

SCoPE VLM: Selective Context Processing for Efficient Document Navigation in Vision-Language Models

Gyubeum Lim¹, Yemo Koo², Vijay Krishna Madiseti¹

¹Georgia Institute of Technology ²Konkuk University
{glim36, vkm}@gatech.edu yemoyemo010831@konkuk.ac.kr

Abstract

Understanding long-context visual information remains a fundamental challenge for vision-language models, particularly in agentic tasks such as GUI control and web navigation. While web pages and GUI environments are inherently structured documents, current VLMs typically neglect decision-oriented document understanding in their training objectives. Existing approaches primarily extend visual embeddings to process long, high-resolution inputs, but these methods are memory-intensive and impractical for locally deployable solutions. To address these issues, we propose SCoPE VLM, a document navigation expert that leverages a novel Chain of Scroll mechanism to selectively and recursively navigate documents, focusing exclusively on relevant segments. We introduce a dedicated data generation pipeline to construct informative Chain of Scroll trajectories and Episodic Group Relative Policy Optimization, a tailored reinforcement learning method to bridge the gap between training and inference. Our method substantially reduces memory usage and effectively models human-like reading behaviors. To the best of our knowledge, SCoPE VLM is the first framework to explicitly model agentic reading patterns in multi-page document question answering, advancing the capabilities of multimodal agents.

1 Introduction

Recent advancements in Vision-Language models (VLMs) (OpenAI, 2023; OpenAI et al., 2024; Liu et al., 2023; Bai et al., 2023; Chen et al., 2023) have demonstrated outstanding capabilities in visual understanding, achieving a tight alignment between text and images to answer complex queries. Despite these successes, handling long-context information remains a substantial challenge, especially for long documents and GUI environments. In fact, tasks involving GUI and web environments (Xie et al., 2024; Koh et al., 2024; Sun et al., 2025)

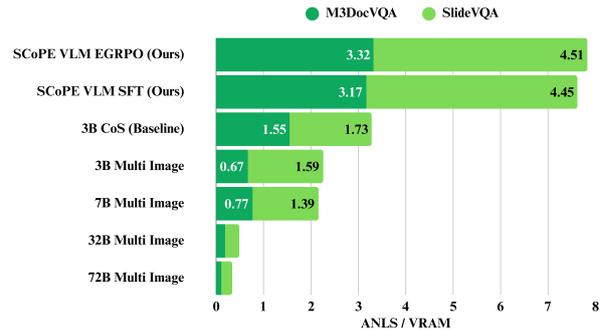


Figure 1: Performance-efficiency comparison of SCoPE VLM variants against Qwen2.5 VL series on M3DocVQA and SlideVQA from Table 3, Section 4, achieving 2.38x higher efficiency than the baseline.

inherently require both sophisticated multimodal document understanding and agentic capabilities. However, conventional approaches neglect action-based document understanding. Furthermore, state-of-the-art VLMs struggle in these environments due to their limited capability to explore and process extensive multimodal contexts. To address this, reinforcement learning offers a promising approach to teach such agentic behavior. While training directly on document-based environments offers a practical and accessible alternative to costly and unstable GUI simulators (Bai et al., 2024), this remains underexplored in the community.

Despite these advantages in action-based document understanding, existing approaches merely extend visual embeddings to process high-resolution images in long-context scenarios (Bai et al., 2025; Ye et al., 2024; Li et al., 2024). These methods rely on processing the entire visual context in a single pass, which is inevitably memory-intensive. As a result, they are not only impractical for local deployment but also fundamentally unscalable to real-world scenarios such as 100-page documents or websites.

To mitigate these limitations, we propose SCoPE VLM, a document navigation expert that leverages

