

Novelty Assessment Report

Paper: Language and Experience: A Computational Model of Social Learning in Complex Tasks

PDF URL: <https://openreview.net/pdf?id=UxDu3RFuDV>

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2026-01-05

Abstract

The ability to combine linguistic guidance from others with direct experience is central to human development, enabling safe and rapid learning in new environments. How do people integrate these two sources of knowledge, and how might AI systems? We present a computational framework that models human social learning as joint probabilistic inference over structured, executable world models given sensorimotor and linguistic data. We make this possible by turning a pretrained language model into a probabilistic model of how humans share advice conditioned on their beliefs, allowing our agents both to generate advice for others and to interpret linguistic input as evidence during Bayesian inference. Using behavioral experiments and simulations across 10 video games, we show how linguistic guidance can shape exploration and accelerate learning by reducing risky interactions and speeding up key discoveries in both humans and models. We further explore how knowledge can accumulate across generations through iterated learning experiments and demonstrate successful knowledge transfer between humans and models—revealing how structured, language-compatible representations might enable human-machine collaborative learning.

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: **Integrating Linguistic Guidance and Direct Experience in Sequential Decision-Making**

A total of **47 papers** were analyzed and organized into a taxonomy with **21 categories**.

Taxonomy Overview

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

- **Language-Guided Agent Learning and Planning**
- **Experience-Based Learning and Adaptation**
- **Integrating Language and Experience**
- **Explanation and Feedback Mechanisms**
- **Evaluation and Benchmarking**
- **Domain-Specific Applications**
- **Theoretical Foundations and Formal Models**
- **Auxiliary Topics**

Complete Taxonomy Tree

- Integrating Linguistic Guidance and Direct Experience in Sequential Decision-Making Survey Taxonomy
- Language-Guided Agent Learning and Planning
 - Task Specification and Instruction Following (3 papers)
 - [4] Using natural language to aid task specification in sequential decision making problems (Goyal, 2022) [View paper](#)
 - [12] Gptvoicetasker: Advancing multi-step mobile task efficiency through dynamic interface exploration and learning (Minh Duc Vu, 2024) [View paper](#)
 - [27] Interactive language instructable robot learning (Mees, 2023) [View paper](#)
 - Hierarchical Planning with Language (3 papers)
 - [3] Describe, explain, plan and select: Interactive planning with large language models enables open-world multi-task agents (Wang Zihao, 2023) [View paper](#)
 - [10] Agent s: An open agentic framework that uses computers like a human (Agashe, 2024) [View paper](#)
 - [15] Learning adaptive planning representations with natural language guidance (Wong, 2023) [View paper](#)
 - Multi-Agent and Cross-Domain Task Automation (2 papers)
 - [6] EndoAgent: A Memory-Guided Reflective Agent for Intelligent Endoscopic Vision-to-Decision Reasoning (Tang Yi, 2025) [View paper](#)
 - [7] MobileSteward: Integrating Multiple App-Oriented Agents with Self-Evolution to Automate Cross-App Instructions (Liu Yuxuan, 2025) [View paper](#)
- Experience-Based Learning and Adaptation
 - Memory and Experience Retrieval (1 papers)
 - [13] RAG-Modulo: Solving Sequential Tasks using Experience, Critics, and Language Models (Jain, 2024) [View paper](#)
 - Self-Evolving Agents (1 papers)
 - [9] Se-vln: A self-evolving vision-language navigation framework based on multimodal large language models (Dong Xiang-yu, 2025) [View paper](#)
 - Reinforcement Learning with Linguistic Feedback (2 papers)
 - [2] Reflexion: Language agents with verbal reinforcement learning (Shinn, 2023) [View paper](#)
 - [22] Learning behaviors via human-delivered discrete feedback: modeling implicit feedback strategies to speed up learning (Robert Loftin, 2016) [View paper](#)

