



# THE UNIVERSITY *of* EDINBURGH

## Edinburgh Research Explorer

### ToM-LM

Delegating theory of mind reasoning to external symbolic executors in large language models

**Citation for published version:**

Tang, W & Belle, V 2024, ToM-LM: Delegating theory of mind reasoning to external symbolic executors in large language models. in *Proceedings of the 18th International Conference for Neurosymbolic Learning and Reasoning*. Proceedings of the International Conference for Neurosymbolic Learning and Reasoning, 18th International Conference on Neural-Symbolic Learning and Reasoning, Barcelona, Spain, 9/09/24. <https://doi.org/10.48550/arXiv.2404.15515>

**Digital Object Identifier (DOI):**

[10.48550/arXiv.2404.15515](https://doi.org/10.48550/arXiv.2404.15515)

**Link:**

[Link to publication record in Edinburgh Research Explorer](#)

**Document Version:**

Peer reviewed version

**Published In:**

Proceedings of the 18th International Conference for Neurosymbolic Learning and Reasoning

**General rights**

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

**Take down policy**

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact [openaccess@ed.ac.uk](mailto:openaccess@ed.ac.uk) providing details, and we will remove access to the work immediately and investigate your claim.



# ToM-LM: Delegating Theory of Mind Reasoning to External Symbolic Executors in Large Language Models

Weizhi Tang<sup>[0009-0006-7444-292X]</sup> and Vaishak Belle<sup>[0000-0001-5573-8465]</sup>

University of Edinburgh, Edinburgh, UK  
{Weizhi.Tang, vbelle}@ed.ac.uk

**Abstract.** Theory of Mind (ToM) refers to the ability of individuals to attribute mental states to others. While Large Language Models (LLMs) have shown some promise with ToM ability, they still struggle with complex ToM reasoning. Our approach leverages an external symbolic executor, specifically the SMCDEL model checker, and fine-tuning to improve the ToM reasoning ability of LLMs. In our approach, an LLM is first fine-tuned through pairs of natural language and symbolic formulation representation of ToM problems and is then instructed to generate the symbolic formulation with a one-shot in-context example. The generated symbolic formulation is then executed by the SMCDEL model checker to perform transparent and verifiable ToM reasoning and give the final result. We demonstrate that our approach, ToM-LM, shows a significant improvement over all the constructed baselines. Our study proposes a novel view about externalizing a particular component of ToM reasoning, mainly reasoning about beliefs, and suggests generalizing it to other aspects of ToM reasoning.

**Keywords:** Large Language Models · Theory of Mind · Reasoning · Neuro-symbolic AI

## 1 Introduction

Theory of Mind (ToM) refers to an individual’s ability to attribute mental states to others [16, 2, 17]. This ability is essential for individuals to understand and reason about others’ beliefs, intentions, emotions, obligations, and so on [16, 5, 17]. For Large Language Models (LLMs), in most scenarios, they are required to interact with users or other agents and to understand their needs to give effective and appropriate responses. Therefore, discovering and enhancing their ToM reasoning ability is necessary and critical. Furthermore, it is worth noting that reasoning about knowledge has been a concern in symbolic philosophy and logic for many decades [4] and certainly in computer science to model security and game theory [1, 6, 7]. To a large extent, it is also the case that most of the machine learning models today have not really supported epistemic modeling and learning. Therefore, the idea that capabilities classically dominated in symbolic

computing could be used in the context of LLMs is interesting and presents intriguing possibilities.

Previous studies have demonstrated the emergence of ToM reasoning ability in LLMs [10, 13, 8, 11]. Nevertheless, most of these studies adopt the ToM evaluation problems originally designed for assessing human ToM reasoning ability, which could be widely exposed to the internet, leading to possible data leakage to the training processes of LLMs [12, 20]. In addition, because of the simplicity of these ToM problems, in order to study whether LLMs can handle more complex ToM reasoning ability, new benchmarks are needed [12, 20]. Recent work has developed a new benchmark called MindGames [20] aiming to give a more complex ToM benchmark specifically designed for evaluating the ToM reasoning ability of LLMs. This benchmark also indicates that despite increases in model size, LLMs still struggle with complex ToM reasoning tasks [20].

