



HAL
open science

Visual Large Language Models Exhibit Human-Level Cognitive Flexibility in the Wisconsin Card Sorting Test

Guangfu Hao, Frédéric Alexandre, Shan Yu

► **To cite this version:**

Guangfu Hao, Frédéric Alexandre, Shan Yu. Visual Large Language Models Exhibit Human-Level Cognitive Flexibility in the Wisconsin Card Sorting Test. *IEEE Transactions on Cognitive and Developmental Systems*, 2025, pp.1 - 11. <10.1109/tcds.2025.3590165>. <hal-05202743>

HAL Id: hal-05202743

<https://inria.hal.science/hal-05202743v1>

Submitted on 7 Aug 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY-NC-ND 4.0 - Attribution - Non-commercial use - No Derivative Works - International License

Visual Large Language Models Exhibit Human-Level Cognitive Flexibility in the Wisconsin Card Sorting Test

Guangfu Hao[✉], Frederic Alexandre[✉], Shan Yu[✉]

Abstract—Cognitive flexibility has been extensively studied in human cognition but remains relatively unexplored in the context of Visual Large Language Models (VLLMs). This study assesses the cognitive flexibility of state-of-the-art VLLMs (GPT-4o, Gemini-1.5 Pro, and Claude-3.5 Sonnet) using the Wisconsin Card Sorting Test (WCST), a classic measure of set-shifting ability. Our results reveal that VLLMs achieve or surpass human-level set-shifting capabilities under chain-of-thought prompting with text-based inputs. However, their abilities are highly influenced by both input modality and prompting strategy. In addition, we find that through role-playing, VLLMs can simulate various functional deficits aligned with patients having impairments in cognitive flexibility, suggesting that VLLMs may possess a cognitive architecture, at least regarding the ability of set-shifting, similar to the brain. This study reveals the fact that VLLMs have already approached the human level on a key component underlying our higher cognition, and highlights the potential to use them to emulate complex brain processes.

Index Terms—Cognitive flexibility, Wisconsin Card Sorting Test, Visual large language models, Prefrontal cortex, Prompting strategy, Cognitive impairment simulation.

I. INTRODUCTION

COGNITIVE flexibility, a key component of executive function, is fundamental to human adaptability and problem-solving [1], [2]. This ability to shift between mental sets or strategies in response to changing environmental demands is crucial for everyday functioning and has been extensively studied in cognitive psychology research [3]. The prefrontal cortex (PFC) is known to be central to this cognitive process [4], [5], facilitating goal-directed behavior and controlled processing.

Recent advancements in artificial intelligence (AI), particularly in visual large language models (VLLMs) [6], have sparked a growing need to assess these systems' cognitive abilities using paradigms analogous to human cognitive assessment [7]–[9]. State-of-the-art VLLMs such as GPT-4o [10], Gemini-1.5 Pro [11], and Claude-3.5 Sonnet [12] have demonstrated remarkable capabilities in processing and interpreting both textual and visual information, excelling in tasks that demand complex reasoning and contextual understanding.

This work was supported in part by the Strategic Priority Research Program of the Chinese Academy of Sciences (CAS)(XDB1010302), CAS Project for Young Scientists in Basic Research, Grant No. YSBR-041 and the International Partnership Program of the Chinese Academy of Sciences (CAS) (173211KY5B20200021). (Corresponding author: shan.yu@nlpr.ia.ac.cn).

Guangfu Hao and Shan Yu are with the Laboratory of Brain Atlas and Brain-inspired Intelligence, Institute of Automation, Chinese Academy of Sciences, Beijing, China, and also with the School of Artificial Intelligence, University of Chinese Academy of Sciences, Beijing, China.

Frederic Alexandre is with the Inria centre at the university of Bordeaux, Neurodegenerative Diseases Institute, Bordeaux, France (e-mail: frederic.alexandre@inria.fr).

Despite these achievements, the extent to which VLLMs exhibit cognitive flexibility, especially in tasks requiring set-shifting and adaptation to changing rules, remains largely unexplored. While these models have demonstrated impressive performance across diverse tasks, their ability to flexibly adapt to changing environmental demands has not been systematically evaluated. This gap in our understanding is particularly significant given the increasing integration of VLLMs into complex, dynamic real-world environments where adaptability is crucial.

The Wisconsin Card Sorting Test (WCST), developed in the 1940s [13] and refined over decades, has emerged as the gold standard for assessing cognitive flexibility in both clinical and research settings [14], [15]. Originally developed to evaluate PFC function, it requires participants to discover sorting rules based on feedback and then flexibly shift to new rules when the criteria change (see Figure 1). The test's sensitivity to PFC function has been consistently demonstrated through lesion studies [16], neuroimaging research [17], and clinical observations, cementing its status as a crucial tool in understanding the cognitive flexibility. While other measures of cognitive flexibility exist, the WCST's established validity make it a core benchmark for evaluating this fundamental cognitive capacity.

This research aims to evaluate the cognitive flexibility in VLLMs using the WCST paradigm and investigate how different input modalities (image-based vs. text-based), prompting strategy (direct vs. chain-of-thought reasoning) and rule description specificity affect their performance. Additionally, we explore the potential of VLLMs to simulate specific patterns of cognitive impairment through role-playing, which enables us to understand human cognitive architecture. By comparing VLLM performance across varied conditions, we aim to elucidate their cognitive flexibility and inherent limitations. This investigation not only advances our understanding of VLLMs but also offers insights into the nature of cognitive flexibility itself.

II. RELATED WORK

A. Cognitive Flexibility and Assessment Methods

Neuroimaging studies have consistently implicated the PFC in cognitive flexibility tasks. The dorsolateral prefrontal cortex (DLPFC) and anterior cingulate cortex (ACC) play critical roles in set-shifting, with the DLPFC maintaining and updating task rules, and the ACC involved in conflict monitoring and error detection [18]. The fronto-parietal network, encompassing these regions, dynamically reconfigures during flexibility-demanding tasks [19]. Cognitive flexibility is closely interre-