- Integrating Language and Experience
 - Bayesian Integration of Language and Reward (1 papers)
 - [16] Integration of Language and Experience via the Instructed Bandit Task (Ho, 2025) [View paper](#)
 - Computational Models of Social Learning ★ (2 papers)
 - [0] Language and Experience: A Computational Model of Social Learning in Complex Tasks (Anon et al., 2026) [View paper](#)
 - [8] Integrating Natural Language in Sequential Decision Problems (Huang, 2025) [View paper](#)
- Explanation and Feedback Mechanisms
 - Causal and Natural Language Explanations (2 papers)
 - [1] Towards reconciling usability and usefulness of policy explanations for sequential decision-making systems (Pradyumna Tambwekar, 2024) [View paper](#)
 - [19] Causal explanations for sequential decision-making in multi-agent systems (Gyevnar, 2023) [View paper](#)
 - Adaptive Explanation Generation (1 papers)
 - [24] SNAPE: A Sequential Non-Stationary Decision Process Model for Adaptive Explanation Generation (Amelie Robrecht, 2023) [View paper](#)
 - Feedback Timing and Sequencing (3 papers)
 - [28] Should you order the feedback sandwich? Efficacy of feedback sequence and timing (Amy J. Henley, 2015) [View paper](#)
 - [36] Examining the Effects of Directive Feedback Compared to Basic Feedback on a Problem Solving Task (Monica M. Garlock, 2015) [View paper](#)
 - [45] A Randomized Controlled Trial to Compare e-Feedback Versus "Standard" Face-to-Face Verbal Feedback to Improve the Acquisition of Procedural Skill. (W. Al-jundi, 2018) [View paper](#)
- Evaluation and Benchmarking (2 papers)
 - [14] MazeEval: A Benchmark for Testing Sequential Decision-Making in Language Models (Einarsson, 2025) [View paper](#)
 - [29] A Tree-Based Search Algorithm with Global Pheromone and Local Signal Guidance for Scientific Chart Reasoning (Min Zhou, 2025) [View paper](#)
- Domain-Specific Applications
 - Medical and Clinical Decision Support (1 papers)
 - [42] Multi-Step Transfer Learning in Natural Language Processing in the Health Domain (Thokozile Manaka, 2023) [View paper](#)
 - Education and Training (4 papers)
 - [11] Reward Feedback Mechanism in Virtual Reality Serious Games in Interventions for Children With Attention Deficits: Pre- and Posttest Experimental Control Group Study (Hao Fang, 2024) [View paper](#)
 - [26] Animating Learning: Effects of AI-Based 3D Visual Feedback on Tackling Skill Acquisition, Motivation, and Performance in Youth Soccer (Mohamed Sami Bouzid, 2025) [View paper](#)
 - [40] Augmented Targeted Training to Improve Balance in Diabetic Neuropathy Patients (Thenmozhi M, 2024) [View paper](#)
 - [47] INTELLIGENT DESIGN: BUILDING IMMERSIVE CASE STUDIES AND SIMULATIONS WITH AI (J ROYE, n.d.) [View paper](#)
 - Human-Computer Interaction and Assistive Systems (1 papers)
 - [30] An empirical study on the effect of a driving companion bot on anger coping behaviors (Youngjae Yoo, 2022) [View paper](#)
- Theoretical Foundations and Formal Models
 - Language-Based Decision Frameworks (2 papers)
 - [17] Sequential Language-based Decisions (Bjorndahl, 2023) [View paper](#)
 - [23] On the Relationship Between Structure in Natural Language and Models of Sequential Decision Processes (Roma Patel, 2020) [View paper](#)
 - Linguistic Decision Analysis (6 papers)
 - [5] Triangular neutrosophic cognitive map for multistage sequential decision-making problems (Salah Hassan Al-subhi, 2021) [View paper](#)
 - [20] Sequential three-way group decision-making for double hierarchy hesitant fuzzy linguistic term set (Nanfeng Luo, 2024) [View paper](#)
 - [21] A sequential three-way decision-based group consensus method under probabilistic linguistic term sets (Xinru Han, 2023) [View paper](#)
 - [32] Linguistic decision analysis: steps for solving decision problems under linguistic information (F. Herrera, 2000) [View paper](#)
 - [33] A sequential selection process in group decision making with a linguistic assessment approach (F Herrera, 1995) [View paper](#)
 - [34] A sequential three-way decision model based on linguistic Z-numbers. (Mao Yi, 2025) [View paper](#)
 - Transfer Learning and Knowledge Accumulation (2 papers)
 - [18] Transfer in sequential decision making (Luketina, 2023) [View paper](#)
 - [25] Language experience changes subsequent learning (Onnis, 2013) [View paper](#)
 - Probabilistic Grammar and Sequence Learning (2 papers)
 - [31] Learning a generative probabilistic grammar of experience: A process-level model of language acquisition (Oren Kolodny, 2015) [View paper](#)
 - [44] Sequential decision strategies for machine interpretation of speech (C. Raymond, 2006) [View paper](#)
- Auxiliary Topics
 - Sequential Pattern Mining (1 papers)
 - [46] Finding Sequential Patterns from Large Sequence Data (Esmaeili Mahdi, 2010) [View paper](#)
 - Peripheral Applications and Studies (6 papers)
 - [35] AirLLM: Diffusion Policy-based Adaptive LoRA for Remote Fine-Tuning of LLM over the Air (Yang, 2025) [View paper](#)
 - [37] Reimagining business education as character formation (K Wong, 2015) [View paper](#)
 - [38] Control of student mediating processes during verbal learning and instruction (Richard C. Anderson, 1970) [View paper](#)
 - [39] The sequential and reflexive achievement of coach participation in the live TV broadcasting of football (Laurent Camus, 2023) [View paper](#)
 - [41] Rational Decision-Making Agent with Learning Internal Utility Judgment (Y Ye, n.d.) [View paper](#)
 - [43] Vicarious and direct reinforcement control of verbal behavior in an interview setting (MARLATT, 1968) [View paper](#)

