Novelty Assessment Report

Paper: Unveiling the Cognitive Compass: Theory-of-Mind–Guided Multimodal Emotion Reasoning
PDF URL: https://openreview.net/pdf?id=8VSrk2CaBr
Venue: ICLR 2026 Conference Submission
Year: 2026
Report Generated: 2026-01-08

Abstract

Despite rapid progress in multimodal large language models (MLLMs), their capability for deep emotional understanding remains limited. We argue that genuine affective intelligence requires explicit modeling of Theory of Mind (ToM), the cognitive substrate from which emotions arise. To this end, we introduce HitEmotion, a ToM-grounded hierarchical benchmark that diagnoses capability breakpoints across increasing levels of cognitive depth. Second, we propose a ToM-guided reasoning chain that tracks mental states and calibrates cross-modal evidence to achieve faithful emotional reasoning. We further introduce TMPO, a reinforcement learning method that uses intermediate mental states as process-level supervision to guide and strengthen model reasoning. Extensive experiments show that HitEmotion exposes deep emotional reasoning deficits in state-of-the-art models, especially on cognitively demanding tasks. In evaluation, the ToM-guided reasoning chain and TMPO improve end-task accuracy and yield more faithful, more coherent rationales. In conclusion, our work provides the research community with a practical toolkit for evaluating and enhancing the cognition-based emotional understanding capabilities of MLLMs.

Core Task Landscape

This paper addresses: **Theory-of-Mind Guided Multimodal Emotion Reasoning**
A total of **20 papers** were analyzed and organized into a taxonomy with **10 categories**.

Taxonomy Overview

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

- **Theory-of-Mind Reasoning Frameworks and Architectures**
- **Multimodal ToM Benchmarks and Evaluation Datasets**
- **Application Domains and Task-Specific Implementations**
- **Cognitive and Clinical Perspectives on ToM**

Complete Taxonomy Tree

- Theory-of-Mind Guided Multimodal Emotion Reasoning Survey Taxonomy
- Theory-of-Mind Reasoning Frameworks and Architectures
  - Bayesian and Probabilistic ToM Reasoning (2 papers)
  - [1] Scaling multimodal theory-of-mind with weak-to-strong bayesian reasoning (C Zhang, 2025) View paper
  - [14] Overcoming Multi-step Complexity in Multimodal Theory-of-Mind Reasoning: A Scalable Bayesian Planner (Zhang Chunhui, 2025) View paper
  - Cognitive Architecture for ToM-Enhanced Emotion Processing ★ (3 papers)
  - [0] Unveiling the Cognitive Compass: Theory-of-Mind–Guided Multimodal Emotion Reasoning (Anon et al., 2026) View paper
  - [11] Consistency, uncertainty or inconsistency detection in multimodal emotion recognition (Alessia Fantini, 2023) View paper
  - [12] Modeling theory of mind in multimodal HCI (Yifan Zhu, 2024) View paper
  - Interpretability and Internal Mechanism Analysis (1 papers)
  - [15] From Black Boxes to Transparent Minds: Evaluating and Enhancing the Theory of Mind in Multimodal Large Language Models (Li Xinyang, 2025) View paper
- Multimodal ToM Benchmarks and Evaluation Datasets
  - Multi-Agent Social Interaction Benchmarks (2 papers)
  - [2] Muma-tom: Multi-modal multi-agent theory of mind (Haojun Shi, 2025) View paper
  - [6] SoMi-ToM: Evaluating Multi-Perspective Theory of Mind in Embodied Social Interactions (Fan Xianzhe, 2025) View paper
  - Video-Based ToM Question Answering (2 papers)
  - [7] MMToM-QA: Multimodal Theory of Mind Question Answering (Cao Jing, 2024) View paper
  - [8] Through the Theory of Mind's Eye: Reading Minds with Multimodal Video Large Language Models (Wang, 2024) View paper
  - Event Causality and Nonverbal Communication Datasets (2 papers)
  - [9] Echo: A visio-linguistic dataset for event causality inference via human-centric reasoning (Kan, 2023) View paper
  - [16] Mind the Motions: Benchmarking Theory-of-Mind in Everyday Body Language (Seungbeen Lee, 2025) View paper
- Application Domains and Task-Specific Implementations
  - Strategic Social Deduction and Negotiation (2 papers)
  - [3] Multimind: Enhancing werewolf agents with multimodal reasoning and theory of mind (Zhang Zheng, 2025) View paper
  - [5] Exploring theory of mind in large language models through multimodal negotiation (Nutchanon Yongsatianchot, 2024) View paper
  - Embodied Agent Decision-Making and Action Generation (1 papers)
  - [13] MindPower: Enabling Theory-of-Mind Reasoning in VLM-based Embodied Agents (Ruoxuan Zhang, 2025) View paper
  - Affective Computing and Emotion Explanation (2 papers)