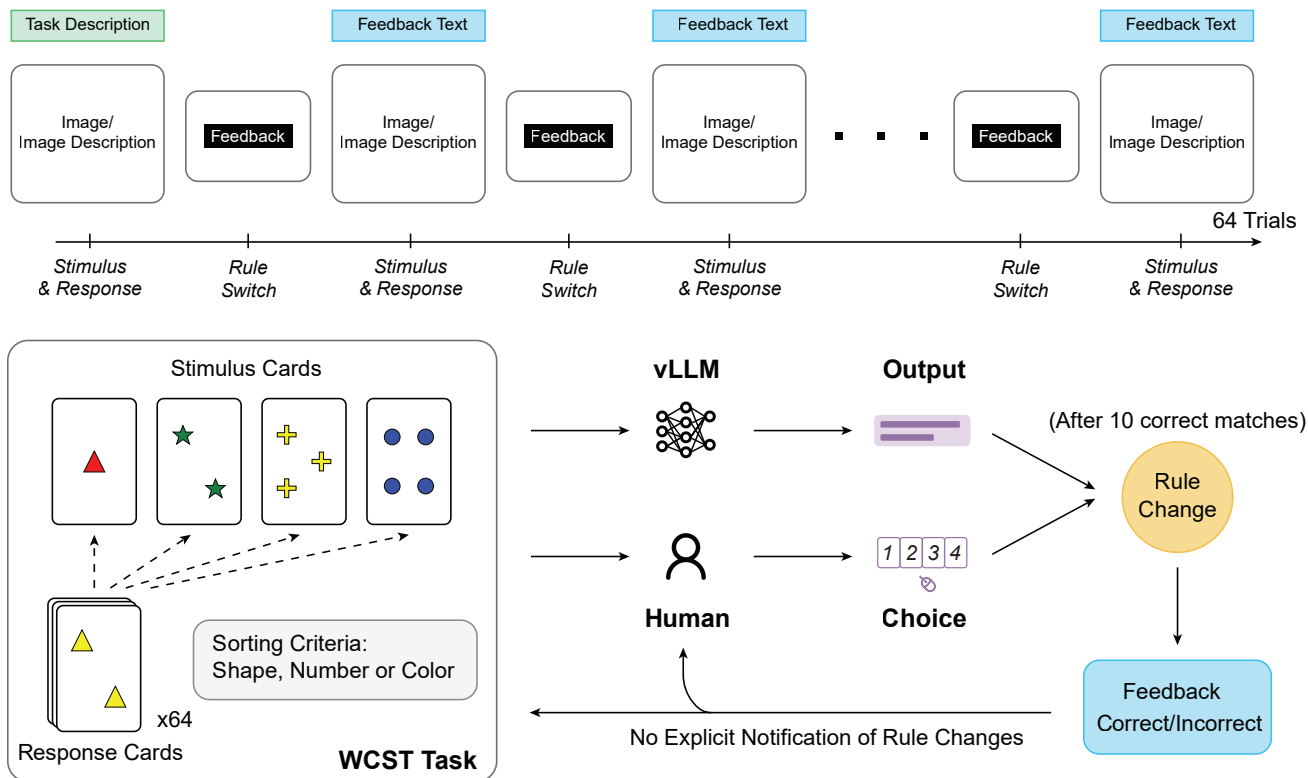


Fig. 1. **WCST Procedure and Sample Stimuli.** The WCST consists of matching response cards to four stimulus cards based on a sorting rule (color, shape, or number) that changes periodically. Participants receive feedback on the accuracy of each match but are not explicitly told the sorting rule or when it changes.

lated with other executive functions: working memory maintains task-relevant information and goals [3], while inhibitory control suppresses previous cognitive sets when rules change [20].

Several tasks have been developed to assess cognitive flexibility in humans, with the WCST being a widely recognized measure of set-shifting ability [21]. The WCST's sensitivity to PFC dysfunction has been extensively validated [22]. Complementary assessments include the Dimensional Change Card Sort (DCCS) task for children [23] and the computerized Intra-Extra Dimensional Set Shift (IED) subtest of the Cambridge Neuropsychological Test Automated Battery (CANTAB) [24], offering targeted measures across different populations and modalities.

B. Multifaceted Evaluation of LLMs

Recent studies have employed diverse assessments to evaluate large language models (LLMs) and VLLMs across various domains and tasks. Models like GPT-4 demonstrated human-level or superior performance on most theory of mind tests [25]. Similarly, research on human creativity found that ChatGPT-assisted ideas were more creative compared to those generated without LLM assistance [26]. However, challenges persist in other areas. The Test of Time (ToT) benchmark exposed difficulties with complex temporal reasoning tasks, particularly those requiring multi-fact integration and intricate

arithmetic operations [27]. Despite strong performance on high-level vision tasks, state-of-the-art VLLMs struggled with basic geometric tasks that are straightforward for humans [28]. A neuropsychological investigation revealed a discontinuous profile in ChatGPT's prefrontal functioning, with performance ranging from superior to impaired across different cognitive tasks [29]. To address the multifaceted nature of artificial intelligence, researchers have proposed new evaluation frameworks. A comprehensive framework for artificial general intelligence (AGI) tests inspired by cognitive science emphasizes the need for multidimensional intelligence assessment [9]. Additionally, the concept of Turing Experiments (TEs) was introduced as a method for evaluating LLMs' ability to simulate human behavior in experimental settings [30].

C. Prompting Strategies

Prompting strategies significantly influence the performance of LLMs [31]. The simplest approach, "Straight-to-Answer" (STA), directly queries the model without additional context. While effective for straightforward tasks, STA often falters on complex problems requiring multi-step reasoning. Chain-of-Thought (CoT) prompting encourages step-by-step reasoning [32], substantially improving performance on complex reasoning tasks [33]. Variations such as zero-shot CoT [34] and self-consistency CoT [35] have further refined this approach, adapting it to scenarios with limited or no task-specific examples. In

multimodal contexts, visual CoT have extended these concepts to VLLMs [36], demonstrating the potential for improved reasoning in tasks that combine textual and visual information. Other task-specific strategies, such as least-to-most prompting address challenges of easy-to-hard generalization [37], while meta-prompting and automatic prompt engineering techniques aim to optimize the prompts themselves [38], [39].

III. METHOD

A. Models and Experimental Procedure

This study focuses on three state-of-the-art VLLMs: GPT-4o, Gemini-1.5 Pro, and Claude-3.5 Sonnet. These models represent the current pinnacle of multimodal LLMs capabilities, demonstrating proficiency in processing both textual and visual inputs (see Table V for details). We employ a standard version of the WCST-64 [15] to assess the cognitive flexibility of VLLMs. Our experimental design incorporates a 2x2 factorial structure, manipulating input modality (Visual Input (VI) / Textual Input (TI)) and prompting strategy (Straight to Answer (STA) / Chain of Thought (CoT)) to comprehensively evaluate VLLMs performance. This design resulted in four experimental conditions: STA-VI, STA-TI, CoT-VI, and CoT-TI. Each VLLM was tested independently across all four conditions, with 10 repetitions per condition. The arrangement of stimulus cards and the sequence of sorting rules were randomized for each repetition.

Algorithm 1 WCST for VLLMs

```

1: Initialize  $r \in \{\text{color, shape, number}\}$ ,  $c \leftarrow 0$ ,  $s \leftarrow 0$ ,
    $t \leftarrow 0$ 
2: Inform model: "{Task description}"
3: while  $t < 64$  do
4: Present card and prompt for sorting (image/text for
   VI/TI, using STA/CoT)
5: Record and parse model's response to extract selection

6: if selection is correct then
7:    $c \leftarrow c + 1$ 
8:   if  $c = 10$  then
9:     Change active rule  $r$ 
10:     $s \leftarrow s + 1$ 
11:     $c \leftarrow 0$ 
12:   end if
13: else
14:    $c \leftarrow 0$ 
15: end if
16: Provide feedback on correctness of current selection
17:  $t \leftarrow t + 1$ 
18: end while

```

Additionally, we collected data from 30 cognitively healthy human participants (aged 20-35) as a baseline for comparison. Human participants interacted with a web-based interface designed to replicate the WCST experience while accommodating human response patterns (supplementary Figure s-5). The interface presented cards sequentially and allowed participants to indicate their sorting choices via button presses.

The language used in instructions was carefully adapted to be more intuitive for human subjects while maintaining the essential structure of the task.

We adapted the WCST for use with VLLMs while maintaining its core principles (see Algorithm 1 for the implementation). The test consists of a series of virtual cards, each featuring shapes (circle, cross, triangle, or star) in varying colors (red, green, yellow, or blue) and quantities (one to four). The models are tasked with sorting these cards according to an undisclosed rule (color, shape, or number), which changes periodically without explicit notification. The sorting rule changed after ten consecutive correct categorizations. The assessment concluded after 64 trials. Detailed descriptions of the task instructions are provided in supplementary Figure s-2. Example stimuli for VI and TI conditions, and prompt templates for STA and CoT strategies are provided in supplementary Figure s-3.

Data collection was fully automated using API calls to each VLLM. Model responses were recorded verbatim for each trial. Human participant data was collected through the web-based interface. All participants provided informed consent, and the study was approved by the institutional review board.

To address potential concerns regarding model memorization of the classic WCST paradigm, we also designed a novel variant called the ALIEN Task that preserves the logical structure of the WCST while using entirely different terminology, stimuli, and framing. In this variant, participants explore an alien civilization by categorizing extraterrestrial astronomical systems. The original WCST dimensions were replaced with thematically distinct alternatives: shape was replaced with planetary orbit types (spiral, elliptical, circular, Z-shaped); color was replaced with atmospheric composition (hydrogen/blue, helium/yellow, nitrogen/purple, oxygen/green); and count was replaced with number of moons (1, 2, 3, 4). The underlying logical structure and rule-switching mechanisms remained identical to the WCST. We implemented this variant under STA-TI and CoT-TI conditions to examine whether performance patterns would persist with entirely novel surface features. The specific instructions and example stimuli for the ALIEN Task are detailed in supplementary Figure s-4.

B. Evaluation Metrics

Performance was primarily assessed using the following metrics which were chosen for their ability to quantify different aspects of cognitive flexibility:

Categories Completed (CC): The number of categories (sets of 10 consecutive correct sorts) completed.

$$CC = \sum_{i=1}^n I(c_i = 10) \quad (1)$$

where n is the total number of trials, c_i is the number of consecutive correct sorts at trial i , and $I(\cdot)$ is the indicator function.