Narrative

Core task: integrating linguistic guidance and direct experience in sequential decision-making. The field structure reflects a broad effort to combine symbolic instruction with trial-and-error learning across diverse problem settings. The taxonomy organizes work into several main branches: Language-Guided Agent Learning and Planning focuses on how natural language instructions shape policy search and

task execution (e.g., Describe Explain Plan[3], Natural Language Task[4]); Experience-Based Learning and Adaptation emphasizes reinforcement learning, imitation, and adaptive mechanisms that refine behavior through interaction; Integrating Language and Experience addresses computational models that explicitly merge these two modalities, including social learning frameworks; Explanation and Feedback Mechanisms explores how agents generate or consume human-readable rationales and corrective signals (e.g., Reflexion[2], Policy Explanations Usability[1]); Evaluation and Benchmarking provides testbeds and metrics; Domain-Specific Applications targets robotics, healthcare, and other specialized contexts (e.g., EndoAgent[6], MobileSteward[7]); and Theoretical Foundations and Formal Models underpins the field with decision-theoretic and linguistic formalisms (e.g., Linguistic Decision Analysis[32], Sequential Linguistic Assessment[33]).

A particularly active line of work examines how agents can leverage verbal feedback or self-generated explanations to accelerate learning, contrasting purely reward-driven methods with those that incorporate structured language (Reflexion[2], Human Discrete Feedback[22]). Another theme concerns the trade-offs between pre-specified linguistic knowledge and emergent strategies learned from environmental interaction, with some studies exploring hybrid architectures that dynamically weight instruction versus experience. The original paper, Language Experience Computational[0], sits within the Computational Models of Social Learning cluster under Integrating Language and Experience, closely aligned with Integrating Natural Language[8]. Its emphasis on social learning mechanisms distinguishes it from purely individual agent frameworks, highlighting how linguistic cues and observational experience jointly shape decision policies—a perspective that complements the more instruction-centric approaches seen in adjacent branches while addressing open questions about scalability and generalization across social contexts.

