

NL2CA: Auto-formalizing Cognitive Decision-Making from Natural Language Using an Unsupervised CriticNL2LTL Framework

Zihao Deng^{1,2}, Yijia Li^{1,2}, Renrui Zhang^{1,2}, Peijun Ye^{1*}

¹Institute of Automation, Chinese Academy of Sciences

²School of Artificial Intelligence, University of Chinese Academy of Sciences
{dengzihao2025, liyijia2023, zhangrenrui2024, peijun.ye}@ia.ac.cn

Abstract

Cognitive computing models offer a formal and interpretable way to characterize human’s deliberation and decision-making, yet their development remains labor-intensive. In this paper, we propose NL2CA, a novel method for auto-formalizing cognitive decision-making rules from natural language descriptions of human experience. Different from most related work that exploits either pure manual or human-guided interactive modeling, our method is fully automated without any human intervention. The approach first translates text into Linear Temporal Logic (LTL) using a fine-tuned large language model (LLM), then refines the logic via an unsupervised Critic Tree, and finally transforms the output into executable production rules compatible with symbolic cognitive frameworks. Based on the resulted rules, a cognitive agent is further constructed and optimized through cognitive reinforcement learning according to the real-world behavioral data. Our method is validated in two domains: (1) NL-to-LTL translation, where our CriticNL2LTL module achieves consistent performance across both expert and large-scale benchmarks without human-in-the-loop feed-backs, and (2) cognitive driving simulation, where agents automatically constructed from human interviews have successfully learned the diverse decision patterns of about 70 trials in different critical scenarios. Experimental results demonstrate that NL2CA enables scalable, interpretable, and human-aligned cognitive modeling from unstructured textual data, offering a novel paradigm to automatically design symbolic cognitive agents.

Introduction

Cognitive architectures such as ACT-R (Anderson and Lebiere 2014) and Soar (Laird 2022) are rule-based models, which can simulate various cognitive tasks such as memory, problem solving, and learning. Such architectures have achieved wide applications in human behavior modeling due to their excellent simulation of human cognitive processes and high interpretability. But such architectures require extensive expert effort to instantiate and depend almost exclusively on human experts’ prior knowledge to design.

Large language models (LLM) such as GPT-4 (Achiam et al. 2023) and Deepseek (Liu et al. 2024), on the other hand, have shown promising capabilities in imitating human

reasoning and behavior after a post-training on corpus with human feedbacks (Kumar et al. 2025). Even though LLMs are still being criticized for hallucination (Rawte et al. 2023) and diminishing returns for scaling (Villalobos et al. 2024), their knowledge of human cognitive process can still be used to automate the instantiation of the decision-making model with cognitive architectures (Kirk, Wray, and Laird 2023) (Niu et al. 2024) (Wray, Kirk, and Laird 2025).

The core component of the cognitive architecture are the declarative memory and procedural memory represented by a set of production rules, each with a precondition and an effect. Cognitive agents operate in perceive–plan–act cycles, dynamically matching environmental features against these rules to determine their actions. There are already some studies on automatically generating production rules using LLM (Zhu and Simmons 2024) (Kirk et al. 2024), but almost all previous work aims to generate the production rules by directly querying the LLM about how the agent should act under an unknown situation. Such approaches construct the cognitive models without human prior knowledge and data, which could lead to a relatively poor alignment with actual human behaviors.

In this work, we propose a novel approach for cognitive agent construction: **NL2CA** (Natural Language to Cognitive Agent), which leverages LLMs to formalize production rules from natural language descriptions of human experiences. As shown in Figure 1, the LLM is used to interpret textual descriptions of human deliberations and convert them into structured representations. Thus, the LLM serves not merely as a rule-generator, but also as a knowledge-extractor. The generated production rules are based both on the human experience document and the LLM, so they are expected to be more human-aligned.

Realizing this vision requires LLM to have both the capability of providing correct knowledge and the capability of formalizing production rules from human experience documents. Since the former capability is already well demonstrated in previous work, our work focuses mainly on improving the LLM’s capability to formalize production rules from human experience documents.