Perseverative Errors (PE): The number of errors where the model persisted with a previously correct but currently incorrect rule.

$$PE = \sum_{i=1}^n I(r_i = r_{\text{prev}} \wedge r_i \neq r_{\text{current}}) \quad (2)$$

where r_i is the rule used by the model at trial i , r_{prev} is the previously correct rule, and r_{current} is the current correct rule.

Non-Perseverative Errors (NPE): All errors that are not perseverative.

$$NPE = \text{Total Errors} - PE \quad (3)$$

NPE captures non-perseverative errors, potentially indicating exploration or random mistakes.

Trials to First Category (TFC): The number of trials required to complete the first category, indicating how quickly the model can deduce and consistently apply the first sorting rule.

$$TFC = \min\{i : c_i = 10\} \quad (4)$$

where i is the trial number and c_i is as defined in CC.

Conceptual Level Responses (CLR): The percentage of responses occurring in runs of three or more correct sorts, indicating conceptual understanding.

$$CLR = \frac{\sum_{i=1}^n I(c_i \geq 3)}{n} \times 100\% \quad (5)$$

where $I(c_i \geq 3)$ is an indicator function that equals 1 if the number of consecutive correct sorts up to and including trial i is 3 or more, and 0 otherwise.

Failure to Maintain Set (FMS): The number of times the model makes an error after five or more consecutive correct sorts but before completing a category.

$$FMS = \sum_{i=1}^{n-1} I(5 \leq c_i < 10) \cdot I(c_{i+1} = 0) \quad (6)$$

where i is the trial number and c_i is as defined in CC.

These metrics collectively provide a comprehensive view of the VLLMs' cognitive flexibility [40], capturing various aspects such as rule learning, set-shifting, perseveration, and conceptual understanding.

IV. RESULTS

A. WCST Task Performance Across Models and Conditions

The WCST performance of GPT-4o, Gemini-1.5 Pro, and Claude-3.5 Sonnet exhibited marked variations across the four conditions (Figure 2). Their cognitive flexibility was measured using the CC metric, standardized on a 0-1 scale, with the human baseline ($\mu = 0.95$, $\sigma = 0.09$). The CoT-TI condition consistently yielded superior outcomes across all VLLMs, followed by CoT-VI, STA-TI, and STA-VI, respectively, underscoring the critical influence of both prompting strategies and input modalities on VLLMs' set-shifting capabilities.

In the STA-VI condition, all VLLMs struggled significantly, with mean performances ranging from 0.02 to 0.04. The transition to STA-TI yielded modest improvements, particularly for Claude-3.5 Sonnet ($\mu = 0.10$, $\sigma = 0.10$). However, the introduction of CoT prompting precipitated a dramatic

performance surge across all models. In the CoT-VI condition, Claude-3.5 Sonnet exhibited remarkable improvement ($\mu = 0.96$, $\sigma = 0.08$), while Gemini-1.5 Pro and GPT-4o also showed substantial gains. This stark contrast between STA and CoT conditions illuminates the pivotal role of explicit reasoning in augmenting VLLMs' cognitive flexibility.

The CoT-TI condition elicited peak performances, with Claude-3.5 Sonnet achieving perfection ($\mu = 1.00$, $\sigma = 0.00$), surpassing even the human baseline. Gemini-1.5 Pro ($\mu = 0.96$, $\sigma = 0.08$) and GPT-4o ($\mu = 0.92$, $\sigma = 0.16$) also demonstrated near-human or human-equivalent performance in this setting. Notably, the performance variability (σ) was generally higher in CoT conditions for Gemini-1.5 Pro and GPT-4o, indicating potential instability in their cognitive processes. The consistent superiority of TI over VI across all conditions suggests a potential advantage in processing textual over visual inputs.

The observed performance gradient, from near-chance levels in STA-VI to human-surpassing in CoT-TI, demonstrates the potential of VLLMs to exhibit human-like cognitive flexibility under appropriate conditions, while also highlighting the critical impact of prompting strategies and input modalities on their performance in tasks requiring set-shifting and rule adaptation.

To investigate whether VLLMs' performance might be influenced by memorization of similar tasks in their training data, we evaluated all three models on our novel ALIEN Task variant under both STA-TI and CoT-TI conditions. The results, presented in Figure s-1 and Table t-1, demonstrate performance patterns remarkably consistent with those observed in the original WCST. Under the STA-TI condition, all models struggled to complete the ALIEN Task, while under the CoT-TI condition, all models demonstrated high performance levels that closely mirrored their achievements on the standard WCST. This consistent performance strongly suggests that the models' capabilities reflect genuine cognitive flexibility rather than memorization of specific task patterns.

B. Detailed Analysis by Evaluation Metric

To offer a comprehensive assessment of the VLLMs' performance on the WCST, we analyzed six key metrics outlined in the previous section. Table I presents the mean scores and standard deviations across all evaluation metrics for each VLLM and condition. This analysis reveals distinct patterns in cognitive flexibility and set-shifting abilities among the models.

PE were most prevalent in the STA-TI condition for all models, with Claude-3.5 Sonnet showing the highest number of errors in this condition ($\mu = 15.90$, $\sigma = 17.53$). In the STA-VI condition, PE were relatively low for all models, as they largely failed to follow the rules at all. However, the transition to CoT conditions reduced PE, with Claude-3.5 Sonnet demonstrated the lowest number of PE in the CoT-TI condition ($\mu = 6.30$, $\sigma = 0.82$). This suggests that Claude-3.5 Sonnet may surpass human performance in adapting to changing rules, especially when provided with explicit reasoning prompts and textual descriptions.

NPE showed a dramatic reduction from STA to CoT conditions across all models, with improvements observed in the

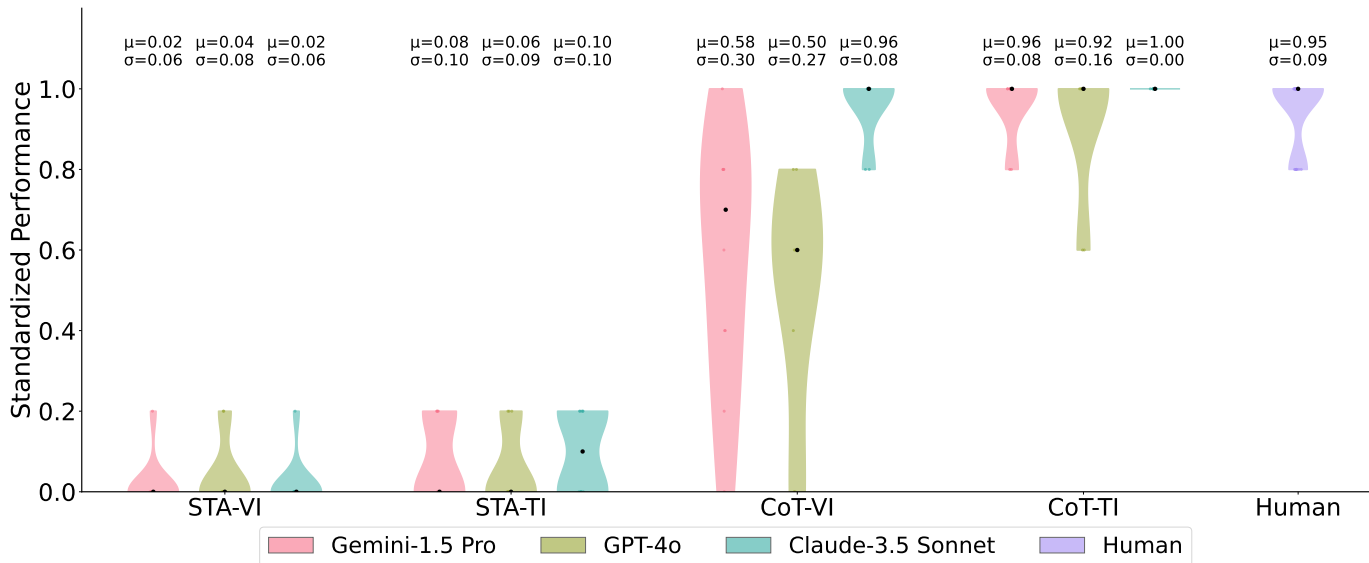


Fig. 2. **WCST Task Performance Across Models and Conditions.** The distribution of standardized Categories Completed (CC) scores for GPT-4o, Gemini-1.5 Pro, and Claude-3.5 Sonnet under four experimental conditions: STA-VI (Straight to Answer - Visual Input), STA-TI (Straight to Answer - Textual Input), CoT-VI (Chain of Thought - Visual Input), and CoT-TI (Chain of Thought - Textual Input), with the human baseline performance.