To reiterate, although ToM is widely recognized and discussed as an inherent ability in LLMs, we propose a new perspective, externalizing and delegating ToM reasoning process to external symbolic executors. Inspired by Logic-LM, which enhances the logical reasoning abilities of LLMs by leveraging external executors that provide explainable and faithful reasoning processes [15], we adopt a similar approach to improve the ToM reasoning ability of LLMs. Specifically, we utilize an external symbolic executor, the SMCDEL model checker [22] designed for handling dynamic epistemic logic problems, to tackle complex ToM reasoning tasks delegated by LLMs, thereby aiming to improve the ToM reasoning ability of LLMs while making the ToM reasoning process transparent and verifiable.

In our approach, an LLM is first given a ToM problem, and then it is asked to convert the problem represented in the natural language to the symbolic formulation. Subsequently, the generated symbolic formulation is then executed by the SMCDEL model checker to perform the ToM reasoning in a transparent and verifiable way and give out the final answer. However, our investigations reveal that this strategy alone does not adequately enhance the ToM reasoning ability of LLMs. Therefore, we employ fine-tuning<sup>1</sup> to further optimize this approach, aiming to substantially improve the ToM reasoning ability of LLMs. The overview of ToM-LM is shown in Figure 1.

Our study demonstrates that ToM-LM significantly improves the ToM reasoning ability of LLMs in terms of accuracy and Area under the ROC Curve (AUC), compared to all constructed baselines. The contributions of this study are outlined as follows:

1. We introduce a new perspective to enhance ToM reasoning ability of LLMs, externalizing and delegating ToM reasoning process to external executors to make it transparent and explainable.
2. Our study also demonstrates that this approach makes substantial enhancements over the constructed baselines.

<sup>1</sup> We fine-tune the chosen LLM, *gpt-3.5-turbo*, through OpenAI Fine-tuning Platform with pairs of natural language and symbolic formulation. More details are discussed later in this paper.

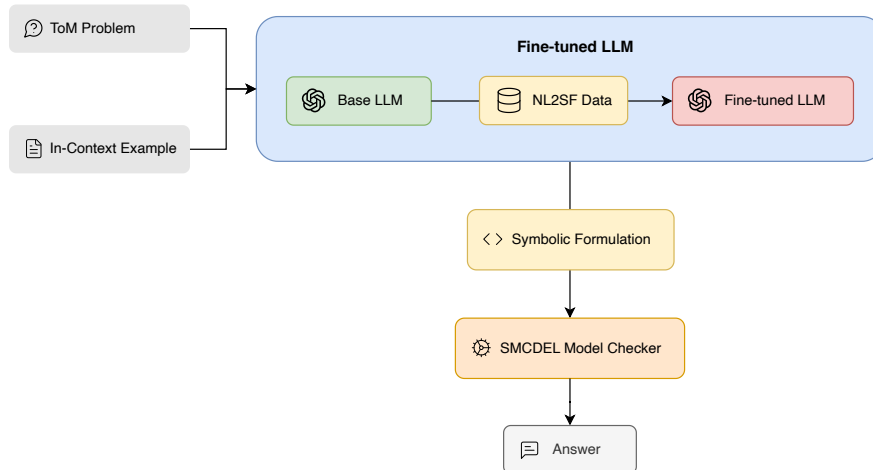


Fig. 1: Overview of ToM-LM, which consists of four principle stages: (1) *Model Fine-tuning* in which a given base model undergoes the fine-tuning process, (2) *Problem Prompting* in which a ToM problem and a one-shot example are constructed and given to the fine-tuned model, (3) *Problem Formalization* in which the model generates the corresponding symbolic formulation, and (4) *External Deterministic ToM Reasoning* in which the generated symbolic formulation is executed by SMCDEL model checker to give out the final result.

## 2 Related Work

### 2.1 ToM In LLMs

Previous works have discussed the emergence of ToM abilities in LLMs [10, 8, 18, 21, 3]. Some works also try in various ways to improve the ToM reasoning ability of LLMs. As in [13], the various in-context learning methods are utilized such as *Let’s think step by step* [9] and the chain-of-thought [23] to enhance the ToM reasoning ability of LLMs. Although this approach brings up an easy and flexible way to improve LLMs’ ToM reasoning ability, we still hold the view that dynamic and complex ToM problems can not be simply represented as a chain-of-thought reasoning process and a more appropriate and robust way to handle them is needed. Thus, we propose an approach to use symbolic formulation to represent a general way to capture the ToM reasoning process and use external symbolic executors to perform the reasoning.