- [4] Emotional Theory of Mind: Bridging Fast Visual Processing with Slow Linguistic Reasoning (Yasaman Etesam, 2024) View paper
- [17] Sarcasm explanation in multimodal dialogues (S Vaibhav, 2022) View paper

- Cognitive and Clinical Perspectives on ToM (4 papers)
  - [10] A conceptual framework of cognitive-affective theory of mind: towards a precision identification of mental disorders (Peng Zhou, 2023) View paper
  - [18] 38 Cognitive and Affective Theory of Mind in Young and Elderly Patients with Multiple Sclerosis (Maxime Montembeault, 2023) View paper
  - [19] A Multimodal Theory of Affect Diffusion. (Kim Peters, 2016) View paper
  - [20] Dynamic neural reconfiguration underpins empathy development in preschoolers: A multimodal EEG study (Ruoxi Wu, n.d.) View paper

Narrative

Core task: Theory-of-Mind guided multimodal emotion reasoning. This emerging field seeks to integrate cognitive models of mental state attribution with multimodal signal processing to enable richer emotion understanding. The taxonomy reveals four main branches: reasoning frameworks and architectures that design computational mechanisms for ToM-enhanced emotion processing; multimodal benchmarks and evaluation datasets that provide standardized testbeds for assessing ToM capabilities across vision, language, and other modalities; application domains spanning negotiation, human-computer interaction, and clinical settings; and cognitive and clinical perspectives that ground computational work in psychological theory. Representative efforts include MMToM QA[7] and Muma ToM[2] for benchmark development, Multimind Werewolf[3] for strategic reasoning in social games, and Emotional Theory Mind[4] for affective state inference. These branches collectively illustrate how ToM reasoning can be operationalized through diverse modalities and task contexts.

Recent work highlights contrasts between symbolic cognitive architectures and data-driven learning approaches, as well as trade-offs between interpretability and scalability. Some studies emphasize Bayesian or probabilistic frameworks for belief modeling, such as Weak to Strong Bayesian[1] and Bayesian Planner ToM[14], while others leverage neural architectures for end-to-end multimodal fusion. Cognitive Compass[0] sits within the cognitive architecture branch, proposing a structured framework for ToM-enhanced emotion processing that integrates multiple reasoning modules. Compared to nearby works like Consistency Uncertainty Detection[11], which focuses on uncertainty estimation in mental state inference, and Modeling ToM HCI[12], which targets interactive system design, Cognitive Compass[0] emphasizes a holistic cognitive architecture that coordinates perception, reasoning, and affective interpretation. This positioning reflects ongoing debates about whether ToM-guided emotion reasoning is best achieved through modular symbolic systems or through tightly integrated neural models.

Related Works in Same Category

The following **2 sibling papers** share the same taxonomy leaf node with the original paper:

1. Consistency, uncertainty or inconsistency detection in multimodal emotion recognition

Authors: Alessia Fantini, Giovanni Pilato, Gianpaolo Vitale | **Year/Venue**: 2023 | **URL**: View paper

Abstract

Humans exploit several sensory channels to recognize emotions and combine the information coming from the different channels into a single perception. Emotion Perception (EP) is also closely related to the Theory of Mind (ToM), which includes processes that capture socially and emotionally related inputs; furthermore, it interprets their meaning and direct responses accordingly. In this paper, we present a first step towards recognizing incoherence in emotions that exploits a three-level cogniti...

Relationship Analysis

Both papers belong to the Cognitive Architecture for ToM-Enhanced Emotion Processing category, focusing on integrating Theory of Mind principles with emotion recognition systems. The original paper presents a comprehensive benchmark (HitEmotion) and optimization framework (TMPO) for evaluating and enhancing MLLMs' ToM-guided emotion reasoning across hierarchical cognitive levels, while the candidate paper proposes a three-level cognitive architecture for detecting consistency, uncertainty, or inconsistency in multimodal emotion recognition using vector representations on Russell's circumplex model. The key difference is that the original paper focuses on large-scale evaluation and training of modern MLLMs with ToM reasoning chains, whereas the candidate paper addresses emotion coherence detection through a conceptual architecture designed for human-robot interaction scenarios.