TABLE I
WCST PERFORMANCE METRICS ACROSS MODELS AND EXPERIMENTAL CONDITIONS

Model	Condition	CC	PE	NPE	TFC	CLR (%)	FMS
Gemini-1.5 Pro	STA-VI	0.10 (0.32)	1.70 (5.38)	40.00 (9.51)	11.00 (-)	5.00 (6.15)	0.20 (0.42)
	STA-TI	0.40 (0.52)	10.10 (13.99)	29.90 (17.90)	19.25 (-)	9.22 (7.88)	0.10 (0.32)
	CoT-VI	2.90 (1.60)	7.40 (4.27)	12.30 (12.37)	19.11 (-)	45.00 (18.95)	0.50 (0.71)
	CoT-TI	4.80 (0.42)	6.80 (1.55)	3.50 (1.72)	13.30 (1.95)	63.12 (4.67)	0.10 (0.32)
GPT-4o	STA-VI	0.20 (0.42)	6.20 (13.21)	32.70 (18.57)	19.00 (-)	10.31 (12.11)	0.80 (1.48)
	STA-TI	0.30 (0.48)	11.60 (18.72)	28.70 (19.81)	12.00 (-)	8.28 (5.90)	0.30 (0.67)
	CoT-VI	2.50 (1.43)	7.60 (5.10)	11.20 (10.03)	17.38 (-)	44.53 (15.51)	1.10 (0.88)
	CoT-TI	4.60 (0.84)	7.60 (1.84)	2.10 (0.88)	12.60 (2.46)	63.28 (5.86)	0.10 (0.32)
Claude-3.5 Sonnet	STA-VI	0.10 (0.32)	3.10 (9.80)	24.50 (10.97)	17.00 (-)	20.47 (12.60)	1.60 (0.97)
	STA-TI	0.50 (0.53)	15.90 (17.53)	22.90 (18.22)	15.80 (-)	8.90 (7.14)	0.10 (0.32)
	CoT-VI	4.80 (0.42)	7.20 (2.82)	2.20 (1.40)	12.70 (1.57)	65.16 (5.32)	0.00 (0.00)
	CoT-TI	5.00 (0.00)	6.30 (0.82)	2.00 (0.82)	12.00 (0.94)	67.50 (2.74)	0.00 (0.00)
Human	STA-VI	4.73 (0.45)	6.87 (1.63)	2.80 (1.69)	12.93 (1.62)	65.15 (4.35)	0.10 (0.31)

transition from VI to TI inputs. In STA conditions, NPE were extremely high, indicating near-random performance. The near-elimination of NPE in CoT-TI (e.g., Claude-3.5 Sonnet: $\mu = 2.00$, $\sigma = 0.82$) suggests that VLLMs can achieve a level of consistent rule application that exceeds human performance. This suggests that explicit reasoning prompts enable VLLMs to maintain a more consistent internal representation of the current sorting rule, reducing random errors.

All models required the fewest trials to complete the first category in the CoT-TI condition, with Claude-3.5 Sonnet performing best ($\mu = 12.00$, $\sigma = 0.94$), followed closely by GPT-4o ($\mu = 12.60$, $\sigma = 2.46$) and Gemini-1.5 Pro ($\mu = 13.30$, $\sigma = 1.95$). Notably, Claude-3.5 Sonnet outperformed the human baseline ($\mu = 12.93$, $\sigma = 1.62$).

CLR patterns showed substantial improvement from STA to

CoT conditions for all VLLMs, with the highest percentages observed in the CoT-TI condition. Claude-3.5 Sonnet achieved the highest CLR in this condition ($\mu = 67.50\%$, $\sigma = 2.74\%$), followed by GPT-4o ($\mu = 63.28\%$, $\sigma = 5.86\%$) and Gemini-1.5 Pro ($\mu = 63.12\%$, $\sigma = 4.67\%$). This indicates that under CoT-TI conditions, VLLMs can maintain conceptual understanding at a level comparable to or exceeding human capability.

FMS were generally low in STA conditions, but this reflects the models' overall poor performance rather than true set maintenance. The transition to CoT conditions led to increased FMS in the VI condition, suggesting that improved overall performance paradoxically led to more instances of set loss after initial successful rule application. However, in the CoT-TI condition, Claude-3.5 Sonnet achieved perfect set maintenance (FMS = 0.00), outperforming the human baseline. This

indicates that VLLMs can maintain exceptional consistency in rule application, potentially surpassing human capabilities in this aspect of cognitive flexibility.

These detailed metrics collectively reinforce the finding that CoT prompting, particularly when combined with textual inputs, substantially enhances VLLMs' cognitive flexibility as measured by WCST performance. While all models showed similar patterns across conditions, Claude-3.5 Sonnet consistently demonstrated superhuman cognitive flexibility in CoT-TI condition. The consistent pattern across all six metrics highlights the robustness of the effects of prompting strategy and input modality, while also revealing subtle differences in the cognitive capabilities of these VLLMs.

C. Analysis of Input modality

To investigate the performance difference between visual and textual input conditions, we conducted a detailed analysis of each model's ability to accurately perceive and interpret the WCST card features. This analysis aimed to determine whether the performance gap was due to limitations in visual processing or differences in cognitive flexibility across modalities. We evaluated the models' accuracy in identifying the three key features of WCST cards: color, shape, and number, comparing their descriptions against actual card features for each trial (detailed in Appendix -E).

The results indicate that all three models demonstrated high accuracy in visual feature recognition (Table II). Claude-3.5 Sonnet demonstrated perfect accuracy across all features, while Gemini-1.5 Pro and GPT-4o showed a decline in visual capabilities, particularly when recognizing how many cards were present in the image and the number of shapes on each card. Notably, GPT-4o almost always misidentified 5 cards as 6 cards.

TABLE II
VISUAL FEATURE RECOGNITION ACCURACY (%)

Model	Count	Color	Shape	Number	Overall
Gemini-1.5 Pro	75	100	100	97.81	96.97
GPT-4o	0	100	100	96.56	89.55
Claude-3.5 Sonnet	100	100	100	100	100

These findings suggest that the performance gap between VI and TI conditions is not solely attributable to limitations in visual feature extraction, but rather to the cascading effects of occasional visual misinterpretations on higher-order cognitive processes. In the VI condition, visual recognition errors can disrupt the model's ability to consistently apply a rule, necessitating re-exploration of the problem space. This phenomenon explains the increased variance observed in model performance under the CoT-VI condition compared to CoT-TI. The textual input's inherent precision eliminates this source of variability, allowing models to demonstrate more consistent cognitive flexibility. These results reveal the complex interplay between visual perception and executive function in VLLMs, highlighting the need for more robust visual processing pipelines.

D. Impact of Explicit Rule Exclusivity

All results in previous analyses were obtained under conditions that included both a general rule statement specifying "The correct answer depends on a rule, which will be based solely on either the number of shapes, the color of the shapes, or the shape type itself" and an explicit rule exclusivity constraint stating "There will be no combination of these characteristics to define the rule". To further investigate the robustness of VLLMs' cognitive flexibility, we conducted an additional study examining performance when explicit rule exclusivity was removed. We examined this under the CoT-TI condition, which had previously demonstrated near-human or superior cognitive flexibility for all models. The results are in Table III.