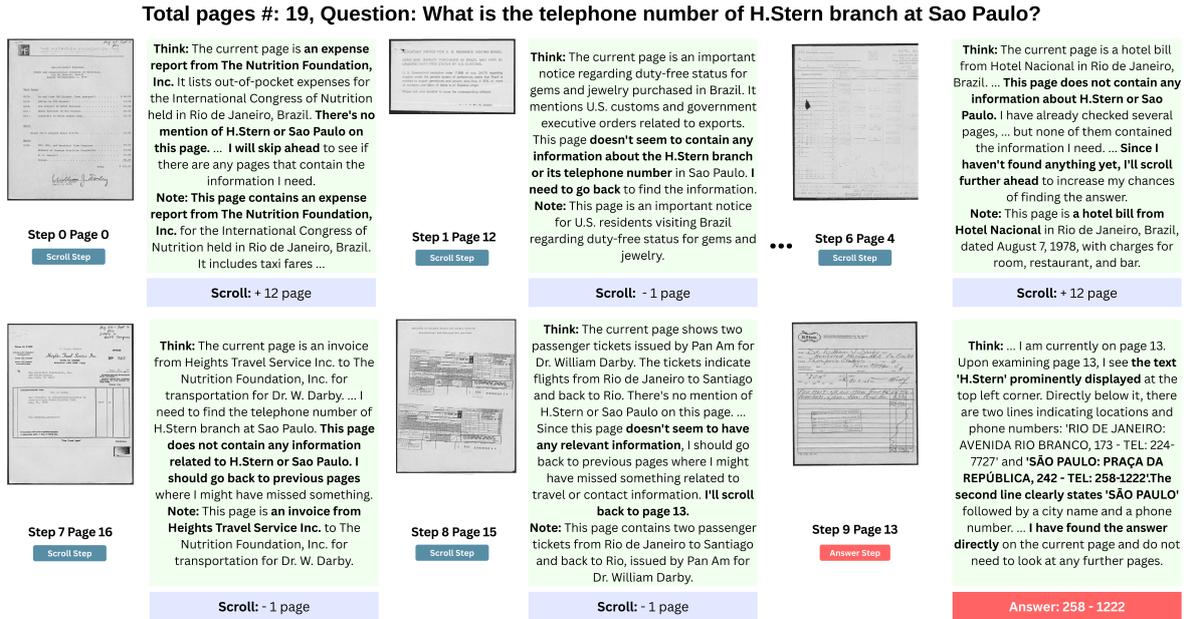


Figure 2: Chain of Scroll navigation trajectories. SCoPE VLM with EGRPO efficiently locates telephone information in a 19-page document, visiting only 9 pages.

the Chain of Scroll framework (CoS), an action-based inference time scaling strategy that allows the model to focus only on relevant document segments to answer the user query. The framework is executed through numerous recursive generations based on query-relevant context, which the model generated in the previous step.

To train such advanced decision-making capabilities in VLMs, we introduce the cold start dataset for Supervised-Fine-Tuning (SFT) CoS Framework. Although SFT is effective in teacher-forcing the output distribution quickly, a gap still presents between the inference and SFT training. Thus, we further enhance the model’s agentic capabilities by introducing Episodic Group Relative Policy Optimization (EGRPO) to further push the limit of SCoPE VLM for effective context-based action selection.

Empirical evaluations illustrate that SCoPE VLM maintains the accuracy of existing baselines while drastically improving memory efficiency. As illustrated in Figure 1, our model exhibits superior ANLS/VRAM performance, with the total scores reflecting the aggregated metrics across both evaluation datasets. Additionally, our method effectively generalizes its learned navigation capabilities to GUI control tasks, demonstrating enhanced adaptability compared to the baseline model. By explicitly modeling agentic behaviors in document-based question answering, SCoPE VLM represents a critical advancement toward more efficient, capable, and locally deployable multimodal agents.

2 Related Work

2.1 Embedding High Resolution Images

To incorporate high-resolution images, conventional VLMs leverage multi-window based vision token processing. LLaVA-Next (Liu et al., 2024) pioneered this approach by splitting high-resolution images into up to five windows, using 5x more visual tokens than standard single window image processing (Liu et al., 2023) to better capture fine details. Furthermore, InternVL 1.5 enhanced this paradigm by scaling the number of tiles up to x40 of the single window, allocating 10496 tokens per image through its tile-based processing system. Qwen2.5 VL (Bai et al., 2025) further extends by exploiting the Naive Dynamic Resolution (Dehghani et al., 2023) which consumes up to 16384 image tokens. They demonstrate that the increased number of visual tokens enables better preservation of fine-grained visual understanding. However, the computational costs scale exponentially with token counts, which results in high memory requirements in both training and inference time. These trade-offs highlight that the recent advancements in high-resolution image processing are not scalable to long document understanding.