To improve the accuracy and reliability of this production rule extraction, we first prompt the LLM to formalize the logical information of the NL description into Linear Temporal Logic (LTL), which is equivalently a powerful formal-

*Corresponding author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

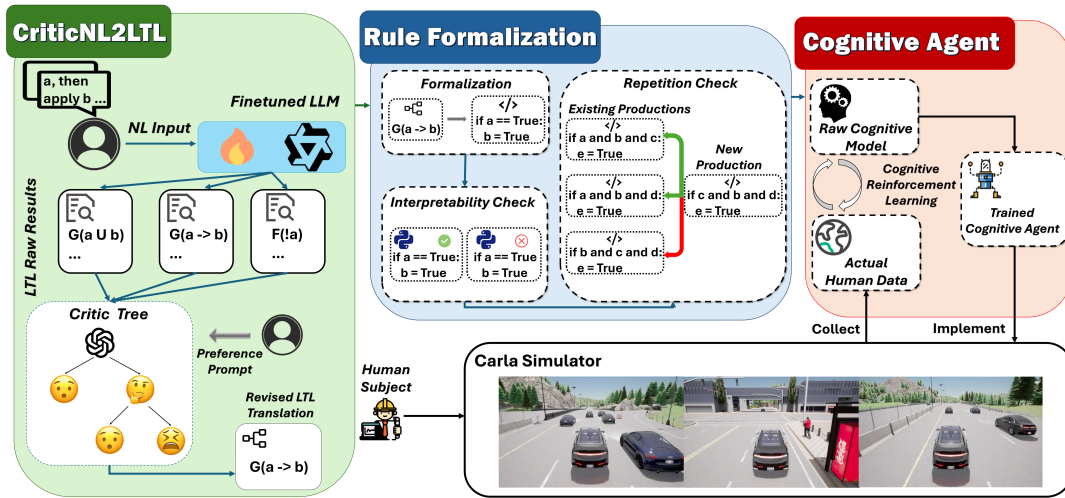


Figure 1: Pipeline of the NL2CA framework.

ism in various domains such as robotics (Wang et al. 2025), electronics design (Browne et al. 1986) and autonomous driving (Maierhofer, Moosbrugger, and Althoff 2022). To facilitate this formalization, We demonstrate a novel framework combining a fine-tuned LLM and a heterogeneous unsupervised Tree-of-Thought (ToT) like Critic Tree structure to improve the NL2LTL translation accuracy. This enhances the generality of our proposed framework, enabling its applicability not only to cognitive modeling, but also to other engineering domains where LTL is operative.

The key contribution of this paper is as follows: 1) we propose a novel framework for fully auto-formalizing cognitive decision-making rules from natural language descriptions of human experience, 2) we propose a more strengthened automated approach to tackle the natural language to linear temporal logic formalization, 3) we examined the validity of our proposed framework in various driving simulation scenarios.

Related Work

Cognitive Architecture Design Automation

Cognitive architectures (e.g., ACT-R, Soar) have proven to be effective in modeling human cognition but require extensive manual effort. To bridge this burden, recent efforts have tried to leverage LLMs for the instantiation of the cognitive model. STARS (Kirk et al. 2024) and Zhu and Simmons (2024) directly query LLMs to generate production rules for specific environmental states. Bajaj et al. (2023) proposes a knowledge extraction mechanism, where LLMs are used to populate symbolic memory structures.

While these approaches enable a faster construction of cognitive systems, they often suffer from limited application in actual human cognitive modeling, as those methods are either task-learning based and not grounded in real human decision traces, or simply for knowledge extraction. To the best of our knowledge, there is still no work on automating the human cognitive modeling pipeline. Our work addresses these limitations by formalizing natural language descriptions of human experience into LTL, preserving both

traceability and human-likeness, and transforming it into executable production rules for cognitive modeling.

Cognitive Driving Model

The idea of combining cognitive architecture to imitate driver behaviors was first brought up by Salvucci (2006). They proposed a general pipeline of how to model driver behaviors by using a specific cognitive architecture and introduced a threaded cognition for driving tasks.

Later, some approaches were conducted to further extend cognitive architectures' application to the construction of virtual drivers. Ritter, Tehranchi, and Oury (2019) elucidated the operation of ACT-R, describing how driver deals with information and generates behaviors through monitoring, decision-making and control processes. Oh, Yun, and Myung (2024) proposed an approach to predict mental workload during take-over situations in autonomous driving. Although numerous attempts were made, due to their lack of learning capability, early approaches were often limited to singular scenarios, thus their applications to autonomous driving are still limited (Chater et al. 2018).

To address these issues, Qi, Hou, and Ye (2024) proposed a cognitive reinforcement learning method to align the cognitive reasoning via human-machine interactions. Li et al. (2024) further refined such method and reported an interpretable cognitive driving model. Their method was reported to outperform standard PPO algorithm, while offering a highly interpretable pattern.

However, these methods require manual specification of production rules, limiting their extensibility. In contrast, our NL2CA framework automates the pipeline by directly formalizing production rules from NL data, making it more extensible to multiple scenarios without manual inspections.

Method

The NL2CA framework consists of three stages: 1) Formalizing logic representations from NL text using a fine-tuned LLM and an unsupervised Critic Tree, 2) Transforming the

resulted temporal logic into executable cognitive production rules, and 3) Training a cognitive agent through cognitive reinforcement learning (CRL) on human data.