---

2. Modeling theory of mind in multimodal HCI

Authors: Yifan Zhu, Hannah VanderHoeven, Kenneth Lai, Mariah Bradford, Christopher Tam, et al. (11 authors total) | **Year/Venue**: 2024 | **URL**: View paper

Abstract

â¦ We believe that models of belief and intent in multimodal HCI â¦ multimodal HCI with ToM also has the potential to inform research in both Affective Computing, such as automatic Emotion â¦

Relationship Analysis

Both papers belong to the Cognitive Architecture for ToM-Enhanced Emotion Processing category, focusing on integrating Theory-of-Mind reasoning with emotion understanding in multimodal contexts. The original paper develops a hierarchical benchmark (HitEmotion) and a reinforcement learning method (TMPO) to enhance MLLMs' emotional reasoning through ToM-guided chains, while the candidate paper applies Simulation Theory of Mind within an embodied HCI framework to track common ground in multimodal dialogues using Dynamic Epistemic Logic. The key difference is that the original paper targets evaluation and optimization of MLLMs for emotion reasoning tasks, whereas the candidate paper focuses on modeling shared beliefs and mental states in human-computer collaborative problem-solving interactions.

Contributions Analysis

Overall novelty summary. The paper introduces HitEmotion, a hierarchical benchmark for diagnosing emotional reasoning capabilities at increasing cognitive depths, alongside a ToM-guided reasoning chain and TMPO reinforcement learning method. It resides in the 'Cognitive Architecture for ToM-Enhanced Emotion Processing' leaf, which contains only three papers total, including this one. This leaf sits within the broader 'Theory-of-Mind Reasoning Frameworks and Architectures' branch, indicating the work addresses architectural design rather than pure benchmarking or application deployment. The sparse population of this specific leaf suggests the integration of ToM principles into cognitive architectures for emotion processing remains an emerging research direction.

The taxonomy reveals neighboring leaves focused on Bayesian probabilistic reasoning and interpretability analysis, while sibling branches address benchmark development and application domains like strategic games and embodied agents. The paper's positioning bridges multiple concerns: it contributes both a benchmark (typically housed in the evaluation branch) and architectural innovations (reasoning chains, TMPO training). This cross-cutting nature distinguishes it from purely benchmark-focused efforts like MMToM-QA or purely application-driven work in negotiation scenarios. The scope notes clarify that this leaf excludes application-specific

implementations and pure evaluation studies, positioning the work as foundational framework development with accompanying diagnostic tools.

Among the 22 candidates examined through semantic search, none clearly refute any of the three contributions. The HitEmotion benchmark examined 2 candidates with no refutations, suggesting limited prior work on hierarchical ToM-grounded emotion evaluation. The ToM-guided reasoning chain and TMPO method each examined 10 candidates with no refutations, indicating these specific technical approaches appear novel within the search scope. However, the analysis explicitly notes this reflects a limited top-K semantic search rather than exhaustive coverage, meaning the absence of refutations should be interpreted cautiously as evidence of novelty within the examined sample rather than definitive proof of field-wide originality.

Given the sparse taxonomy leaf and absence of refutations among 22 examined candidates, the work appears to occupy relatively unexplored territory at the intersection of ToM cognitive architectures and emotion reasoning. The hierarchical benchmark structure and process-level supervision via mental states represent distinctive technical choices. However, the limited search scope means potentially relevant work in adjacent areas—such as emotion explanation systems or multi-agent ToM benchmarks—may not have been fully captured, and broader literature on reinforcement learning from intermediate reasoning steps could provide additional context for assessing the TMPO contribution's novelty.

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

Contribution 1: HitEmotion: ToM-grounded hierarchical benchmark for multimodal emotion understanding

Description: The authors present HitEmotion, a benchmark that systematically organizes 24 emotion-related tasks into three hierarchical levels (Emotion Perception and Recognition, Emotion Understanding and Analysis, and Emotion Cognition and Reasoning) grounded in Theory of Mind principles. This structure enables precise measurement of model capability breakpoints at different cognitive depths.