### 2.2 Delegating to External Executors

Our approach is significantly inspired by [19] which demonstrates delegating complex functionalities to external tools and by [15, 14] which leverage symbolic formulation generation of LLMs and external symbolic executors to enhance

their logical reasoning abilities. We follow a similar methodology and algorithmic strategy and hypothesize that it should be also suitable for the enhancement of complex ToM reasoning.

### 3 Methodology

#### 3.1 ToM-LM

The overview of ToM-LM is shown in Figure 1. It consists of four principle stages, each of which is delineated in detail as follows.

**Model Fine-tuning** During this stage, we fine-tune the LLM, specifically the *gpt-3.5-turbo*, using our newly constructed *Symbolic Formulation Generation Prompting Fine-tuning Dataset*<sup>2</sup> which consists of pairs of ToM problems represented in both natural language and symbolic formulation. The fine-tuning was being processed on the OpenAI Fine-tuning Platform with 3 epochs. This process aims to improve the performance of the generation of symbolic formulations.

**Problem Prompting** Subsequently, a ToM problem and a one-shot in-context example are constructed as a prompt and given to the fine-tuned LLM to instruct it to give an appropriate symbolic formulation<sup>3</sup>.

**Problem Formalization** The fine-tuned LLM is then expected to generate the symbolic formulation to appropriately represent the problem described in the natural language. The generated symbolic formulations are anticipated to be executable<sup>4</sup>.

**External Deterministic ToM Reasoning** Finally, the generated symbolic formulation is then executed by the external symbolic executor, specifically the SMCDEL model checker, to perform the ToM reasoning in a transparent and verifiable way and give out the final result.

#### 3.2 Dataset

We select MindGames [20] as our evaluation and fine-tuning dataset. It is a dataset designed specifically for evaluating the ToM reasoning ability of the LLMs. Each item<sup>5</sup> in the dataset represents a ToM problem including a premise which describes the ToM problem situation and context, a hypothesis which

<sup>2</sup> *Symbolic Formulation Generation Prompting Fine-tuning Dataset* is explained in Section 3.2.

<sup>3</sup> Details of the prompting template are described in Appendix A.2.

<sup>4</sup> An example of the symbolic formulation is given in Appendix B.

<sup>5</sup> An example is given in Appendix B.

describes a situation needing LLMs to judge its validity, a Boolean label which is the ground truth of the validity of the hypothesis, and a corresponding symbolic formulation to the premise and hypothesis. Each problem in MindGames can be categorized as either a true-belief or false-belief problem. LLMs are expected to handle them correctly to show their ToM reasoning ability.

We then randomly sample 200 items from the test set of MindGames dataset, as the evaluation dataset, to ensure an even distribution of labels. It means the ground truth labels in this sampled evaluation dataset are meticulously balanced to achieve an exact distribution of 50% *True* label and 50% *False* label.

Additionally, we randomly sample 150 items from the MindGames train set for fine-tuning. The dataset is preprocessed in two ways for different purposes: (1) *Direct Prompting Fine-tuning Dataset*, in which each problem is formatted using the *Direct Prompting Template*<sup>6</sup> to generate prompts prepared for the DP<sub>FT</sub> baseline; (2) *Symbolic Formulation Generation Prompting Fine-tuning Dataset*, in which each problem is formatted using the *Symbolic Formulation Generation Prompting Template*<sup>7</sup> to generate prompts prepared for ToM-LM.

### 3.3 Baselines

We choose *gpt-3.5-turbo* as the base model and form three baselines to compare with our approach. The baselines are then differentiated as follows: (1) *Direct Prompting* or *Symbolic Formulation Generation Prompting*, and (2) whether or not Fine-tuning is applied. Our approach, ToM-LM, can be recognized as *Symbolic Formulation Generation Prompting with Fine-tuning*.

**DP** It stands for *Direct Prompting Without Fine-tuning*. In this setting, we directly ask the base model to answer the problem in "TRUE", "FALSE", or "I DON'T KNOW", with a one-shot in-context example provided. Answering in "TRUE" means it considers the hypothesis is valid in the given situation, "FALSE" means it thinks the hypothesis is invalid, and "I DON'T KNOW" means it has no idea about the validity of the hypothesis.

**SFGP** It stands for *Symbolic Formulation Generation Prompting Without Fine-tuning*. In this setting, the workflow is the same as ToM-LM but the model is not fine-tuned.