Related Works in Same Category

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

1. Integrating Natural Language in Sequential Decision Problems

Authors: S Huang | **Year/Venue:** 2025 | **URL:** [View paper](#)

Abstract

â promising avenue for integrating language into sequential decision-making systems. However, â the challenges of leveraging LLMs in sequential decision-making by evaluating existing â

Relationship Analysis

Both papers belong to the Computational Models of Social Learning category, focusing on integrating linguistic guidance with experience in sequential decision-making tasks. While the original paper presents a comprehensive Bayesian framework that models social learning through joint probabilistic inference over structured world models using both sensorimotor and linguistic data (with extensive human experiments across 10 video games), the candidate paper appears to evaluate existing approaches for leveraging LLMs in sequential decision problems without presenting a novel computational framework for social learning. The original paper's core contribution is a probabilistic model of human advice-giving and interpretation, whereas the candidate focuses on evaluating LLM integration challenges in sequential tasks.

Contributions Analysis

Overall novelty summary. The paper presents a computational framework modeling human social learning through joint probabilistic inference over structured world models, integrating linguistic advice with sensorimotor experience. It resides in the 'Computational Models of Social Learning' leaf under 'Integrating Language and Experience', which contains only two papers total. This sparse positioning suggests the work addresses a relatively underexplored niche within the broader taxonomy of 47 papers. The sibling paper focuses on integrating natural language hints with reward signals, indicating the leaf targets explicit language-experience fusion rather than purely instruction-following or purely experiential learning approaches.

The taxonomy reveals that most related work clusters in adjacent branches: 'Language-Guided Agent Learning and Planning' contains systems using instructions for task specification and hierarchical planning, while 'Experience-Based Learning and Adaptation' emphasizes memory retrieval and self-evolution mechanisms. The 'Bayesian Integration of Language and Reward' leaf sits nearby, focusing on probabilistic bandit models rather than structured world models. The paper's emphasis on social learning—modeling how humans share and interpret advice—distinguishes it from these neighboring directions, which typically address single-agent instruction following or reward-based adaptation without explicit social communication modeling.

Among 24 candidates examined across three contributions, the literature search found limited prior overlap. The core framework contribution (10 candidates examined, 0 refutable) and the language-accelerated inference contribution (10 candidates, 0 refutable) appear relatively novel within this search scope. However, the speaker model contribution (4 candidates examined, 1 refutable) shows more substantial prior work, suggesting that using language models to generate or interpret advice has precedent. The analysis explicitly notes this is a limited top-K semantic search, not exhaustive coverage, meaning the novelty assessment reflects only the examined subset of the field.

Given the sparse taxonomy leaf and limited refutation signals across most contributions, the work appears to occupy a distinct position within the examined literature. The social learning framing and structured world model approach differentiate it from instruction-following systems and bandit-based integration methods. However, the restricted search scope (24 candidates) and the presence of at least one overlapping prior work for the speaker model component suggest caution in claiming comprehensive novelty without broader literature coverage.

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

Contribution 1: Computational framework for social learning via joint inference over structured world models

Description: The authors propose a Bayesian framework that treats linguistic guidance and direct experience as complementary evidence sources for inferring executable, program-like world models. This enables agents to integrate both modes of knowledge acquisition during learning in complex tasks.

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. Language-Informed Synthesis of Rational Agent Models for Grounded Theory-of-Mind Reasoning On-The-Fly

URL: [View paper](#)

Brief Assessment

Language-Informed Rational Agents[50] focuses on social inference from multimodal inputs (language + vision) to infer agent goals and beliefs, not on learning executable world models from experience and linguistic guidance for interactive task learning.

2. VERA-V: Variational Inference Framework for Jailbreaking Vision-Language Models

URL: [View paper](#)

Brief Assessment

VERA-V[56] focuses on adversarial attacks against vision-language models through variational inference, not on social learning or world model inference from linguistic and experiential data. The domains and objectives are entirely different.

3. On the Arrow of Inference

URL: [View paper](#)

Brief Assessment

Arrow of Inference[52] focuses on a general cognitive principle (context-content uncertainty) governing inference across perception, memory, and language. It does not address social learning in interactive tasks or joint inference from linguistic and sensorimotor data in the manner proposed by the original paper.