This contribution was assessed against **2 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Affective-CoT: Decomposing Multimodal Emotion Reasoning through a Hierarchical Cognitive Workflow

URL: View paper

Brief Assessment

Affective CoT[32] focuses on a hierarchical framework for emotion reasoning through perception-reasoning decoupling, not on creating a ToM-grounded hierarchical benchmark. The candidate addresses interpretable emotion analysis methodology rather than systematic benchmark construction with Theory of Mind principles.

---

2. Directions for Computational Theory of Mind: Data, Metrics, Models

URL: View paper

Brief Assessment

Computational Theory Mind[31] discusses multimodal ToM benchmarks in general terms but does not present a hierarchical emotion understanding benchmark with three cognitive levels. The candidate focuses on broader computational theory of mind directions rather than the specific emotion-focused hierarchical structure of HitEmotion.

Contribution 2: ToM-guided reasoning chain for faithful emotional reasoning

Description: The authors develop a structured reasoning approach based on Theory of Mind that explicitly tracks mental states and integrates multimodal evidence. This method aims to shift models from superficial pattern matching to deeper mental state simulation for more faithful emotional understanding.

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. From Coarse to Nuanced: Cross-Modal Alignment of Fine-Grained Linguistic Cues and Visual Salient Regions for Dynamic Emotion Recognition

URL: View paper

Brief Assessment

Cross Modal Alignment[28] focuses on cross-modal alignment of linguistic cues and visual regions for dynamic emotion recognition in videos, not on Theory of Mind-based reasoning chains or mental state tracking for emotional understanding.

---

2. Integrating emotion dynamics in mental health: A trimodal framework combining ecological momentary assessment, physiological measurements, and speech â¦

URL: View paper

Brief Assessment

Emotion Dynamics Trimodal[25] focuses on ecological momentary assessment, physiological measurements, and speech for mental health monitoring. It does not present a Theory of Mind-based reasoning framework for multimodal emotional understanding in MLLMs.

---

3. Husformer: A Multimodal Transformer for Multimodal Human State Recognition

URL: View paper

Brief Assessment

Husformer[26] focuses on multi-modal fusion for human state recognition (emotion and cognitive load) using cross-modal transformers. It does not address Theory of Mind reasoning, mental state tracking, or structured reasoning chains for emotional understanding.

---

4. Inference-enabled tracking of acute mental stress via multi-modal wearable physiological sensing: A proof-of-concept study

URL: View paper

Brief Assessment

Mental Stress Tracking[23] focuses on physiological signal processing for binary stress classification using wearable sensors, not on Theory of Mind reasoning chains or multimodal evidence integration for emotional understanding in language models.

---

5. Behavioral and physiological signals-based deep multimodal approach for mobile emotion recognition

URL: View paper

Brief Assessment

Deep Multimodal Emotion[30] focuses on multimodal sensor fusion (facial, speech, keystroke, physiological signals) for mobile emotion recognition using attention-based LSTM, not on Theory of Mind reasoning chains or mental state tracking for emotional understanding.

---

6. Emotional Intelligence in Artificial Agents: Leveraging Deep Multimodal Big Data for Contextual Social Interaction and Adaptive Behavioral Modelling

URL: View paper

Brief Assessment

Emotional Intelligence Agents[22] focuses on user-state modeling through multimodal big data analysis for social agents, not on structured Theory of Mind reasoning chains with explicit mental state tracking for emotional understanding in MLLMs.

---

7. Multimodal mental state analysis

URL: View paper

Brief Assessment

Mental State Analysis[21] focuses on sentiment classification and mental health monitoring through multimodal integration, not on Theory of Mind-based reasoning chains for emotional understanding.

---

8. Multimodal large language models meet multimodal emotion recognition and reasoning: A survey

URL: View paper

Brief Assessment

Multimodal LLM Emotion[27] is a survey paper that reviews existing methods for emotion recognition using MLLMs. It does not propose a specific ToM-guided reasoning framework with mental state tracking and cross-modal evidence integration as described in the original contribution. The survey discusses general approaches but does not present a novel structured reasoning method based on Theory of Mind principles.

---

9. Online Learning Platform of Modern Chinese Course Based on Multimodal Emotion-Aware Adaptive Learning

URL: View paper

Brief Assessment

Emotion Aware Learning[29] focuses on adaptive learning systems that detect learners' emotional states through multimodal fusion (facial expressions, eye-tracking, interaction data) for educational content personalization. It does not address Theory of Mind reasoning chains, mental state tracking across cognitive depths, or the structured reasoning approach for emotional understanding proposed in the original paper.