In the first stage, an input NL segment is formalized as a Linear Temporal Logic (LTL) formula. The Critic Tree module iteratively refines the formula to better align with the meaning of the original sentence. The refined LTL is then grounded with variables from a predefined knowledge base and transformed into production rules compatible with cognitive modeling frameworks.

In the final stage, a cognitive symbolic-driven agent is constructed based on the generated production rules. The agent is trained on real human behavioral data using cognitive reinforcement learning method, adjusting the rules' utilities based on how closely its outputs match human behaviors. Figure 1 illustrates the NL2CA pipeline.

CriticNL2LTL Framework

To enable an accurate and unsupervised NL to LTL formalization, we introduce an CriticNL2LTL module, which consists of two main components: a fine-tuned LLM for initial translation and a heterogeneous Critic Tree for iterative refinement, which is shown in Algorithm 1.

Given an NL input, the fine-tuned LLM first generates a raw LTL formula. This formula is then passed to the Critic Tree, a self-organizing revision structure with user-configurable parameters like maximum depth and number of critics. The Critic Tree consists of a revisor LLM, which is responsible for revising the current formula, and critic LLMs, which evaluate the formula's logical correctness and its alignment with user-defined preferences (e.g., grounded knowledge base, desired format). When a critic detects a mismatch, it appends a feedback message to the context of the revisor, instructing the revisor to revise the formula and split a new child node, the child node stores the revised LTL formula and the context of the revisor. The process repeats until all critics approve and no new child nodes can be created or the maximum depth is reached. We construct the critics using heterogeneous LLMs to provide diverse suggestions for revision, since standard Self-refine approaches are reported to fail to improve the accuracy of NL-to-LTL translation (Fang et al. 2025).

Logic to Code Transformation

After the logic formalization, the LTL formulas are transformed into executable production rules compatible with **pyactr**, an open source Python library used to create and run cognitive models. We leverage LLM to align the propositions of the LTL formulas with computational variables and to ensure the formalization. This step also includes a post-process to ensure the logical correctness, interpretability, and non-redundancy.

To address common issues such as Function interface mismatch, Over-constraining/generalization and Rule duplication, reported by Zhu and Simmons (2024), we implement the following mechanisms:

- **Function interface mismatch:** The generated code is immediately interpreted using **pyactr**. If execution fails,

Algorithm 1: Critic Tree

Input: Natural language text T , Initial translation A_{init}

Parameters: Number of critics δ , Maximum depth D

Output: Final translation A_{final}

```

1:  $(A_{\text{root}}, C_{\text{root}}) \leftarrow \text{Revisor}(T, A_{\text{init}})$ 
2: Initialize root node  $v_0 = (A_{\text{root}}, C_{\text{root}})$ 
3: Initialize node list  $\mathcal{N} \leftarrow \{v_0\}$ 
4: for  $d = 0$  to  $D$  do
5:    $\mathcal{N}_{\text{new}} \leftarrow \emptyset$ 
6:   for each node  $v = (A, C)$  in  $\mathcal{N}$  do
7:     if  $v$  already has children then
8:       continue
9:     end if
10:     $results \leftarrow \emptyset$ 
11:    for  $k = 1$  to  $\delta$  do
12:       $r_k \leftarrow \text{Critic}(A, T)$ 
13:       $results.append(r_k)$ 
14:      if all  $results$  indicate approval then
15:         $A_{\text{final}} \leftarrow A$ 
16:        return  $A_{\text{final}}$ 
17:      end if
18:       $C' \leftarrow C \cup r_k$ 
19:       $(A', C'') \leftarrow \text{Revisor.revise}(C')$ 
20:      Add new node  $v' = (A', C'')$  to  $\mathcal{N}_{\text{new}}$ 
21:    end for
22:  end for
23:   $\mathcal{N} \leftarrow \mathcal{N}_{\text{new}}$ 
24: end for
25:  $A_{\text{final}} \leftarrow A_{\text{root}}$ 
26: return  $A_{\text{final}}$ 

```

the error message and the faulty rule are returned to the LLM for revision via an error-guided correction loop.

- **Over-constraining / over-generalization:** Most logical inconsistencies are filtered by the CriticNL2LTL module. However, ambiguous or contradictory input texts may still yield problematic rules. These are handled during the cognitive reinforcement learning stage when rules are evaluated against actual human behaviors.
- **Rule duplication:** To avoid redundant rules, each rule is assigned a name based on its preconditions and effects. We transform the rule names into BERT embeddings and use cosine similarity to measure affinity. Each newly generated rule will be compared with existing rules with top 5 similarity scores. If a newly generated rule matches an existing one semantically, it is discarded.

These measures ensure that only interpretable, distinct, and executable production rules are retained for cognitive agent modeling.