4. NEUSIS: A Compositional Neuro-Symbolic Framework for Autonomous Perception, Reasoning, and Planning in Complex UAV Search Missions

URL: [View paper](#)

Brief Assessment

NEUSIS[53] focuses on UAV search missions using neuro-symbolic perception and planning for autonomous navigation, not on social learning or integrating linguistic guidance with experiential learning through Bayesian inference over program-like world models.

5. Generative models for sequential dynamics in active inference

URL: [View paper](#)

Brief Assessment

Generative Sequential Dynamics[51] focuses on sequential dynamics in active inference for motor control, perception, and planning using continuous and discrete generative models. It does not address joint probabilistic inference over structured world models from both sensorimotor and linguistic data for social learning tasks.

6. Modeling the Mental World for Embodied AI: A Comprehensive Review

URL: [View paper](#)

Brief Assessment

Mental World Embodied[57] focuses on reviewing mental world modeling approaches for embodied AI broadly, not on proposing a specific Bayesian framework for joint inference from linguistic and sensorimotor data in social learning contexts.

7. Inferring Place-Object Relationships by Integrating Probabilistic Logic and Multimodal Spatial Concepts

URL: [View paper](#)

Brief Assessment

Place-Object Relationships[54] focuses on spatial reasoning for robotic object search using probabilistic logic and multimodal concepts, not on social learning or integrating linguistic guidance with experiential learning in interactive tasks.

8. Development of compositionality through interactive learning of language and action of robots

URL: [View paper](#)

Brief Assessment

Compositionality Interactive Learning[49] focuses on developing linguistic compositionality through robot sensorimotor learning using predictive coding and active inference. It does not address social learning, joint inference from linguistic guidance and experience, or Bayesian inference over structured world models for interactive tasks.

9. Symbolic learning and reasoning with noisy data for probabilistic anchoring

URL: [View paper](#)

Brief Assessment

Symbolic Probabilistic Anchoring[55] focuses on robotic perceptual anchoring and object tracking using probabilistic logic, not on social learning from linguistic guidance and experience in interactive tasks.

10. A transformer-based representation-learning model with unified processing of multimodal input for clinical diagnostics

URL: [View paper](#)

Brief Assessment

Transformer Multimodal Diagnostics[48] focuses on clinical diagnostics using medical images and clinical text, not on social learning or world model inference from linguistic and sensorimotor data.

Contribution 2: Language models as probabilistic speaker models for advice interpretation and generation

Description: The framework leverages language models to approximate the probability that a human with specific beliefs would produce particular advice. This enables the model to both evaluate received advice under different world hypotheses and generate human-interpretable guidance for others.

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. Unified pragmatic models for generating and following instructions

URL: [View paper](#)

Prior Art Analysis

Unified Pragmatic Models[60] demonstrates prior work that uses language models as probabilistic speaker models for both interpreting and generating instructions. The candidate paper explicitly models speaker behavior using probabilistic inference, where a rational speaker reasons about listener interpretation to generate instructions, and a rational listener reasons about speaker intent to interpret instructions. This is directly analogous to the original paper's contribution of using language models to approximate human speaker probabilities for advice interpretation and generation.

Evidence

Evidence 1 - **Rationale:** Both papers demonstrate using the same probabilistic speaker model for both generating and interpreting instructions, establishing this as a unified framework rather than a novel contribution. - **Original:** the agent generates linguistic advice for future players by sampling from the same speaker model used to evaluate language likelihood during inference - effectively

translating its map theorymap into natural language - **Candidate:** the primary contribution of this work is to show how existing models of pragmatic reasoning can be extended to support instruction following and generation for challenging, multi-step, interactive tasks.

Evidence 2 - **Rationale:** Both papers describe using the same model bidirectionally for generation and interpretation through Bayesian theory of mind reasoning about speaker and listener beliefs. - **Original:** by using the lm both to interpret linguistic guidance and to generate it, the agent captures key aspects of bayesian theory of mind - modeling how humans communicate their beliefs and how they interpret the beliefs of others through language. - **Candidate:** a rational speaker s1 chooses a description that has a high chance of causing the listener modeled by l0 to follow the given actions... conversely, a rational listener l1 follows a description by choosing an action sequence which has high probability of having caused the speaker, modeled by s0, to p...