2.2 Vision Token Compression

A prominent strategy for mitigating extensive vision token processing is vision token compression. Honeybee (Cha et al., 2023) refines token compression by proposing locality-enhanced projectors, thus im-

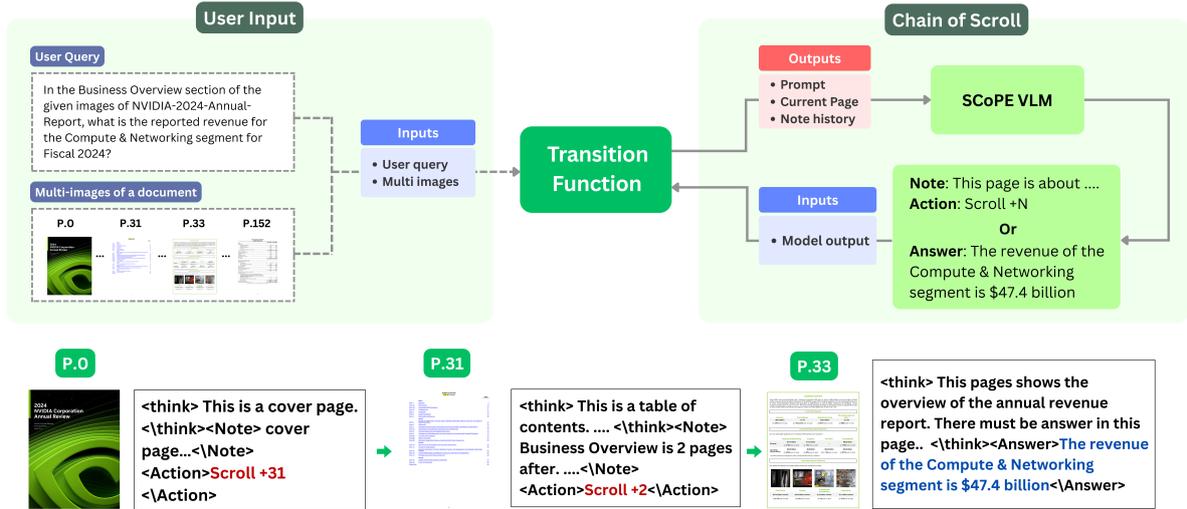


Figure 3: Overview of the Chain of Scroll framework. The model iteratively decides whether to scroll to new pages or generate an answer based on accumulated context and relevance signals, emulating human-like selective document navigation.

proving token relevance in dense visual data. In video understanding, where every frame contributes a dense set of visual tokens, visual token compression becomes more critical. PVC (Yang et al., 2024) and LongVU (Shen et al., 2024) both attempt to mitigate this issue by dynamically compressing tokens, preserving critical information while discarding less relevant details. However, these methods frequently struggle with performance degradation when tasks demand high resolution.

2.3 Multimodal Retrieval-Augmented Generation (RAG)

Multimodal Retrieval-Augmented Generation (RAG) has emerged as a promising framework for addressing the challenges of integrating extensive external multimodal contexts. These systems combine a retrieval mechanism to search relevant multimodal contexts with a generation model (Yasunaga et al., 2023). Despite its strengths, RAG systems struggle with processing on-demand user queries involving long-context documents. A primary constraint is the reliance on a pre-structured database, which must be prepared in advance (Lewis et al., 2021). This requirement introduces challenges in dynamic or real-time scenarios where the user query may involve newly introduced data.

3 Proposed Methods

3.1 Chain of Scroll

In real-world document question answering, queries are often localized to specific sections of a document rather than requiring holistic processing.

Algorithm 1 Chain of Scroll (CoS)

Input:
 q ▷ user’s query
 $imgs = \{Image_0, \dots, Image_N\}$ ▷ input images
 $maxSteps$ ▷ max steps allowed to scroll

Output: $answer$

Initialization:
1: $answer \leftarrow \emptyset$ ▷ no answer yet
2: $page \leftarrow 0$ ▷ start at first page
3: $scroll \leftarrow 0$ ▷ number of pages to move
4: $notes \leftarrow \emptyset$ ▷ all collected notes
5: $visited \leftarrow [False]^{|images|}$ ▷ store visited pages

Run Chain of Scroll:
6: **for** $step \leftarrow 1$ **to** $\min(maxSteps, |images|)$ **do**
7: $c \leftarrow (q, imgs, page, scroll, notes, visited)$
8: $prompt, cur_img, page, visited \leftarrow \text{TRAN_FN}(c)$
9: $response \leftarrow \pi_\theta(cur_img, prompt)$
10: $(cur_note, scroll, answer) \leftarrow \text{PARSE}(response)$
11: $notes \leftarrow notes \cup \{cur_note\}$
12: **if** $answer \neq \emptyset$ **then**
13: **break** ▷ answer found, early exit
14: **end if**
15: **end for**
16: **return** $answer$

While each page carries its own content, its relevance to the query varies; some pages provide critical information, while others are entirely irrelevant. Humans naturally adapt to this structure by skipping or answering based on contextual signals. This selective nature leads to efficient context processing with a minimal number of pages visited. To emulate this adaptive behavior, we propose the Chain of Scroll, an action-based document navigation strategy to mimic human-like reading behavior by making informed decisions at each step.