Cognitive Reinforcement Learning

The production rules extracted from natural language may still contain inconsistencies or errors due to ambiguities in the input. To improve reliability, we apply cognitive reinforcement learning to adjust each rule's utility based on its alignment with human decision-making behaviors.

Strategy	Model	Dataset ₁		Dataset ₂	
		Acc* (%)	BLEU	Acc (%)	BLEU
Interactive	GPT3.5	75.0	0.814	63.3	0.916
	GPT4	94.4	0.911	78.3	0.978
	Gemini	88.9	0.884	70.4	0.960
Fine-tuning	Flan-t5	8.3	0.181	90.8	0.987
	Qwen3-0.6B	22.2	0.246	97.5	0.993
End-to-End	Deepseek R1	63.9	0.546	–	–
	+Self-Refine	63.9	0.490	–	–
Unsupervised	CriticNL2LTL (ours)	<u>80.6</u>	<u>0.517</u>	<u>90.2</u>	<u>0.945</u>

Table 1: Comparative Performance Analysis

As mentioned above, cognitive decision-making is about dynamically matching relevant environmental features to the production rules and applying their effects. Some states may fire multiple production rules, thus leading to many reasoning branches. To diminish the ambiguity, the activation probability for each production rule p_i at step t will be determined by calculating its normalized utility values u_i :

$$p(q_t = p_i) = \frac{e^{u_i/\sigma}}{\sum_j e^{u_j/\sigma}} \quad (1)$$

where σ is a super parameter. During the cognitive agent training process, we are mainly concerned about whether its final decision aligns with the actual human decision in the same situations. So, each time the model outputs a same result as actual human decision, its reasoning chain receives a reward R_+ , otherwise a negative reward R_- is assigned, the reward will be decomposed into all activated production rules on the previous reasoning path:

$$r_i = R - \beta\Delta t \quad (2)$$

where Δt represents the time gap between the activation of production rule p_i and the moment the reward is received, while β denotes the decay coefficient, reflecting that production rules triggered earlier receive smaller portions of the reward. Subsequently, each production rule’s utility is updated based on its allocated reward:

$$u_i(n) = u_i(n-1) + \alpha[r_i(n) - u_i(n-1)] \quad (3)$$

where n denotes the number of iterations, and α denotes the learning rate. If the reward a production rule receives exceeds its current utility, the utility will be increased, thereby enhancing its probability of being activated under similar conditions. Conversely, if the reward is lower, the utility will be reduced.

LTL and Production Rules

In this section, we will clarify how we ensure the formalization of LTL to production rules.

LTL is built upon classical propositional logic, but with temporal operators introduced, it can also describe sequences of events and temporal properties. Unlike CTL (Computation Tree Logic), LTL only focuses on linear time

evolutions along individual execution paths. The syntax of LTL is as follows:

$$\varphi ::= true \mid a \mid \varphi_1 \wedge \varphi_2 \mid \neg\varphi \mid \bigcirc\varphi \mid \varphi_1 \cup \varphi_2 \quad (4)$$

where *true* refers to a logical constant that is always satisfied, *a* refers to an atomic proposition representing a basic statement about the system, \wedge stands for conjunction operator, \neg stands for negation operator, \bigcirc is a temporal operator which means that the proposition must be true in the next state or moment in time, $\varphi_1 \cup \varphi_2$ expresses that φ_1 must remain true continuously until φ_2 becomes true at some point in the future.

Production rules, on the other hand, are commonly described in the “*IF* < condition > *THEN* < action >” format, expressing a simple conditional logic: whenever the condition is met, the associated action is triggered. With that, the production rules can be seen as a subset of LTL. An LTL formula of “*G*($a \rightarrow b$)”, describing the logic of “whenever a holds, b holds as well”, can be formalized as “*IF* *a* *THEN* *b*”, describing the same logic.

But not all LTL formulas can be formalized as production rules. The operators like “*F*” (finally), “*U*” (until) can’t find their synonyms in production rules. The meaning of such temporal operators go beyond what production rules can represent. If a segment of NL text is formalized into LTL formula containing these temporal operators, it implies that the text involves overly complex logic and is therefore not suitable for further formalization into production rules. So, in our work, the LTL formulas using such operators will be marked as **Inference Error**, and only LTL formulas in the format of “*G*(< propositions > \rightarrow < propositions >” are formalized into production rules.

Introducing LTL as an intermediate representation offers several benefits: 1) It transforms the unstructured natural language descriptions of human experience into well-defined logical formula, making it easier to extract and formalize corresponding production rules. 2) It allows for explicit representation of complex logic that can’t be expressed as production rules, preventing inappropriate results. 3) LTL abstracts away from domain-specific actions, it can serve as a reusable, task-agnostic layer that facilitates generalization across different tasks or scenarios.