TABLE III
IMPACT OF EXPLICIT RULE EXCLUSIVITY ON CC

Model	Normal	w/o Constraints	Decline
Gemini-1.5 Pro	4.8 (0.42)	2.6 (2.01)	2.2
GPT-4o	4.6 (0.84)	3.5 (1.27)	1.1
Claude-3.5 Sonnet	5.0 (0.00)	4.7 (0.67)	0.3

Gemini-1.5 Pro exhibited the most pronounced sensitivity to the absence of explicit rule exclusivity, with mean CC decreasing from 4.8 ($\sigma = 0.42$) with the constraint to 2.6 ($\sigma = 2.01$) without it, representing a 2.2 decline. GPT-4o demonstrated moderate sensitivity, with performance dropping from 4.6 ($\sigma = 0.84$) to 3.5 ($\sigma = 1.27$) categories. Claude-3.5 Sonnet showed the most robust performance, maintaining high functionality even without the explicit exclusivity statement, with only a marginal decline from perfect performance ($\mu = 5.0$, $\sigma = 0.00$) to near-perfect ($\mu = 4.7$, $\sigma = 0.67$). The observed increases in standard deviations across all models when the exclusivity constraint was removed indicate that explicit rule exclusivity not only enhances performance but also promotes more consistent cognitive flexibility.

The differential declines observed among models reflect disparities in their ability to maintain simple rule structures in the absence of explicit constraints against more complex possibilities. Claude-3.5 Sonnet's robust performance suggests a superior ability to infer and adhere to simpler rule structures, even when the possibility of more complex rules is not explicitly excluded.

E. Simulating Cognitive Impairment

To explore the potential of VLLMs in modeling human cognitive impairment without modifying the models, we employed role-playing prompts to simulate three key aspects of the PFC function commonly impaired in various neurological conditions [41], [42]: goal maintenance, inhibitory control, and adaptive updating. This method leverages the models' ability to imagine and simulate different cognitive states, allowing us to study how they conceptualize and perform under various impairment conditions without modifying the underlying model architecture.

We focused on the CoT-TI condition, as it consistently yielded the best performance across all models in our previous

TABLE IV
WCST PERFORMANCE UNDER NORMAL AND SIMULATED IMPAIRMENT CONDITIONS (CoT-TI)

Model	Condition	CC	PE	NPE	TFC	CLR (%)	FMS
Gemini-1.5 Pro	Normal	4.80 (0.42)	6.80 (1.55)	3.50 (1.72)	13.30 (1.95)	63.12 (4.67)	0.10 (0.32)
	Goal Maint. (↓)	1.90 (1.60)	8.60 (7.83)	12.00 (12.38)	16.86 (-)	37.03 (22.39)	0.60 (0.70)
	Inhib. Ctrl. (↓)	1.70 (1.49)	6.40 (5.66)	20.10 (14.04)	32.00 (-)	30.63 (18.91)	0.90 (0.99)
	Adapt. Upd. (↓)	3.90 (0.74)	8.30 (2.71)	6.70 (4.57)	17.50 (6.20)	56.09 (9.53)	0.00 (0.00)
GPT-4o	Normal	4.60 (0.84)	7.60 (1.84)	2.10 (0.88)	12.60 (2.46)	63.28 (5.86)	0.10 (0.32)
	Goal Maint. (↓)	3.50 (1.65)	9.80 (2.66)	4.50 (3.78)	18.10 (9.46)	52.34 (16.49)	0.80 (1.03)
	Inhib. Ctrl. (↓)	4.20 (1.03)	10.10 (6.66)	3.70 (2.41)	14.00 (3.50)	57.97 (11.66)	0.30 (0.48)
	Adapt. Upd. (↓)	4.30 (0.82)	8.30 (3.13)	2.80 (2.49)	13.10 (1.79)	61.56 (9.24)	0.10 (0.32)
Claude-3.5 Sonnet	Normal	5.00 (0.00)	6.30 (0.82)	2.00 (0.82)	12.00 (0.94)	67.50 (2.74)	0.00 (0.00)
	Goal Maint. (↓)	3.20 (1.40)	12.50 (5.64)	5.50 (4.79)	17.50 (7.20)	47.19 (17.44)	0.60 (0.84)
	Inhib. Ctrl. (↓)	1.50 (1.65)	12.80 (13.82)	18.70 (19.82)	18.83 (-)	23.59 (20.75)	0.40 (0.52)
	Adapt. Upd. (↓)	3.60 (1.26)	8.60 (4.93)	7.50 (9.35)	20.60 (12.55)	51.56 (14.91)	0.00 (0.00)

experiments. The specific role-playing prompts and analysis methods for this component are detailed in supplementary Figure s-6. The results presented in Table IV, reveal that all three VLLMs demonstrated significant performance decrements under simulated impairment conditions, with patterns that align with neuropsychological observations of patients with prefrontal dysfunction.

Gemini-1.5 Pro exhibited the highest sensitivity to simulated impairments, with substantial declines in CC across all conditions. The most severe impact was observed under inhibitory control impairment (CC reduced from 4.80 to 1.70), accompanied by a marked increase in NPE from 3.50 to 20.10.

GPT-4o demonstrated greater resilience, maintaining relatively stable performance across impairment conditions. The model's CC decreased modestly from 4.60 to 3.50-4.30, with PE showing consistent increases across conditions. Notably, NPE remained stable, indicating a robust ability to maintain overall response consistency even under simulated impairments. This stability suggests that GPT-4o's decision-making processes may be more resistant to perturbation.

Claude-3.5 Sonnet, despite showing the highest baseline performance (CC = 5.00), exhibited significant vulnerability to simulated impairments. The model showed increases in both PE and NPE under impairment conditions, particularly for inhibitory control (PE: 12.80, NPE: 18.70). This pattern suggests that Claude-3.5 Sonnet's high baseline performance may rely on finely tuned cognitive processes that are more susceptible to disruption when specific aspects of executive function are impaired.

Across all models, inhibitory control impairment consistently produced the most severe performance decrements and led to increased NPE and FMS, aligning with observations in patients with orbitofrontal damage [43]. Models frequently mentioned irrelevant card features, simulating distraction and impulsivity. Goal maintenance impairment primarily affected CLR and FMS, reflecting difficulties in consistently applying rules. This pattern is consistent with observations in patients with dorsolateral prefrontal cortex lesions [44]. Adaptive updating impairment had a more moderate impact, mainly

affecting CC and CLR, while having less effect on FMS, consistent with difficulties in switching to new rules [45]. These distinct patterns of impairment across models suggest that while VLLMs can simulate aspects of cognitive dysfunction, the underlying mechanisms of their decision-making processes may differ.

V. DISCUSSION AND CONCLUSION

This study demonstrates that state-of-the-art VLLMs can achieve, and in some cases surpass, human-level cognitive flexibility as measured by the WCST, suggesting a potential for emulating and exceeding human set-shifting abilities in specific contexts. The observed performance gradient underscores the complex interplay between input modalities and prompting strategies. The consistent performance across both the standard WCST and our novel ALIEN Task variant provides compelling evidence that these models demonstrate genuine cognitive flexibility rather than relying on memorized patterns. Despite the complete transformation of surface features, the models exhibited virtually identical performance patterns across conditions, suggesting engagement with the underlying logical structure of set-shifting tasks rather than merely recalling similar examples from their training data. The performance gap between VI and TI conditions indicates that current VLLMs may rely more heavily on language-based reasoning pathways, even when processing visual information. Explicit reasoning prompts enable VLLMs to maintain more stable internal representations of task rules.

Our analysis of explicit rule exclusivity reveals a critical dependence of VLLMs on precise task instructions. The significant performance decline observed when specific rule constraints were removed highlights the models' reliance on explicit information to guide their decision-making processes. This finding suggests that VLLMs' impressive performance in structured tasks may not fully generalize to more ambiguous real-world scenarios without careful consideration of instruction design.

The simulation of cognitive impairments through role-playing prompts demonstrates the potential of VLLMs to

model complex behavioral patterns of executive dysfunction. The distinct performance profiles observed under simulated goal maintenance, inhibitory control, and adaptive updating impairments closely mirror behavioral patterns seen in human neuropsychological research. However, as recent research has shown that language models can form biased associations that mirror societal stereotypes, it is important to acknowledge that VLLMs' simulations of cognitive impairments may reflect stereotypical representations of clinical populations rather than authentic mechanisms underlying the disorders themselves.