**DP<sub>FT</sub>** It stands for *Direct Prompting With Fine-tuning*. In this setting, the workflow is the same as the DP baseline but the model is fine-tuned based on the *Direct Prompting Fine-tuning Dataset*.

<sup>6</sup> Details of the prompting template are described in Appendix A.1.

<sup>7</sup> Details of the prompting template are described in Appendix A.2.

## 4 Results

The responses given by DP, SFGP,  $DP_{FT}$ , and ToM-LM can be categorized and summarized into three types in general: (1) *TRUE*, (2) *FALSE*, or (3) *I DON'T KNOW*.

For *Direct Prompting* setting, *TRUE* simply means the LLM answers in "TRUE" literally, *FALSE* means the LLM answers in "FALSE" literally, and answering in "I DON'T KNOW" literally or other contents are considered as answering in *I DON'T KNOW*.

For *Symbolic Formulation Generation Prompting* setting, *TRUE* means the LLM generates an executable symbolic formulation and the model checker executes it to get *TRUE* value towards the problem, *FALSE* means the LLM generates an executable symbolic formulation and the model checker executes it to get *FALSE* value towards the problem, and *I DON'T KNOW* in this case means the LLM fails to generate an executable symbolic formulation.

### 4.1 Metrics

The evaluation metrics consist of three aspects, (1) Execution Rate, (2) Accuracy, and (3) Area under the ROC Curve (AUC). Table 1 shows all the three metrics for DP, SFGP,  $DP_{FT}$ , and ToM-LM.

Table 1: The Metrics of DP, SFGP,  $DP_{FT}$ , and ToM-LM

Approach	Execution Rate(%)	Accuracy(%)	AUC
DP	99.50	58.00	0.58
SFGP	78.00	49.00	0.60
$DP_{FT}$	<b>100</b>	76.00	0.76
ToM-LM	94.50	<b>91.00</b>	<b>0.94</b>

**Execution Rate** This metric is used to measure how successfully the LLM can generate available or executable answers. In the case of *Direct Prompting* setting, it is *TRUE* or *FALSE* and in the case of *Symbolic Formulation Generation Prompting* setting, it is executable symbolic formulation. A higher value of execution rate means higher available or executable answers generated by the LLM. But it does not mean the generated answers are correct.

The execution rates of baselines in *Direct Prompting* setting are almost around 100%, which means the LLM can generate available answers at most of time in this setting. Instead, in the case of *Symbolic Formulation Generation Prompting* setting, the execution rate is 78% without fine-tuning but can be high up to 94.5% in ToM-LM.

**Accuracy** As demonstrated in Table 1, ToM-LM reaches up to 91% accuracy which significantly surpasses all the baselines. It achieves the 33% improvement over the DP baseline, the 42% increase compared to the SFGP baseline, the 15% enhancement over the DP<sub>FT</sub> baseline.

**AUC** Additionally, AUC is utilized to provide a more holistic indicator and assessment of the performance of ToM-LM and the baselines. As shown in Table 1, while the AUC for DP and SFGP are around 0.60, ToM-LM reaches up to the 0.94 AUC and outperforms all the baselines.

## 4.2 Output Distributions

We have also illustrated the distributions of outputs in Figure 2 to provide an intuitive visualization of the results and performance. Given that the evaluation set is evenly distributed across the ground truth labels, the distributions of outputs can intuitively reflect performance differences. While the output distributions for all baseline models exhibit significant bias, the output distribution of ToM-LM is notably more balanced.




Fig. 2: Output Distributions of DP, SFGP, DP<sub>FT</sub>, and ToM-LM

## 5 Discussion

### 5.1 Results Analysis

The results demonstrate that ToM-LM has successfully enhanced the ToM reasoning ability of LLMs, surpassing the other three baselines in terms of both accuracy, which reaches as high as 91%, and AUC which attains 0.94.

From the results of the DP baseline, the accuracy is 58% and AUC is 0.58, meaning although the LLM itself may possess a modest degree of ToM reasoning ability, it still struggles with it and its answers are highly dependent on random



guessing. From the SFGP baseline, although AUC is shown with a marginal improvement by 0.02 and reaches 0.6, indicating that it modestly enhances the performance of the LLM against random guessing, the gain comes at the cost of reduced accuracy and a highly increased incidence of non-executable symbolic formulations.