Experiment

To validate the proposed NL2CA framework, we need to examine its capabilities in two aspects:

- Can CriticNL2LTL framework extract accurate logic information from natural language texts? (see section **NL-to-LTL Experiment**)
- Can NL2CA construct a cognitive agent which is capable of imitating human behaviors? (see section **Cognitive Driving Simulation**)

NL-to-LTL Experiment

We conduct NL-to-LTL experiments on two datasets: *Dataset₁*, a challenging open-source benchmark for neural translation approach (Cosler et al. 2023), and *Dataset₂*, a larger-scale public synthetic dataset (Xu, Feng, and Miao 2024), containing samples across eight LTL expression patterns and various domains.

We fine-tune a Qwen3-0.6B model on the training samples on *Dataset₂* to generate initial LTL translations, which are then refined using our Critic Tree constructed with Deepseek models (number of critics = 2, max depth = 2), the revisor is constructed using DeepseekR1. For critics, each critic is constructed with a 50% probability based on DeepseekR1 and a 50% probability based on DeepseekV3. Performance is assessed using the match accuracy (*ACC*) and *BLEU* score. For *Dataset₁*, due to its smaller scale and high level of difficulty, accuracy is evaluated through manual semantic verification (denoted as *ACC**). To address generality, we only use the *minimal* prompt that only includes minimal domain knowledge (details in Appendix). In practice, it can be replaced with more domain-specific content to further boost the framework’s performance.

The experimental results are shown in Table 1. For LLM-based methods, the performance on *Dataset₂* is measured on 143 samples randomly picked from the test dataset.

Our approach achieves high consistency across both datasets, outperforming purely fine-tuned and end-to-end approaches on *Dataset₁*, while maintaining competitive results on *Dataset₂*. Interactive approaches that benefit from human oversight perform outstandingly on *Dataset₁* but show a significant performance drop on *Dataset₂* due to domain shift. Notably, conventional self-Refine (Madaan et al. 2023) fails to improve the performance of end-to-end baselines, highlighting the effectiveness of our proposed Critic Tree structure.

Cognitive Driving Simulation

To verify NL2CA’s capability of constructing proper cognitive models from human experience texts, we conduct several cognitive driving simulation experiments. We construct 3 representative driving scenarios including highway, signalized intersection and lane changing interference using CARLAUE4, as shown in Figure 2:

- **Highway:** A three-lane highway scenario, there is a group of vehicles ahead of the ego vehicle occupying all three lanes. During the driving process, an NPC vehicle to the right of the ego-vehicle will attempt to cut in to the ego-vehicle’s lane.

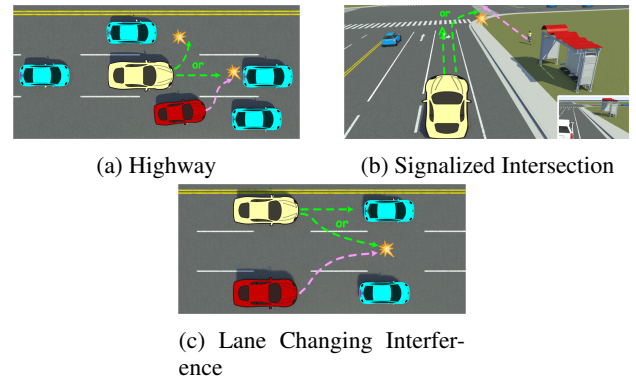


Figure 2: Examples of Driving Simulation Scenarios.

Parameter	Description	Value
α	Learning rate	2e-4
β	Attenuation coefficient	0.01
σ	Utility coefficient	$\sqrt{2}$
u_0	Initial utility of production	0
R_+	Positive reward	10
R_-	Negative reward	0

Table 2: Experimental Parameters

- **Signalized Intersection:** A signalized intersection scenario, the ego-vehicle needs to make a right turn through the intersection. During the process, a pedestrian NPC will attempt to cross the road, hindering the ego-vehicle.
- **Lane Changing Interference:** A three-lane highway scenario, the ego-vehicle is driving in the leftmost lane, with a slow-driving NPC vehicle ahead. If the ego-vehicle attempts to overtake, an NPC vehicle in the rightmost lane will attempt to make a left lane change.

These scenarios are carefully designed to represent corner cases in urban and highway driving, where decision-making is highly sensitive to subtle environmental changes and driver intent. By focusing on these cases, we aim to assess whether the cognitive agents constructed through NL2CA can capture a general human reasoning patterns.