Future research should focus on elucidating the underlying mechanisms that enable VLLMs to perform set-shifting tasks and investigating the generalizability of these abilities to other domains of executive function. Additionally, the development of more sophisticated visual processing capabilities and the integration of multimodal information processing warrant further exploration. The potential of VLLMs to simulate specific patterns of cognitive impairment also opens up new possibilities for creating realistic models of neuropsychological conditions, which could have applications in both clinical research and AI safety. By analyzing VLLMs' internal representations during simulated impairments, we could potentially decode the computational principles underlying various cognitive functions, complementing traditional neuroscience methods.

In conclusion, this study provides insights into our understanding of cognitive flexibility in VLLMs, revealing capabilities that match or exceed human performance and important limitations that depend on task framing and input modalities. As these models continue to evolve, a deeper understanding of their cognitive processes will be crucial for harnessing their potential while addressing their constraints, ultimately leading to more adaptive and robust AI systems that can flexibly navigate complex, real-world environments.

MODEL DETAILS

This study employs three state-of-the-art visual language models: Gemini-1.5 Pro, GPT-4o, and Claude-3.5 Sonnet. Table A1 presents a comparative overview of these models' key characteristics.

A. Gemini-1.5 Pro

Gemini-1.5 Pro, created by Google, employs a Mixture of Experts architecture, allowing for efficient processing of both textual and visual inputs. While specific architectural details are proprietary, the model demonstrates strong performance across various tasks. For image processing, Gemini-1.5 Pro uses a standardized approach where each image is equivalent to 258 tokens, regardless of size. Large images are scaled down to a maximum of 3072x3072 pixels, while small images are scaled up to 768x768 pixels, both preserving aspect ratio.

B. GPT-4o

GPT-4o, developed by OpenAI, represents an advanced iteration of the GPT series. It utilizes a transformer-based architecture and incorporates visual processing capabilities. GPT-4o offers adaptive image processing with low and high resolution modes, allowing for a balance between processing speed and detail level. In low resolution mode, it processes a 512px x 512px version of the image, representing it with 85 tokens. The high resolution mode initially processes a low-res image, then creates detailed 512px x 512px crops, each represented by 170 tokens.

C. Claude-3.5 Sonnet

Developed by Anthropic, Claude-3.5 Sonnet builds upon previous Claude models, incorporating enhanced visual understanding capabilities. The model utilizes a transformer-based architecture optimized for multimodal inputs. Claude-3.5 Sonnet balances multi-image processing by resizing images that exceed 1568 pixels on the long edge or approximately 1,600 tokens. It calculates token usage based on image dimensions (tokens = (width px * height px)/750) and emphasizes image clarity and text legibility for optimal performance.

PROMPTS AND TOKEN USAGE

D. Detailed prompts

The WCST setup consists of four stimulus cards, each featuring unique combinations of color (red, green, yellow, blue), shape (triangle, star, cross, circle), and number (one,

TABLE V
COMPARISON OF VISUAL LANGUAGE MODELS

Characteristic	Gemini-1.5 Pro	GPT-4o	Claude-3.5 Sonnet
Multimodal Capabilities	Text, audio, image, video	Text, audio, image	Text, image
API Version	gemini-1.5-pro	gpt-4o-2024-05-13	claude-3-5-sonnet-20240620
Access Method	aistudio.google.com	platform.openai.com	console.anthropic.com
Context Window	2M tokens	128K tokens	200K tokens
Maximum Out-Tokens	8,192 tokens	4,096 tokens	4,096 tokens
Knowledge Cutoff	November 2023	October 2023	April 2024
Release Date	May 2024	May 2024	June 2024
Model Ranking (till 2024-08-31)	LMSYS	#4	#2
	OpenCompass	#5	#2
	Benchmarks	#4	#1

two, three, four) of symbols. A series of 64 response cards is used, each sharing properties with the stimulus cards but in different combinations. The sorting rules are based on three possible categories: color, shape, or number. In our implementation, each trial is presented to the VLLMs as a single image containing two rows. The top row displays the four stimulus cards, while the bottom row shows one response card. This format is consistent across all 64 trials, providing a standardized visual input for the models to process. For the text-based conditions (TI), detailed descriptions of these images are provided instead.

The test procedure begins with the VLLM being instructed to match each response card to one of the stimulus cards based on a rule that it must deduce from feedback. After each match attempt, the VLLM receives feedback (correct or incorrect) without explicit mention of the current sorting rule. The sorting rule changes after 10 consecutive correct matches, without notification to the VLLM. The test concludes after all 64 cards have been presented.

We implemented four distinct experimental conditions to assess the VLLMs' performance: STA-VI (Straight to Answer with Original Image input), STA-TI (Straight to Answer with Original Text description input), CoT-VI (Chain of Thought reasoning with Original Image input), and CoT-TI (Chain of Thought reasoning with Original Text description input). Supplementary Figure s-2 provides a visual representation of the WCST procedure and sample stimuli used in our study.

The prompting strategies for each condition are illustrated in supplementary Figure s-3. For the CoT conditions, VLLMs were explicitly instructed to verbalize their reasoning process, including their observations, hypotheses about the current rule, and justification for their sorting decisions.

Supplementary Figure s-5 presents the web-based interface developed for human participants. This interface was designed to closely mimic the experience of VLLMs while accommodating human interaction patterns. It features a clear presentation of stimulus and response cards, along with intuitive controls for participants to indicate their sorting choices. The interface also provides immediate feedback on sorting decisions, mirroring the feedback mechanism used with VLLMs.

To explore the VLLMs' capacity to simulate cognitive impairments, we introduced role-playing scenarios as described in supplementary Figure s-6. This figure outlines the specific instructions given to models for simulating various prefrontal cortex dysfunctions, including impaired goal maintenance, inhibitory control deficits, and adaptive updating impairments. These scenarios were carefully designed to mimic specific cognitive deficits commonly observed in neurological conditions, allowing us to assess the models' ability to flexibly adapt their behavior to simulate human-like cognitive impairments.

Supplementary Figure s-7 provides a detailed example of a typical VLLM interaction during the WCST. This example illustrates how models process the presented cards, articulate their reasoning (in CoT-TI conditions), and make decisions. Supplementary Figure s-8 showcases the visual processing capabilities of VLLMs.

E. Visual Accuracy Calculation

The visual accuracy of VLLMs was assessed using a comprehensive scoring system that evaluated their ability to correctly identify key features of the WCST cards across 64 trials. The system encompassed five distinct measures: Card Count Accuracy, Color Accuracy, Shape Accuracy, Number Accuracy, and Overall Accuracy. For each trial, models were evaluated on their ability to correctly identify the presence of five cards and accurately describe the color, shape, and number of symbols on each card. Detailed descriptions of the Visual instructions are provided in supplementary Figure s-3. Supplementary Figure s-8 showcases the visual processing capabilities of VLLMs.

Count Accuracy was calculated as the proportion of trials where the model correctly identified the presence of five cards:

$$ACC_{count} = \frac{\sum_{i=1}^{64} I(c_i = 5)}{64} \times 100\% \quad (7)$$

where $I(c_i = 5)$ is an indicator function that equals 1 if the model correctly counted 5 cards in trial i , and 0 otherwise.

Color Accuracy, Shape Accuracy, and Number Accuracy were calculated similarly, assessing the model's performance across all cards in all trials:

$$ACC_{feature} = \frac{\sum_{i=1}^{64} \sum_{j=1}^5 I(f_{ij} = f_{ij}^*)}{64 * 5} \times 100\% \quad (8)$$

where $feature \in color, shape, number$, f_{ij} is the model's identification of the feature for card j in trial i , and f_{ij}^* is the correct feature.

The Overall Accuracy ($ACC_{overall}$) was computed as a composite score, incorporating all correct identifications while applying a penalty for overcounting cards. First, we define a penalty function P for overcounting:

$$P = \frac{\sum_{i=1}^{64} 0.5 \times \max(0, c_i - 5)}{64} \times 100\% \quad (9)$$

where c_i is the number of cards counted by the model in trial i . This penalty deducts 0.5 points for each card counted beyond the correct number of 5 in any given trial.