---

10. Multimodal temporal context network for tracking dynamic changes in emotion

URL: View paper

Brief Assessment

Temporal Context Emotion[24] focuses on time series modeling for tracking dynamic emotional changes across temporal contexts, not on Theory of Mind-based mental state tracking or structured reasoning chains for emotional understanding.

Contribution 3: TMPO: Theory-of-Mind preference optimization method

Description: The authors propose TMPO, a novel reinforcement learning framework that leverages intermediate mental states from ToM-based reasoning chains as process-level supervision. This method combines supervised fine-tuning with group-wise reward policy optimization to transform reasoning from a general emergent ability into a domain-acquired skill.

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. Reflexion: Language agents with verbal reinforcement learning

URL: View paper

Brief Assessment

Reflexion[37] focuses on verbal reinforcement learning through self-reflection for sequential decision-making tasks, not on Theory-of-Mind-based reasoning chains or mental state supervision for multimodal emotion understanding. The candidate uses linguistic feedback from task outcomes, while the original contribution specifically leverages ToM-based mental states as process-level supervision signals.

---

2. Mental modeling of reinforcement learning agents by language models

URL: View paper

Brief Assessment

Mental Modeling Agents[39] focuses on evaluating whether LLMs can build mental models of RL agents by reasoning about agent behavior and state transitions in sequential decision-making tasks. This is fundamentally different from TMPO, which uses intermediate mental states from ToM-based reasoning chains as process-level supervision in a reinforcement learning framework for multimodal emotion reasoning.

---

3. Relational deep reinforcement learning

URL: View paper

Brief Assessment

Relational Deep RL[42] focuses on relational reasoning in reinforcement learning through structured perception and attention mechanisms for navigation and planning tasks, not on theory-of-mind reasoning chains or mental state supervision for emotional understanding.

---

4. Learning to reason without external rewards

URL: View paper

Brief Assessment

Reason Without Rewards[38] focuses on reinforcement learning from internal feedback (RLIF) using self-certainty as an intrinsic reward signal for general reasoning tasks, not on Theory-of-Mind-based mental state supervision for emotional reasoning as proposed in TMPO.

---

5. SuperRL: Reinforcement Learning with Supervision to Boost Language Model Reasoning

URL: View paper

Brief Assessment

SuperRL[36] focuses on alternating between RL and SFT for general reasoning tasks with sparse rewards, not on theory-of-mind-based mental state supervision for emotional reasoning. The technical approaches and application domains differ fundamentally.

---

6. Multiagent inverse reinforcement learning via theory of mind reasoning

URL: View paper

Brief Assessment

Multiagent Inverse RL[35] focuses on inferring reward functions in multiagent settings through theory of mind reasoning, not on using mental states as process-level supervision for reinforcement learning optimization in emotion reasoning tasks.

---

7. Learning only with images: Visual reinforcement learning with reasoning, rendering, and visual feedback

URL: View paper

Brief Assessment

Visual RL Reasoning[34] focuses on visual reasoning tasks (image-to-code generation for charts/web interfaces) using rendering-based feedback, not theory-of-mind reasoning for emotional understanding. The technical domains and supervision signals are fundamentally different.

---

8. Grounded Reinforcement Learning for Visual Reasoning

URL: View paper

Brief Assessment

Grounded Visual Reasoning[33] focuses on visual reasoning tasks with spatial grounding in image coordinates, not on Theory-of-Mind-based emotional reasoning or mental state supervision as in TMPO. The candidate uses reinforcement learning for visual attention and spatial anchoring, while TMPO uses intermediate mental states from ToM reasoning chains as process-level supervision for emotional understanding.

---

9. Visual Reinforcement Learning With Self-Supervised 3D Representations

URL: View paper

Brief Assessment

Self Supervised 3D[40] focuses on visual reinforcement learning with 3D representations for motor control tasks, not on theory-of-mind reasoning or mental state supervision for multimodal emotion understanding. The technical domains are entirely distinct.

---

10. Improving model-based reinforcement learning with internal state representations through self-supervision

URL: View paper

Brief Assessment

Internal State Representations[41] focuses on model-based RL with reconstruction losses for environment dynamics in control tasks, not on Theory-of-Mind reasoning chains or mental state supervision for multimodal emotion understanding.