As anticipated, the  $DP_{FT}$  baseline markedly surpasses both of the DP and SFGP baseline in terms of both accuracy and AUC. When comparing it to our approach, ToM-LM shows significant improvement over it with both of accuracy and AUC. In addition, within the same fine-tuning settings, we also observe that the training loss of the  $DP_{FT}$  baseline has a violent oscillation and does not converge at the end but the training loss of the ToM-LM has a very slighter oscillation and it converges at the end of fine-tuning<sup>8</sup>.

Furthermore, the distributions of outputs indicate a significant bias among the three baselines, with all tending to respond *TRUE* to the given ToM problems. This pattern reveals their inability to accurately handle false-belief tasks, underscoring a deficiency in ToM reasoning ability. And as shown in Figure 2, it is noteworthy that even in the SFGP baseline, the LLM fails to generate appropriate executable symbolic formulations for false-belief tasks, while predominantly generating symbolic formulations for true-belief tasks instead. However, ToM-LM significantly improves the balance of results compared to all the three baselines.

## 5.2 Limitations and Future Work

While ToM is supposed to encompass a broad range of aspects, such as emotions, goals, intentions, desires, and beyond, our study only focuses on handling false-belief and true-belief of ToM due to the expressive limitations of the current external executor. Nevertheless, this approach has the potential to be generalized to handle other ToM tasks. Future research should explore additional components of ToM using external executors and may also aim to develop a more expressive and capable external executor or language to handle a diverse array of ToM tasks.

## 6 Conclusion

In this study, we propose a novel approach, ToM-LM, to externalize the ToM reasoning process of LLMs by employing the external symbolic executor, specifically the SMCDEL model checker. It aims to improve the ToM reasoning ability of LLMs while also making the ToM reasoning process transparent and explainable. Our approach demonstrates significant improvements over the three constructed baselines. We also anticipate that this research can offer new insights about improving other aspects of ToM reasoning ability of LLMs.

<sup>8</sup> Loss curves for both  $DP_{FT}$  and ToM-LM are given in Appendix C.

## References

1. Abraham, I., Alvisi, L., Halpern, J.Y.: Distributed computing meets game theory: combining insights from two fields. *ACM SIGACT News* **42**(2), 69–76 (Jun 2011). <https://doi.org/10.1145/1998037.1998055>
2. Baron-Cohen, S., Leslie, A.M., Frith, U.: Does the autistic child have a “theory of mind”? *Cognition* **21**(1), 37–46 (Oct 1985). [https://doi.org/10.1016/0010-0277\(85\)90022-8](https://doi.org/10.1016/0010-0277(85)90022-8)
3. Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P., Lee, Y.T., Li, Y., Lundberg, S., et al.: Sparks of artificial general intelligence: Early experiments with gpt-4. arXiv preprint arXiv:2303.12712 (2023)
4. Fagin, R., Halpern, J.Y., Moses, Y., Vardi, M.: Reasoning about knowledge. MIT press (2004)
5. Frith, C.D., Frith, U.: The neural basis of mentalizing. *Neuron* **50**(4), 531–534 (May 2006). <https://doi.org/10.1016/j.neuron.2006.05.001>
6. Halpern, J.Y.: Substantive rationality and backward induction. *Games and Economic Behavior* **37**(2), 425–435 (Nov 2001). <https://doi.org/10.1006/game.2000.0838>
7. Halpern, J.Y., Pass, R., Raman, V.: An epistemic characterization of zero knowledge. In: Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge. p. 156–165. TARK ’09, Association for Computing Machinery, New York, NY, USA (Jul 2009). <https://doi.org/10.1145/1562814.1562837>, <https://dl.acm.org/doi/10.1145/1562814.1562837>
8. Jamali, M., Williams, Z.M., Cai, J.: Unveiling theory of mind in large language models: A parallel to single neurons in the human brain
9. Kojima, T., Gu, S.S., Reid, M., Matsuo, Y., Iwasawa, Y.: Large language models are zero-shot reasoners (arXiv:2205.11916) (Jan 2023), <http://arxiv.org/abs/2205.11916>, arXiv:2205.11916 [cs]
10. Kosinski, M.: Evaluating large language models in theory of mind tasks (2023). <https://doi.org/10.48550/ARXIV.2302.02083>, <https://arxiv.org/abs/2302.02083>
11. Li, H., Chong, Y.Q., Stepputtis, S., Campbell, J., Hughes, D., Lewis, M., Sycara, K.: Theory of mind for multi-agent collaboration via large language models (arXiv:2310.10701) (Oct 2023), <http://arxiv.org/abs/2310.10701>, arXiv:2310.10701 [cs]
12. Ma, Z., Sansom, J., Peng, R., Chai, J.: Towards a holistic landscape of situated theory of mind in large language models (arXiv:2310.19619) (Oct 2023), <http://arxiv.org/abs/2310.19619>, arXiv:2310.19619 [cs]
13. Moghaddam, S.R., Honey, C.J.: Boosting theory-of-mind performance in large language models via prompting
14. Olausson, T., Gu, A., Lipkin, B., Zhang, C., Solar-Lezama, A., Tenenbaum, J., Levy, R.: Linc: A neurosymbolic approach for logical reasoning by combining language models with first-order logic provers. In: Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing. p. 5153–5176. Association for Computational Linguistics, Singapore (2023). <https://doi.org/10.18653/v1/2023.emnlp-main.313>, <https://aclanthology.org/2023.emnlp-main.313>
15. Pan, L., Albalak, A., Wang, X., Wang, W.Y.: Logic-lm: Empowering large language models with symbolic solvers for faithful logical reasoning (arXiv:2305.12295) (Oct 2023), <http://arxiv.org/abs/2305.12295>, arXiv:2305.12295 [cs]