The Overall Accuracy is then calculated as:

$$ACC_{overall} = ACC_{count} + \sum ACC_{feature} - P \quad (10)$$

This assessment of the VLLMs' visual processing capabilities enables detailed comparisons across models and features. By evaluating multiple aspects of visual perception, from basic counting to complex feature recognition, the system offered insights into the strengths and limitations of each model's visual cognition in the context of the WCST.

F. Token Usage

To provide insight into the computational resources required, we list token usage across models and conditions (Table VI). Across all models, VI conditions consistently required more tokens than TI conditions, reflecting the additional computational demand of processing visual information. CoT conditions consumed significantly more tokens than STA conditions, indicating the increased computational cost of

TABLE VI
TOKEN USAGE AND COST ANALYSIS

Model	Condition	Task	Last Token Avg	Session Token Avg	Session Price Avg	Total Tokens	Total Price
Gemini-1.5 Pro	STA-VI	WCST	18,221	604,260	\$2.12	6,042,605	\$21.17
	STA-TI	WCST	6,631	227,378	\$0.8	2,273,782	\$7.95
	CoT-VI	WCST	19,898	658,489	\$2.32	6,584,887	\$23.18
		WCST	9,885	338,747	\$1.21	3,387,469	\$12.09
		WCST w/o restriction	12,795	421,199	\$1.52	4,211,991	\$15.19
	CoT-TI	WCST Goal Maint	8,970	311,489	\$1.11	3,114,890	\$11.08
		WCST Inhib Ctrl	10,447	355,556	\$1.27	3,555,564	\$12.74
		WCST Adapt Upd	8,423	295,753	\$1.05	2,957,531	\$10.49
	GPT-4o	STA-VI	WCST	7,373	251,712	\$1.26	2,517,120
STA-TI		WCST	6,831	233,900	\$1.17	2,338,995	\$11.66
CoT-VI		WCST	20,210	670,972	\$3.49	6,709,718	\$34.87
		WCST	20,216	672,265	\$3.5	6,722,651	\$34.96
		WCST w/o restriction	23,093	749,301	\$3.91	7,493,007	\$39.12
CoT-TI		WCST Goal Maint	18,414	619,910	\$3.22	6,199,099	\$32.19
		WCST Inhib Ctrl	18,642	624,323	\$3.24	6,243,230	\$32.43
		WCST Adapt Upd	18,824	634,350	\$3.29	6,343,505	\$32.93
Claude-3.5 Sonnet		STA-VI	WCST	27,404	903,104	\$2.72	9,031,040
	STA-TI	WCST	7,073	242,113	\$0.73	2,421,131	\$7.34
	CoT-VI	WCST	43,704	1,426,502	\$4.48	14,265,023	\$44.78
		WCST	19,257	641,037	\$2.08	6,410,367	\$20.77
		WCST w/o restriction	20,718	675,461	\$2.2	6,754,606	\$21.96
	CoT-TI	WCST Goal Maint	23,806	778,877	\$2.54	7,788,771	\$25.42
		WCST Inhib Ctrl	24,550	802,087	\$2.62	8,020,867	\$26.18
		WCST Adapt Upd	24,774	799,338	\$2.62	7,993,378	\$26.16

explicit reasoning processes. Among the models, Claude-3.5 Sonnet showed the highest token usage across all conditions, suggesting a more computationally intensive approach to task processing. These token usage patterns provide valuable insights into the relative efficiency and resource requirements of different VLLMs and experimental conditions in cognitive flexibility tasks.

CODE AVAILABILITY

The complete code and experimental data are available at: <https://drive.google.com/file/d/1uUyPn6fP4JDI50zRcdMeGJKwRaI-JTUs/view?usp=sharing>.

REFERENCES

- [1] D. R. Dajani and L. Q. Uddin, "Demystifying cognitive flexibility: Implications for clinical and developmental neuroscience," *Trends in neurosciences*, vol. 38, no. 9, pp. 571–578, 2015.
- [2] L. Q. Uddin, "Cognitive and behavioural flexibility: neural mechanisms and clinical considerations," *Nature Reviews Neuroscience*, vol. 22, no. 3, pp. 167–179, 2021.
- [3] T. Ionescu, "Exploring the nature of cognitive flexibility," *New ideas in psychology*, vol. 30, no. 2, pp. 190–200, 2012.
- [4] T. Spellman, M. Svei, J. Kaminsky, G. Manzano-Nieves, and C. Liston, "Prefrontal deep projection neurons enable cognitive flexibility via persistent feedback monitoring," *Cell*, vol. 184, no. 10, pp. 2750–2766, 2021.
- [5] S. Funahashi and J. M. Andreau, "Prefrontal cortex and neural mechanisms of executive function," *Journal of Physiology-Paris*, vol. 107, no. 6, pp. 471–482, 2013.
- [6] L. Chen, B. Li, S. Shen, J. Yang, C. Li, K. Keutzer, T. Darrell, and Z. Liu, "Large language models are visual reasoning coordinators," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [7] I. Momennejad, H. Hasanbeig, F. Vieira Frujeri, H. Sharma, N. Jovic, H. Palangi, R. Ness, and J. Larson, "Evaluating cognitive maps and planning in large language models with cogeval," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [8] Y. Chang, X. Wang, J. Wang, Y. Wu, L. Yang, K. Zhu, H. Chen, X. Yi, C. Wang, Y. Wang *et al.*, "A survey on evaluation of large language models," *ACM Transactions on Intelligent Systems and Technology*, vol. 15, no. 3, pp. 1–45, 2024.
- [9] Y. Qu, C. Wei, P. Du, W. Che, C. Zhang, W. Ouyang, Y. Bian, F. Xu, B. Hu, K. Du *et al.*, "Integration of cognitive tasks into artificial general intelligence test for large models," *Iscience*, vol. 27, no. 4, 2024.
- [10] OpenAI, "Hello gpt-4o," <https://openai.com/index/hello-gpt-4o/>, 2024, accessed: 2024-05-18.
- [11] M. Reid, N. Savinov, D. Teplyashin, D. Lepikhin, T. Lillicrap, J.-b. Alayrac, R. Soricut, A. Lazaridou, O. Firat, J. Schrittwieser *et al.*, "Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context," *arXiv preprint arXiv:2403.05530*, 2024.
- [12] Anthropic, "Announcements: Claude 3.5 sonnet," <https://www.anthropic.com/news/claude-3-5-sonnet>, 2024, accessed: 2024-06-21.
- [13] E. A. Berg, "A simple objective technique for measuring flexibility in thinking," *The Journal of general psychology*, vol. 39, no. 1, pp. 15–22, 1948.
- [14] S. Miles, C. A. Howlett, C. Berryman, M. Nedeljkovic, G. L. Moseley, and A. Phillipou, "Considerations for using the wisconsin card sorting test to assess cognitive flexibility," *Behavior research methods*, vol. 53, no. 5, pp. 2083–2091, 2021.
- [15] K. W. Greve, "The wcst-64: A standardized short-form of the wisconsin card sorting test," *The Clinical Neuropsychologist*, vol. 15, no. 2, pp. 228–234, 2001.
- [16] K. Jodzio and D. Biechowska, "Wisconsin card sorting test as a measure of executive function impairments in stroke patients," *Applied neuropsychology*, vol. 17, no. 4, pp. 267–277, 2010.
- [17] C.-H. Lie, K. Specht, J. C. Marshall, and G. R. Fink, "Using fmri to decompose the neural processes underlying the wisconsin card sorting test," *Neuroimage*, vol. 30, no. 3, pp. 1038–1049, 2006.
- [18] C. Kim, N. F. Johnson, S. E. Cilles, and B. T. Gold, "Common and distinct mechanisms of cognitive flexibility in prefrontal cortex," *Journal of Neuroscience*, vol. 31, no. 13, pp. 4771–4779, 2011.
- [19] L. Qiao, M. Xu, X. Luo, L. Zhang, H. Li, and A. Chen, "Flexible adjustment of the effective connectivity between the fronto-parietal and visual regions supports cognitive flexibility," *NeuroImage*, vol. 220, p. 117158, 2020.
- [20] A. Diamond, "Executive functions," *Annual Review of Psychology*, vol. 64, no. 1, pp. 135–168, 2013.
- [21] P. Eling, K. Derckx, and R. Maes, "On the historical and conceptual

background of the wisconsin card sorting test," *Brain and cognition*, vol. 67, no. 3, pp. 247–253, 2008.