Appendix: Text Similarity Detection

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

References

- [0] Unveiling the Cognitive Compass: Theory-of-Mind–Guided Multimodal Emotion Reasoning View paper
- [1] Scaling multimodal theory-of-mind with weak-to-strong bayesian reasoning View paper
- [2] Muma-tom: Multi-modal multi-agent theory of mind View paper
- [3] Multimind: Enhancing werewolf agents with multimodal reasoning and theory of mind View paper
- [4] Emotional Theory of Mind: Bridging Fast Visual Processing with Slow Linguistic Reasoning View paper
- [5] Exploring theory of mind in large language models through multimodal negotiation View paper
- [6] SoMi-ToM: Evaluating Multi-Perspective Theory of Mind in Embodied Social Interactions View paper
- [7] MMToM-QA: Multimodal Theory of Mind Question Answering View paper
- [8] Through the Theory of Mind's Eye: Reading Minds with Multimodal Video Large Language Models View paper
- [9] Echo: A visio-linguistic dataset for event causality inference via human-centric reasoning View paper
- [10] A conceptual framework of cognitive-affective theory of mind: towards a precision identification of mental disorders View paper
- [11] Consistency, uncertainty or inconsistency detection in multimodal emotion recognition View paper
- [12] Modeling theory of mind in multimodal HCI View paper
- [13] MindPower: Enabling Theory-of-Mind Reasoning in VLM-based Embodied Agents View paper
- [14] Overcoming Multi-step Complexity in Multimodal Theory-of-Mind Reasoning: A Scalable Bayesian Planner View paper
- [15] From Black Boxes to Transparent Minds: Evaluating and Enhancing the Theory of Mind in Multimodal Large Language Models View paper
- [16] Mind the Motions: Benchmarking Theory-of-Mind in Everyday Body Language View paper
- [17] Sarcasm explanation in multimodal dialogues View paper
- [18] 38 Cognitive and Affective Theory of Mind in Young and Elderly Patients with Multiple Sclerosis View paper
- [19] A Multimodal Theory of Affect Diffusion. View paper
- [20] Dynamic neural reconfiguration underpins empathy development in preschoolers: A multimodal EEG study View paper
- [21] Multimodal mental state analysis View paper
- [22] Emotional Intelligence in Artificial Agents: Leveraging Deep Multimodal Big Data for Contextual Social Interaction and Adaptive Behavioral Modelling View paper
- [23] Inference-enabled tracking of acute mental stress via multi-modal wearable physiological sensing: A proof-of-concept study View paper
- [24] Multimodal temporal context network for tracking dynamic changes in emotion View paper
- [25] Integrating emotion dynamics in mental health: A trimodal framework combining ecological momentary assessment, physiological measurements, and speech â View paper
- [26] Husformer: A Multimodal Transformer for Multimodal Human State Recognition View paper
- [27] Multimodal large language models meet multimodal emotion recognition and reasoning: A survey View paper
- [28] From Coarse to Nuanced: Cross-Modal Alignment of Fine-Grained Linguistic Cues and Visual Salient Regions for Dynamic Emotion Recognition View paper
- [29] Online Learning Platform of Modern Chinese Course Based on Multimodal Emotion-Aware Adaptive Learning View paper

- [30] Behavioral and physiological signals-based deep multimodal approach for mobile emotion recognition View paper
- [31] Directions for Computational Theory of Mind: Data, Metrics, Models View paper
- [32] Affective-CoT: Decomposing Multimodal Emotion Reasoning through a Hierarchical Cognitive Workflow View paper
- [33] Grounded Reinforcement Learning for Visual Reasoning View paper
- [34] Learning only with images: Visual reinforcement learning with reasoning, rendering, and visual feedback View paper
- [35] Multiagent inverse reinforcement learning via theory of mind reasoning View paper
- [36] SuperRL: Reinforcement Learning with Supervision to Boost Language Model Reasoning View paper
- [37] Reflexion: Language agents with verbal reinforcement learning View paper
- [38] Learning to reason without external rewards View paper
- [39] Mental modeling of reinforcement learning agents by language models View paper
- [40] Visual Reinforcement Learning With Self-Supervised 3D Representations View paper
- [41] Improving model-based reinforcement learning with internal state representations through self-supervision View paper
- [42] Relational deep reinforcement learning View paper