2. Language models are bounded pragmatic speakers

URL: [View paper](#)

Brief Assessment

Bounded Pragmatic Speakers[59] focuses on modeling language models as pragmatic speakers in communication games using Bayesian inference, not on advice interpretation in reinforcement learning contexts. The candidate addresses formal vs. functional competence in language generation, while the original paper specifically addresses social learning through advice in interactive tasks.

3. A comparison of human and GPT-4 use of probabilistic phrases in a coordination game

URL: [View paper](#)

Brief Assessment

Probabilistic Phrases Coordination[58] focuses on comparing human and GPT-4 use of probabilistic phrases in coordination games, not on using language models as speaker models for advice interpretation and generation in reinforcement learning contexts.

4. Use of probabilistic phrases in a coordination game: human versus GPT-4

URL: [View paper](#)

Brief Assessment

Probabilistic Phrases GPT[61] focuses on evaluating how GPT-4 estimates probabilities from natural language phrases (e.g., 'likely') in advice contexts, not on using language models as speaker models for Bayesian inference over world models or for generating/interpreting advice in interactive learning tasks.

Contribution 3: Language-accelerated Bayesian inference through targeted proposal distributions

Description: The authors introduce a mechanism where language models bias the proposal of game rules during inference by converting linguistic advice into probability distributions over candidate rules. This accelerates convergence by directing inference toward theories most compatible with received guidance.

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. Aligning Language Models with Human Preferences via a Bayesian Approach

URL: [View paper](#)

Brief Assessment

Aligning Bayesian Preferences[67] focuses on using language models within a Bayesian framework to model disagreement in human preferences for alignment tasks. It does not address using linguistic advice to bias proposal distributions during Bayesian inference over game rules or world models, which is the core mechanism in the original paper.

2. Ditar: Diffusion transformer autoregressive modeling for speech generation

URL: [View paper](#)

Brief Assessment

Ditar[71] focuses on speech generation using diffusion transformers and autoregressive modeling, not on Bayesian inference or language-guided proposal distributions for game rule learning.

3. Training Language Models with Language Feedback at Scale

URL: [View paper](#)

Brief Assessment

Language Feedback Scale[65] focuses on using language feedback to refine language model outputs through iterative generation and selection, not on accelerating Bayesian inference over game rules or world models through targeted proposal distributions as in the original paper.

4. Using Large Language Models to Suggest Informative Prior Distributions in Bayesian Statistics

URL: [View paper](#)

Brief Assessment

LLMs Informative Priors[64] uses LLMs to suggest hyperparameters for prior distributions in Bayesian statistics, not to bias proposal distributions during inference. The candidate focuses on eliciting prior knowledge before inference, while the original contribution describes using language models to guide the proposal mechanism during the inference process itself.

5. A Llama Sunk My Battleship! Asking Rational Questions with LLMs via Bayesian Inference

URL: [View paper](#)

Brief Assessment

Llama Battleship Bayesian[69] uses LLMs to bias question proposals in a Battleship game, not to accelerate Bayesian inference over world models through targeted proposal distributions for game rules as in the original paper. The candidate focuses on question-asking rather than world model inference acceleration.

6. Red Teaming Language Models with Language Models

URL: [View paper](#)

Brief Assessment

Red Teaming Language[62] focuses on using language models to generate adversarial test cases for detecting harmful behaviors in deployed chatbots, not on accelerating Bayesian inference through targeted proposal distributions for learning structured world models.

7. From word models to world models: Translating from natural language to the probabilistic language of thought

URL: [View paper](#)

Brief Assessment

[Final Audit Failure] The model insisted on a refutation claim but failed to provide verifiable evidence after multiple retries. Marked as cannot_refute for safety. Please manually verify the candidate text.

8. Large Language Bayes

URL: [View paper](#)

Brief Assessment

Large Language Bayes[68] focuses on using LLMs to generate formal probabilistic models and bias proposal distributions in Bayesian model averaging, not on accelerating inference in reinforcement learning contexts through linguistic advice as targeted proposals for game rules.