We also invite 34 different drivers and conduct about 70 trials in total for each virtual scenario using Virtual-Reality device. They are asked to drive from the predefined start point to the destination, along with complying with the traffic rules and avoiding any potential collisions. During the driving process, their control input (e.g., throttle, steering angle, etc.) and the environmental state information are recorded as the actual human data used for cognitive reinforcement learning. After the driving process, we interview each driver about the driving decisions they made and the corresponding reasons, thus collecting a set of human driving experience texts for logic information extraction.

We adopt the same parameters as in the NL-to-LTL experiment to formalize production rules from NL text. Each interview transcript segment is first input into CriticNL2LTL to get an intermediate LTL representation. The resulted LTL

Scenario	Model	Success Rate (%) \uparrow	RSR (%) \uparrow	ACR (%) \uparrow	AST (s)	ALD (cm) \downarrow
Highway	Human Data	53.73	–	83.98	27.79	27.86
	Manual	25.00	98.67	89.06	26.29	25.78
	M_{lit}	58.70	87.62	82.02	26.41	9.43
	M_{sup}	12.00	93.07	77.58	22.29	10.70
Signalized Intersection	Human Data	84.13	–	99.56	16.08	31.93
	Manual	75.00	100.00	92.41	10.19	22.99
	M_{lit}	85.71	99.46	99.84	7.25	20.15
	M_{sup}	76.60	99.62	99.77	8.83	22.26
Lane Changing Interference	Human Data	85.92	–	97.59	28.71	39.88
	Manual	95.92	99.73	97.81	22.42	25.54
	M_{lit}	91.30	100	89.73	30.87	35.05
	M_{sup}	63.82	100	89.70	19.19	21.65

Table 3: Comparative Driving Simulation Analysis




Figure 3: Logic Extraction and Formalization Result.

formula is then fed to another LLM and transformed into executable production rule with grounded variables.

Following the standard cognitive driving pipeline, we divide the driving decision-making process into longitudinal and lateral decision-makings. For each intermediate LTL representation, it is formalized into two production rules describing longitudinal and lateral decision-making processes. Moreover, since not every text segment contains complete descriptions of both decision-making processes, we also prompt the LLM to output “pass” when it determines that the decision-related information in an LTL formula is incomplete or missing.

After the extraction of logical information and its formalization into production rules, we also analyze the results of our methodology on the interview transcripts across three driving scenarios. To clarify the results, we define the following four types of formalization results and recorded their frequencies across different scenario datasets:

- **Format Mismatch:** The formalized production rule cannot be parsed by the inference engine.
- **Duplicated Content:** The generated production rule contains the same content as an existing production rule.

- **Inference Error:** The generated production rule contains undefined variables or operators.
- **Viable:** None of the above three types of errors appear in the generated production rule.

As shown in Figure 3, most formalized production rules are viable. The suffixes “supply” and “literal” refer to two different generation methods, which will be clarified later.

Using the generated production rules, we construct cognitive driving agents for 3 scenarios and train them through the cognitive reinforcement learning pipeline with the collected human driving data. The key of this training is to update the utility of each production in the driving models. To facilitate the training, the inputs of the model are precisely annotated scene elements instead of original sensory data such as images and point clouds, and the outputs are the control signals of steering wheel, throttle and brake. A positive reward R_+ is given when the model predicts an action consistent with human choice, otherwise a negative reward R_- is assigned. More detailed parameter setting is shown in Table 2.

To distinguish our methodology from current approaches that directly query the LLM how to react under unknown situations, we also collect a set of production rules by prompting the LLM to actively infer and fill in potentially missing environmental information in the interview transcripts. This was done to explore whether further incorporating the LLM’s prior knowledge to discover hidden logic could improve the automated cognitive modeling. Such approach is noted as M_{sup} , and the original approach of formalizing production rules literally from the interview transcripts is noted as M_{lit} . The cognitive driving models constructed using both methods go through the same cognitive reinforcement learning procedure on the actual human datasets.

After optimizing the cognitive driving models, we evaluate the performance of each model against the test results collected from human subjects in the three driving scenarios. We measure key driving performance metrics including RSR (reasoning success rate), ACR (average completion rate), AST (average simulation time) and ALD (average lane deviation), reported metrics are the average results for 50 test runs. For more comparative results, test experiments using




Figure 4: JS Divergence Curves of Cognitive Models.

a manually designed interpretable cognitive driving baseline model (Li et al. 2024) are also carried out, shown in Table 3.

As shown in the experimental results, M_{lit} demonstrates competitive performance against the manually constructed baseline that is incorporated with experts’ prior knowledge, but M_{sup} shows no advantage against these two methods. The manual baseline, benefiting from experts’ dedicated design, exhibits high reasoning completeness and thus holds a distinct advantage in RSR metric. Across all three test scenarios, it achieves an average RSR of 99.47%, along with highest average ACR of 93.09%. In contrast, M_{lit} achieves an average RSR of 95.69% and an average ACR of 90.53%. However, the manual baseline still shows a significant gap compared to M_{lit} in terms of Success Rate(SR) and ALD. Specifically, M_{lit} achieves the highest average SR of 78.57% and an average ALD of 21.54cm, while the manual baseline scores 65.30% and 24.77cm, respectively.

