- [34] Information gathering and decisionmaking under stress [View paper](#)
- [35] Influence of experience on information-gathering strategies. [View paper](#)
- [36] The addition of anticipated regret to attitudinally based, goal-directed models of information search behaviours under conditions of uncertainty and risk [View paper](#)
- [37] BIG: An agent for resource-bounded information gathering and decision making [View paper](#)
- [38] Information gathering and motivated decision making in health and disease [View paper](#)
- [39] On Plaintiffs Strategic Information Acquisition and Disclosure during Discovery [View paper](#)
- [40] Information Regulation in WorkLife: Applying the Comprehensive Model of Information Seeking to Organizational Networks [View paper](#)
- [41] The Role of Information-Gathering on Students Satisfaction in Engineering Majors [View paper](#)
- [42] Decisions for information or information for decisions? Optimizing information gathering in decision-intensive processes [View paper](#)
- [43] Question storming: The power of Questions-A method to increase options for action in uncertainty. [View paper](#)
- [44] CARING: Communication-Aware Robotic Information Gathering [View paper](#)
- [45] Toward a goal-oriented, business intelligence decision-making framework [View paper](#)
- [46] Appellate Courts And The Adversary System: Information-gathering And Appellate Decision-making. [View paper](#)
- [47] External and internal information gathering during decision making [View paper](#)
- [48] Optimal sequential decision making with probabilistic digital twins [View paper](#)
- [49] Crisis decision making: Lessons from a SWAT team [View paper](#)
- [50] Towards Asking Clarification Questions for Information Seeking on Task-Oriented Dialogues [View paper](#)
- [51] BED-LLM: Intelligent Information Gathering with LLMs and Bayesian Experimental Design [View paper](#)
- [52] Doing experiments and revising rules with natural language and probabilistic reasoning [View paper](#)
- [53] Bayesian Computations in the 21st Century [View paper](#)
- [54] Reverse-Annealed Sequential Monte Carlo for Efficient Bayesian Optimal Experiment Design [View paper](#)
- [55] Manipulative underspecification [View paper](#)
- [56] Autonomous agents for collaborative task under information asymmetry [View paper](#)
- [57] The Traitors: Deception and Trust in Multi-Agent Language Model Simulations [View paper](#)
- [58] Prompt, information, and game theory: A strategic guide to existence [View paper](#)
- [59] Multi-Turn Puzzles: Evaluating Interactive Reasoning and Strategic Dialogue in LLMs [View paper](#)
- [60] Human-agent cooperation in games under incomplete information through natural language communication [View paper](#)
- [61] Conversation as action under uncertainty [View paper](#)
- [62] Steering language models with game-theoretic solvers [View paper](#)
- [63] Human-Agent Coordination in Games under Incomplete Information via Multi-Step Intent [View paper](#)
- [64] Communication and Verification in LLM Agents towards Collaboration under Information Asymmetry [View paper](#)
- [65] Student translators web-based vs. GenAI-based information-seeking behavior in translation process: A comparative study [View paper](#)
- [66] BEARCUBS: A benchmark for computer-using web agents [View paper](#)
- [67] DiscipLink: Unfolding Interdisciplinary Information Seeking Process via Human-AI Co-Exploration [View paper](#)
- [68] Mind2Web 2: Evaluating Agentic Search with Agent-as-a-Judge [View paper](#)
- [69] ChatGPTs Information Seeking Strategy: Insights from the 20-Questions Game [View paper](#)
- [70] ClarQ-LLM: A Benchmark for Models Clarifying and Requesting Information in Task-Oriented Dialog [View paper](#)
- [71] On the motivations to seek information from artificial intelligence agents versus humans: A risk information seeking and processing perspective [View paper](#)
- [72] Long-Horizon Prediction for Human-Robot Collaboration [View paper](#)
- [73] The turn: Integration of information seeking and retrieval in context [View paper](#)
- [74] Thought-Augmented Planning for LLM-Powered Interactive Recommender Agent [View paper](#)