9. Controllable text generation for large language models: A survey

URL: [View paper](#)

Brief Assessment

Controllable Text Generation[70] focuses on text generation control methods for large language models, not on Bayesian inference acceleration for game rule discovery or social learning tasks.

10. Large Language Models to Enhance Bayesian Optimization

URL: [View paper](#)

Brief Assessment

LLMs Bayesian Optimization[63] focuses on using language models to enhance Bayesian optimization for hyperparameter tuning by conditioning on desired objective values. The original paper's contribution involves using language models to bias proposal distributions during Bayesian inference over game rules, which is a fundamentally different application domain and mechanism.

Appendix: Text Similarity Detection

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

References

- [0] Language and Experience: A Computational Model of Social Learning in Complex Tasks [View paper](#)
- [1] Towards reconciling usability and usefulness of policy explanations for sequential decision-making systems [View paper](#)
- [2] Reflexion: Language agents with verbal reinforcement learning [View paper](#)
- [3] Describe, explain, plan and select: Interactive planning with large language models enables open-world multi-task agents [View paper](#)
- [4] Using natural language to aid task specification in sequential decision making problems [View paper](#)
- [5] Triangular neotrophic cognitive map for multistage sequential decision-making problems [View paper](#)
- [6] EndoAgent: A Memory-Guided Reflective Agent for Intelligent Endoscopic Vision-to-Decision Reasoning [View paper](#)
- [7] MobileSteward: Integrating Multiple App-Oriented Agents with Self-Evolution to Automate Cross-App Instructions [View paper](#)
- [8] Integrating Natural Language in Sequential Decision Problems [View paper](#)
- [9] Se-vln: A self-evolving vision-language navigation framework based on multimodal large language models [View paper](#)
- [10] Agent s: An open agentic framework that uses computers like a human [View paper](#)
- [11] Reward Feedback Mechanism in Virtual Reality Serious Games in Interventions for Children With Attention Deficits: Pre- and Posttest Experimental Control Group Study [View paper](#)
- [12] Gptvoicetasker: Advancing multi-step mobile task efficiency through dynamic interface exploration and learning [View paper](#)
- [13] RAG-Modulo: Solving Sequential Tasks using Experience, Critics, and Language Models [View paper](#)
- [14] MazeEval: A Benchmark for Testing Sequential Decision-Making in Language Models [View paper](#)
- [15] Learning adaptive planning representations with natural language guidance [View paper](#)
- [16] Integration of Language and Experience via the Instructed Bandit Task [View paper](#)
- [17] Sequential Language-based Decisions [View paper](#)
- [18] Transfer in sequential decision making [View paper](#)
- [19] Causal explanations for sequential decision-making in multi-agent systems [View paper](#)
- [20] Sequential three-way group decision-making for double hierarchy hesitant fuzzy linguistic term set [View paper](#)
- [21] A sequential three-way decision-based group consensus method under probabilistic linguistic term sets [View paper](#)
- [22] Learning behaviors via human-delivered discrete feedback: modeling implicit feedback strategies to speed up learning [View paper](#)
- [23] On the Relationship Between Structure in Natural Language and Models of Sequential Decision Processes [View paper](#)
- [24] SNAPE: A Sequential Non-Stationary Decision Process Model for Adaptive Explanation Generation [View paper](#)
- [25] Language experience changes subsequent learning [View paper](#)
- [26] Animating Learning: Effects of AI-Based 3D Visual Feedback on Tackling Skill Acquisition, Motivation, and Performance in Youth Soccer [View paper](#)
- [27] Interactive language instructable robot learning [View paper](#)
- [28] Should you order the feedback sandwich? Efficacy of feedback sequence and timing [View paper](#)
- [29] A Tree-Based Search Algorithm with Global Pheromone and Local Signal Guidance for Scientific Chart Reasoning [View paper](#)
- [30] An empirical study on the effect of a driving companion bot on anger coping behaviors [View paper](#)
- [31] Learning a generative probabilistic grammar of experience: A process-level model of language acquisition [View paper](#)
- [32] Linguistic decision analysis: steps for solving decision problems under linguistic information [View paper](#)
- [33] A sequential selection process in group decision making with a linguistic assessment approach [View paper](#)
- [34] A sequential three-way decision model based on linguistic Z-numbers. [View paper](#)
- [35] AirLLM: Diffusion Policy-based Adaptive LoRA for Remote Fine-Tuning of LLM over the Air [View paper](#)
- [36] Examining the Effects of Directive Feedback Compared to Basic Feedback on a Problem Solving Task [View paper](#)
- [37] Reimagining business education as character formation [View paper](#)
- [38] Control of student mediating processes during verbal learning and instruction [View paper](#)