To fully evaluate the cognitive models’ alignment with human behaviors, we further computed the Jensen-Shannon (JS) Divergence between the decision distributions P generated by each cognitive model and the corresponding distributions Q observed in actual human behaviors. The definition of JS divergence is as follows:

$$JS(P||Q) = \frac{1}{2} KL(P||M) + \frac{1}{2} KL(Q||M) \quad (5)$$

where KL refers to Kullback-Leibler (KL) Divergence and $M = (P + Q)/2$. Specifically, we select the top 10 environmental inputs with the largest number of human decision samples. For each of these inputs, the cognitive decision models are executed n times to generate the corresponding decision distributions. n denotes the number of samples with a specific environmental input. During the training process of the cognitive models, we continuously track the corresponding JS Divergence. The results are shown in Figure 4.

As shown in the results. M_{lit} demonstrates the best alignment with human behavior by achieving the lowest JS Divergence on average, while performing a decent learning capacity. In contrast, the JS Divergence of the manual baseline shows limited change as the number of training iterations increases in Signalized Intersection and Lane Changing Interference scenarios. This indicates that the performance of

manual baseline relies more on experts’ subjective design, rather than learning from human driving behavior.

M_{sup} demonstrates a certain degree of learning capacity. However, its ability to approximate human decision distributions remains relatively limited. This may be attributed to the fact that the preconditions of its production rules are extended by the LLM, potentially resulting in Over-constraining. Such constraints can lead to a reduced number of decision branches of the cognitive model (fewer decision options for a given environmental input), limiting the model’s capacity to fit the human decision distributions.

Discussion

We propose a method to automatically build cognitive decision rules from natural language texts without human intervention, which can benefit both AI and cognitive computing. For AI, it offers a feasible approach to auto-formalize unstructured texts to formal logic, providing an alternative paradigm different from reasoning all relative texts through a “black box” network, yielding traceable and understandable decision-makings to avoid the hallucination of LLMs caused by their probabilistic nature (Wang et al. 2023) (Zhang et al. 2023). For cognitive computing, this method opens a path for low labor-cost automatic modeling based on human natural language, which can help investigating human performances in various human-machine interactive tasks, such as cooperative driving or other critical interactive systems.

Conclusion

This paper presents a framework for automatically formalizing cognitive decision-making from human experience documents and constructing cognitive models for imitation of human behaviors. We demonstrate how this framework could formalize production rules from texts and be applied to driving simulations. This work provides brand new insights on automated cognitive modeling based on NL data.

Acknowledgments

This work is supported in part by National Natural Science Foundation of China under Grant 62476270 and Grant T2192933.