[22] D. T. Stuss and B. Levine, "Adult clinical neuropsychology: lessons from studies of the frontal lobes," *Annual review of psychology*, vol. 53, no. 1, pp. 401–433, 2002.

[23] P. D. Zelazo, "The dimensional change card sort (dccs): A method of assessing executive function in children," *Nature protocols*, vol. 1, no. 1, pp. 297–301, 2006.

[24] A. Heinzl, G. Northoff, H. Boeker, P. Boesiger, and S. Grimm, "Emotional processing and executive functions in major depressive disorder: dorsal prefrontal activity correlates with performance in the intra–extra dimensional set shift," *Acta neuropsychiatrica*, vol. 22, no. 6, pp. 269–279, 2010.

[25] J. W. Strachan, D. Albergo, G. Borghini, O. Pansardi, E. Scality, S. Gupta, K. Saxena, A. Rufo, S. Panzeri, G. Manzi *et al.*, "Testing theory of mind in large language models and humans," *Nature Human Behaviour*, pp. 1–11, 2024.

[26] B. C. Lee and J. Chung, "An empirical investigation of the impact of chatgpt on creativity," *Nature Human Behaviour*, pp. 1–9, 2024.

[27] B. Fatemi, M. Kazemi, A. Tsitsulin, K. Malkan, J. Yim, J. Palowitch, S. Seo, J. Halcrow, and B. Perozzi, "Test of time: A benchmark for evaluating llms on temporal reasoning," *arXiv preprint arXiv:2406.09170*, 2024.

[28] P. Rahmazadehgervi, L. Bolton, M. R. Taesiri, and A. T. Nguyen, "Vision language models are blind," *arXiv preprint arXiv:2407.06581*, 2024.

[29] R. Loconte, G. Orru, M. Tribastone, P. Pietrini, and G. Sartori, "Challenging chatgpt'intelligence' with human tools: a neuropsychological investigation on prefrontal functioning of a large language model," *Intelligence*, 2023.

[30] G. V. Aher, R. I. Arriaga, and A. T. Kalai, "Using large language models to simulate multiple humans and replicate human subject studies," in *International Conference on Machine Learning*. PMLR, 2023, pp. 337–371.

[31] P. Liu, W. Yuan, J. Fu, Z. Jiang, H. Hayashi, and G. Neubig, "Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing," *ACM Computing Surveys*, vol. 55, no. 9, pp. 1–35, 2023.

[32] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou *et al.*, "Chain-of-thought prompting elicits reasoning in large language models," *Advances in neural information processing systems*, vol. 35, pp. 24 824–24 837, 2022.

[33] Z. Chu, J. Chen, Q. Chen, W. Yu, T. He, H. Wang, W. Peng, M. Liu, B. Qin, and T. Liu, "A survey of chain of thought reasoning: Advances, frontiers and future," *arXiv preprint arXiv:2309.15402*, 2023.

[34] T. Kojima, S. S. Gu, M. Reid, Y. Matsuo, and Y. Iwasawa, "Large language models are zero-shot reasoners," *Advances in neural information processing systems*, vol. 35, pp. 22 199–22 213, 2022.

[35] X. Wang, J. Wei, D. Schuurmans, Q. Le, E. Chi, S. Narang, A. Chowdhery, and D. Zhou, "Self-consistency improves chain of thought reasoning in language models," *arXiv preprint arXiv:2203.11171*, 2022.

[36] Q. Chen, L. Qin, J. Zhang, Z. Chen, X. Xu, and W. Che, "M $\hat{3}$ cot: A novel benchmark for multi-domain multi-step multi-modal chain-of-thought," *arXiv preprint arXiv:2405.16473*, 2024.

[37] D. Zhou, N. Schärli, L. Hou, J. Wei, N. Scales, X. Wang, D. Schuurmans, C. Cui, O. Bousquet, Q. Le *et al.*, "Least-to-most prompting enables complex reasoning in large language models," *arXiv preprint arXiv:2205.10625*, 2022.

[38] R. Pryzant, D. Iter, J. Li, Y. T. Lee, C. Zhu, and M. Zeng, "Automatic prompt optimization with" gradient descent" and beam search," *arXiv preprint arXiv:2305.03495*, 2023.

[39] Y. Zhou, A. I. Muresanu, Z. Han, K. Paster, S. Pitis, H. Chan, and J. Ba, "Large language models are human-level prompt engineers," *arXiv preprint arXiv:2211.01910*, 2022.

[40] D. J. Schretlen, *Modified Wisconsin Card Sorting Test®: M-WCST; Professional Manual*. PAR, 2010.

[41] E. K. Miller and J. D. Cohen, "An integrative theory of prefrontal cortex function," *Annual review of neuroscience*, vol. 24, no. 1, pp. 167–202, 2001.

[42] D. T. Stuss and M. P. Alexander, "Is there a dysexecutive syndrome?" *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 362, no. 1481, pp. 901–915, 2007.

[43] D. T. Stuss, D. Benson, W. Weir, M. Naeser, I. Lieberman, and D. Ferrill, "The involvement of orbitofrontal cerebrum in cognitive tasks," *Neuropsychologia*, vol. 21, no. 3, pp. 235–248, 1983.

[44] D. Stuss, B. Levine, M. Alexander, J. Hong, C. Palumbo, L. Hamer, K. Murphy, and D. Izukawa, "Wisconsin card sorting test performance in patients with focal frontal and posterior brain damage: effects of lesion location and test structure on separable cognitive processes," *Neuropsychologia*, vol. 38, no. 4, pp. 388–402, 2000.

[45] B. Milner, "Effects of different brain lesions on card sorting: The role of the frontal lobes," *Archives of neurology*, vol. 9, no. 1, pp. 90–100, 1963.



Guangfu Hao received the Bachelor degree of Engineering in Mechatronic Engineering from Beijing Institute of Technology, Beijing, China, in 2021. He is currently pursuing a Ph.D. degree in pattern recognition and intelligent systems at the National Laboratory of Pattern Recognition, Institute of Automation, Chinese Academy of Sciences, Beijing, China. His research interests lie in computational neuroscience and artificial intelligence, particularly in brain-inspired deep learning. He is focused on developing more robust neural networks by drawing inspiration from the prefrontal cortex and hippocampus. Currently, his work centers on visual reasoning tasks, reinforcement learning, and exploring prefrontal cortex functions in large language models (LLMs).



Frederic Alexandre is a Director of Research at Inria and Head of the Mnemosyne Team, based at Inria Bordeaux, France. His research focuses on computational neuroscience, cognitive modeling, and artificial intelligence, exploring how intelligent behavior emerges from neural networks. He investigates how brain regions like the prefrontal cortex, hippocampus, and basal ganglia interact to support functions such as memory, decision-making, and reasoning. Additionally, his work extends to autonomous robotics and machine learning, aiming to build models of cognitive systems that integrate sensory, emotional, and motor functions. His team works closely with neuroscience and medical communities, applying their models to study brain mechanisms, neurodegenerative diseases, and industrial applications, while also contributing to the development of Artificial General Intelligence.



Shan Yu received the B.S. and Ph.D. degrees in biology from the University of Science and Technology of China, Hefei, China, in 2000 and 2005, respectively. From 2005 to 2014, he conducted Post-doctoral Research with the Max-Planck Institute of Brain Research, Germany (2005–2008), and the National Institute of Mental Health, USA (2008–2014). In 2014, he was with the Institute of Automation (IA), Chinese Academy of Sciences (CASIA), as a recipient of the "One Hundred Talents" program of CAS. He is currently a Professor and the Director of the Brain Mapping and Brain-Inspired Intelligence Laboratory at IA, CASIA. Since 2018, he has also served as the Deputy Director with the National Laboratory of Pattern Recognition (NLPR). He has authored and coauthored more than 60 peer-reviewed papers in neuroscience and other interdisciplinary fields at leading international journals such as the *Nature Machine Intelligence*, the *Journal of Neuroscience*, *eLife*, etc. His current research interests include neuronal information processing, brain-inspired computing, and artificial intelligence.