- [39] The sequential and reflexive achievement of coach participation in the live TV broadcasting of football [View paper](#)
- [40] Augmented Targeted Training to Improve Balance in Diabetic Neuropathy Patients [View paper](#)
- [41] Rational Decision-Making Agent with Learning Internal Utility Judgment [View paper](#)
- [42] Multi-Step Transfer Learning in Natural Language Processing in the Health Domain [View paper](#)
- [43] Vicarious and direct reinforcement control of verbal behavior in an interview setting [View paper](#)
- [44] Sequential decision strategies for machine interpretation of speech [View paper](#)
- [45] A Randomized Controlled Trial to Compare e-Feedback Versus "Standard" Face-to-Face Verbal Feedback to Improve the Acquisition of Procedural Skill. [View paper](#)
- [46] Finding Sequential Patterns from Large Sequence Data [View paper](#)
- [47] INTELLIGENT DESIGN: BUILDING IMMERSIVE CASE STUDIES AND SIMULATIONS WITH AI [View paper](#)
- [48] A transformer-based representation-learning model with unified processing of multimodal input for clinical diagnostics [View paper](#)
- [49] Development of compositionality through interactive learning of language and action of robots [View paper](#)
- [50] Language-Informed Synthesis of Rational Agent Models for Grounded Theory-of-Mind Reasoning On-The-Fly [View paper](#)
- [51] Generative models for sequential dynamics in active inference [View paper](#)
- [52] On the Arrow of Inference [View paper](#)
- [53] NEUSIS: A Compositional Neuro-Symbolic Framework for Autonomous Perception, Reasoning, and Planning in Complex UAV Search Missions [View paper](#)
- [54] Inferring Place-Object Relationships by Integrating Probabilistic Logic and Multimodal Spatial Concepts [View paper](#)
- [55] Symbolic learning and reasoning with noisy data for probabilistic anchoring [View paper](#)
- [56] VERA-V: Variational Inference Framework for Jailbreaking Vision-Language Models [View paper](#)
- [57] Modeling the Mental World for Embodied AI: A Comprehensive Review [View paper](#)
- [58] A comparison of human and GPT-4 use of probabilistic phrases in a coordination game [View paper](#)
- [59] Language models are bounded pragmatic speakers [View paper](#)
- [60] Unified pragmatic models for generating and following instructions [View paper](#)
- [61] Use of probabilistic phrases in a coordination game: human versus GPT-4 [View paper](#)
- [62] Red Teaming Language Models with Language Models [View paper](#)
- [63] Large Language Models to Enhance Bayesian Optimization [View paper](#)
- [64] Using Large Language Models to Suggest Informative Prior Distributions in Bayesian Statistics [View paper](#)
- [65] Training Language Models with Language Feedback at Scale [View paper](#)
- [66] From word models to world models: Translating from natural language to the probabilistic language of thought [View paper](#)
- [67] Aligning Language Models with Human Preferences via a Bayesian Approach [View paper](#)
- [68] Large Language Bayes [View paper](#)
- [69] A Llama Sunk My Battleship! Asking Rational Questions with LLMs via Bayesian Inference [View paper](#)
- [70] Controllable text generation for large language models: A survey [View paper](#)
- [71] Ditar: Diffusion transformer autoregressive modeling for speech generation [View paper](#)