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Anderson, J. R.; and Lebiere, C. J. 2014. *The atomic components of thought*. Psychology Press.
- Bajaj, G.; Pearce, K.; Kennedy, S.; Larue, O.; Hough, A.; King, J.; Myers, C.; and Parthasarathy, S. 2023. Generating chunks for cognitive architectures. In *Proceedings of the AAAI Symposium Series*, volume 2, 246–252.
- Browne; Clarke; Dill; and Mishra. 1986. Automatic verification of sequential circuits using temporal logic. *IEEE Transactions on Computers*, 100(12): 1035–1044.
- Chater, N.; Misyak, J.; Watson, D.; Griffiths, N.; and Mouzakis, A. 2018. Negotiating the traffic: Can cognitive science help make autonomous vehicles a reality? *Trends in cognitive sciences*, 22(2): 93–95.
- Cosler, M.; Hahn, C.; Mendoza, D.; Schmitt, F.; and Trippel, C. 2023. nl2spec: Interactively translating unstructured natural language to temporal logics with large language models. In *International Conference on Computer Aided Verification*, 383–396. Springer.
- Fang, Y.; Jin, Z.; An, J.; Chen, H.; Chen, X.; and Zhan, N. 2025. Enhancing Transformation from Natural Language to Signal Temporal Logic Using LLMs with Diverse External Knowledge. *arXiv preprint arXiv:2505.20658*.
- Kirk, J. R.; Wray, R. E.; and Laird, J. E. 2023. Exploiting language models as a source of knowledge for cognitive agents. In *Proceedings of the AAAI Symposium Series*, volume 2, 286–294.
- Kirk, J. R.; Wray, R. E.; Lindes, P.; and Laird, J. E. 2024. Improving knowledge extraction from llms for task learning through agent analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 18390–18398.
- Kumar, K.; Ashraf, T.; Thawakar, O.; Anwer, R. M.; Cholakkal, H.; Shah, M.; Yang, M.-H.; Torr, P. H.; Khan, F. S.; and Khan, S. 2025. Llm post-training: A deep dive into reasoning large language models. *arXiv preprint arXiv:2502.21321*.
- Laird, J. E. 2022. Introduction to soar. *arXiv preprint arXiv:2205.03854*.
- Li, Y.; Qi, H.; Zhu, F.; Lv, Y.; and Ye, P. 2024. Interpretable Autonomous Driving Model Based on Cognitive Reinforcement Learning. In *2024 IEEE Intelligent Vehicles Symposium (IV)*, 515–520. IEEE.
- Liu, A.; Feng, B.; Xue, B.; Wang, B.; Wu, B.; Lu, C.; Zhao, C.; Deng, C.; Zhang, C.; Ruan, C.; et al. 2024. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*.
- Madaan, A.; Tandon, N.; Gupta, P.; Hallinan, S.; Gao, L.; Wiegrefe, S.; Alon, U.; Dziri, N.; Prabhunoye, S.; Yang, Y.; et al. 2023. Self-refine: Iterative refinement with self-feedback. *Advances in Neural Information Processing Systems*, 36: 46534–46594.
- Maierhofer, S.; Moosbrugger, P.; and Althoff, M. 2022. Formalization of intersection traffic rules in temporal logic. In *2022 IEEE Intelligent Vehicles Symposium (IV)*, 1135–1144. IEEE.
- Niu, Q.; Liu, J.; Bi, Z.; Feng, P.; Peng, B.; Chen, K.; Li, M.; Yan, L. K.; Zhang, Y.; Yin, C. H.; et al. 2024. Large language models and cognitive science: A comprehensive review of similarities, differences, and challenges. *arXiv preprint arXiv:2409.02387*.
- Oh, H.; Yun, Y.; and Myung, R. 2024. Driver behavior and mental workload for takeover safety in automated driving: ACT-R prediction modeling approach. *Traffic injury prevention*, 25(3): 381–389.
- Qi, H.; Hou, E.; and Ye, P. 2024. Cognitive reinforcement learning: An interpretable decision-making for virtual driver. *IEEE Journal of Radio Frequency Identification*, 8: 627–631.
- Rawte, V.; Chakraborty, S.; Pathak, A.; Sarkar, A.; Tonmoy, S. I.; Chadha, A.; Sheth, A.; and Das, A. 2023. The troubling emergence of hallucination in large language models-an extensive definition, quantification, and prescriptive remediations. Association for Computational Linguistics.
- Ritter, F. E.; Tehranchi, F.; and Oury, J. D. 2019. ACT-R: A cognitive architecture for modeling cognition. *Wiley Interdisciplinary Reviews: Cognitive Science*, 10(3): e1488.
- Salvucci, D. D. 2006. Modeling driver behavior in a cognitive architecture. *Human factors*, 48(2): 362–380.
- Villalobos, P.; Ho, A.; Sevilla, J.; Besiroglu, T.; Heim, L.; and Hobbhahn, M. 2024. Position: Will we run out of data? Limits of LLM scaling based on human-generated data. In *Forty-first International Conference on Machine Learning*.
- Wang, J.; Sundarsingh, D. S.; Deshmukh, J. V.; and Kantaros, Y. 2025. ConformalNL2LTL: Translating Natural Language Instructions into Temporal Logic Formulas with Conformal Correctness Guarantees. *arXiv preprint arXiv:2504.21022*.
- Wang, W.; Xie, J.; Hu, C.; Zou, H.; Fan, J.; Tong, W.; Wen, Y.; Wu, S.; Deng, H.; Li, Z.; et al. 2023. Drivemlm: Aligning multi-modal large language models with behavioral planning states for autonomous driving. *arXiv preprint arXiv:2312.09245*.
- Wray, R. E.; Kirk, J. R.; and Laird, J. E. 2025. Applying Cognitive Design Patterns to General LLM Agents. *arXiv preprint arXiv:2505.07087*.
- Xu, Y.; Feng, J.; and Miao, W. 2024. Learning from failures: Translation of natural language requirements into linear temporal logic with large language models. In *2024 IEEE 24th International Conference on Software Quality, Reliability and Security (QRS)*, 204–215. IEEE.
- Zhang, H.; Song, H.; Li, S.; Zhou, M.; and Song, D. 2023. A survey of controllable text generation using transformer-based pre-trained language models. *ACM Computing Surveys*, 56(3): 1–37.
- Zhu, F.; and Simmons, R. 2024. Bootstrapping cognitive agents with a large language model. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 655–663.