16. Premack, D., Woodruff, G.: Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences* **1**(4), 515–526 (Dec 1978). <https://doi.org/10.1017/S0140525X00076512>
17. Quesque, F., Rossetti, Y.: What do theory-of-mind tasks actually measure? theory and practice. *Perspectives on Psychological Science* **15**(2), 384–396 (Mar 2020). <https://doi.org/10.1177/1745691619896607>
18. Sap, M., LeBras, R., Fried, D., Choi, Y.: Neural theory-of-mind? on the limits of social intelligence in large lms (arXiv:2210.13312) (Apr 2023), <http://arxiv.org/abs/2210.13312>, arXiv:2210.13312 [cs]
19. Schick, T., Dwivedi-Yu, J., Dessì, R., Raileanu, R., Lomeli, M., Zettlemoyer, L., Cancedda, N., Scialom, T.: Toolformer: Language models can teach themselves to use tools (arXiv:2302.04761) (Feb 2023), <http://arxiv.org/abs/2302.04761>, arXiv:2302.04761 [cs]
20. Sileo, D., Lerneuld, A.: Mindgames: Targeting theory of mind in large language models with dynamic epistemic modal logic (arXiv:2305.03353) (Nov 2023), <http://arxiv.org/abs/2305.03353>, arXiv:2305.03353 [cs]
21. Trott, S., Jones, C., Chang, T., Michaelov, J., Bergen, B.: Do large language models know what humans know? *Cognitive Science* **47**(7), e13309 (2023)
22. Van Benthem, J., Van Eijck, J., Gattinger, M., Su, K.: Symbolic model checking for dynamic epistemic logic — s5 and beyond\*. *Journal of Logic and Computation* **28**(2), 367–402 (Mar 2018). <https://doi.org/10.1093/logcom/exx038>
23. Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E., Le, Q., Zhou, D.: Chain-of-thought prompting elicits reasoning in large language models (2022). <https://doi.org/10.48550/ARXIV.2201.11903>, <https://arxiv.org/abs/2201.11903>

## A Prompting Templates

### A.1 Direct Prompting Template

Figure 3 shows the template used for prompting to directly generate the prediction of the answer. The texts within the braces serve as placeholders to be replaced with actual example or problem contents. The placeholders beginning with "example" indicate that they are for the one-shot in-context ToM example, while the placeholders "context" and "hypothesis" are for the ToM problem.

```

You will be given a problem with a context and hypothesis. You should judge whether the hypothesis is "TRUE", "FALSE", or "I
DONT KNOW" in the given context. "TRUE" means the hypothesis is true in the context, "FALSE" means the hypothesis is
false in the context, and "I DONT KNOW" means you cannot determine the truth value of the hypothesis in the context. You
should provide your answer in the form of "TRUE", "FALSE", or "I DONT KNOW".

-----

Example:

Context:
{example_context}
Hypothesis:
{example_hypothesis}
Answer:
{example_answer}

-----

Context:
{context}
Hypothesis:
{hypothesis}
Answer:

```

Fig. 3: Direct Prompting Template

### A.2 Symbolic Formulation Generation Prompting Template

Figure 4 demonstrates the template used for prompting to generate the prediction of the symbolic formulation. The texts within the braces serve as placeholders to be replaced with actual example or problem contents. The placeholders beginning with "example" indicate that they are for the one-shot in-context ToM example while the placeholders starting with "problem" mean they are for the ToM problem.

```

You should transform the natural language representation to the correct SMODEL symbolic formulation. You should only output
the transformed symbolic formulation as shown in the example.
-----
=== Example Begin ===
Natural Language Representation:
Context:
{example_context}
Hypothesis:
{example_hypothesis}

Symbolic Formulation:
{example_sf}
=== Example End ===
-----
Natural Language Representation:
Context:
{problem_context}
Hypothesis:
{problem_hypothesis}

Symbolic Formulation:

```

Fig. 4: Symbolic Formulation Generation Prompting Template

## B Problem Example

An example of a ToM problem consisting of a premise, a hypothesis, and a symbolic formulation given in MindGames [20] is illustrated in Figure 5.

The premise describes a scenario involving four agents, each drawing a face unrevealed card, and someone may reveal their card to others. The hypothesis to be validated based on this given scenario is "Cara can now know whether Conrad picked a red card". The corresponding symbolic formulation representation of the premise and hypothesis is also given. To briefly explain<sup>9</sup>, "VARS 1,2,3,4" defines four atomic propositions which represent the status of each agent, such as that "2" represents "Conrad picked a red card". "LAW Top" indicates that all possible Boolean states of these statuses should be considered. "OBS" simply denotes which agent observes which status given in "VARS", such as that "Agentc:1" means Cara observes the status that "Vasiliki picked a red card". The "VALID?" line queries an agent's belief about a certain status, in which "!" means a public announcement to all agents and "|" is a logical *OR*, such as that "!(1|2|3|4)" means "It is publicly announced that someone picked a red card". It represents and encapsulates the main part of the premise and hypothesis.

<sup>9</sup> Refer to <https://github.com/jrclogic/SMODEL> for more syntax details.

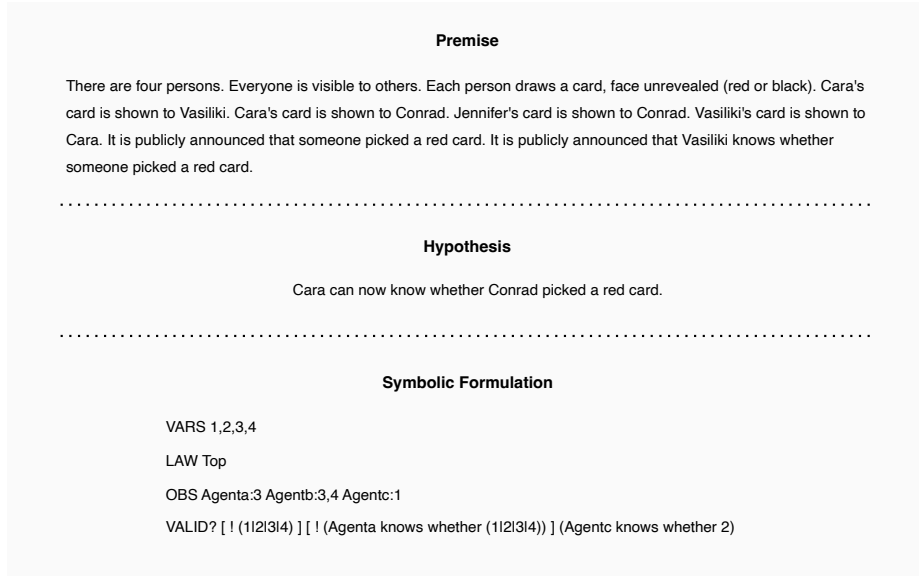


Fig. 5: An example of premise, hypothesis, and symbolic formulation in MindGames.

## C Fine-tuning Loss Curves

### C.1 Loss Curve of Direct Prompting With Fine-tuning

As the loss curve illustrated in Figure 6, the training loss of the  $DP_{FT}$  baseline is high with the range of loss between 0 and 9 as shown in the y-axis, while also exhibits severe oscillations and failing to converge by the end of training.




Fig. 6: Loss Curve of Direct Prompting With Fine-tuning

### C.2 Loss Curve of ToM-LM

As shown in Figure 7, the training loss of the ToM-LM is low with the range of loss between 0 and 0.9 and converges at the end of training.




Fig. 7: Loss Curve of ToM-LM