The Chain of Scroll framework allows the model

to make recursive decisions—whether to scroll or answer—based on contextual cues accumulated throughout the episode. By modeling the semantic coherence between the query and document segments, the model learns to predict which regions are likely to contain relevant information, how far to skip, and when to stop. This selective navigation not only reduces unnecessary computation but also mirrors human-like reading behavior, where not all content is equally attended to. The overall flow of the proposed CoS framework is illustrated in Figure 3 and formally outlined in Algorithm 1. In addition, Figure 2 demonstrates the human-like document navigation to answer the user query, while processing half of the document. See Section E for the full trajectory. For the formal representation of the CoS framework, including the action space and transition function, see Section C in the appendix.

Action space: In the CoS framework, the model operates within a discrete action space consisting of two steps: Scroll and Answer. At the beginning of each CoS trajectory, the model is given the first page of the document along with the user query. After performing Chain of Thought (CoT) reasoning (Wei et al., 2022; Wang et al., 2023; Mu et al., 2023; Lu et al., 2022; Guo et al., 2024), the model may choose to take notes and return a scroll value relative to the current page number. These generated notes are accumulated and used as context for subsequent steps via the defined transition function. Once the model has gathered sufficient information to address the query, it selects the Answer action following another CoT step. Thus, the action space comprises (1) a scroll value range bounded by the total number of pages in both positive and negative directions and all possible notes with query-relevant information, and (2) all possible answers. Both notes and answers are represented within the model’s token vocabulary.

Transition function: To emulate action-based context processing, it is essential to design a function that maps a previous state to the next, enabling sequential decision-making without a pre-defined environment. RL4VLM (Zhai et al., 2024) introduced a post-processing function to integrate VLMs into the decision-making loop; however, it still depended on fixed environments. In contrast, we propose a transition function that enables step-by-step progression for targeted document navigation as a Markov Decision Process (MDP). The major role of the transition function is to transform user queries into a structured CoS prompt (see an example in

Figure 5 in the appendix), bring the chosen page by the action from the previous step, validate scroll actions, manage visited page histories, and accumulate evolving notes that summarize key observations. At each step, it processes the model’s output, maps to the next step, and contextualizes based on the previous responses. As shown in Algorithm 1, the transition function plays the role of an environment by directing states and propagating contextual information. This function is used during both training and inference in our framework. It also serves as an exception handler, returning the context based on a random page when given an illegal scroll action.

Input prompt: As shown in Figure 5 in the appendix, the input prompt dynamically updates at each step with accumulated notes that preserve query-related context. This enables the model to maintain a continuous memory of its exploration based on globally informed decisions about whether to scroll further or generate an answer. Figure 2 illustrates how context from previous steps can effectively persist across exploration. In contrast to traditional multi-turn inference, our framework functions through multiple iterations of single-turn inference, with previous context preserved in accumulated notes. This design considerably improves memory efficiency while also allowing analysis of multi-page documents in high-resolution, with memory requirements comparable to those of processing a single image. See Section A in the appendix for the CoS prompt and an example.

3.2 SCoPE Dataset

Name	M3DocVQA	MP-DocVQA	SlideVQA	MMLongBench-Doc
#Pages / question	8.95	18.65	20.00	47.67
#Questions	8K	8K	4K	1K
Evidence page	Multi	Single	Multi	Multi

Table 1: Training datasets for supervised fine-tuning and reinforcement learning stages.

The SFT dataset for CoS is designed to capture the reasoning and decision-making capabilities of large-scale VLMs as a baseline. To generate high-quality annotations similar to Figure 2 at every sampled steps, we employ a pseudo-labeling ensemble comprising Gemini 1.5 Pro, Gemini Flash 2.0, and Gemini Flash 2.0 Thinking (Team et al., 2024), optimizing for generation time, cost, and quality. This ensemble is applied to a comprehensive corpus of 21K multi-page document question-answering samples, as shown in Table 1 (Cho et al., 2024; Tito et al., 2023; Tanaka et al., 2023; Ma et al., 2024).

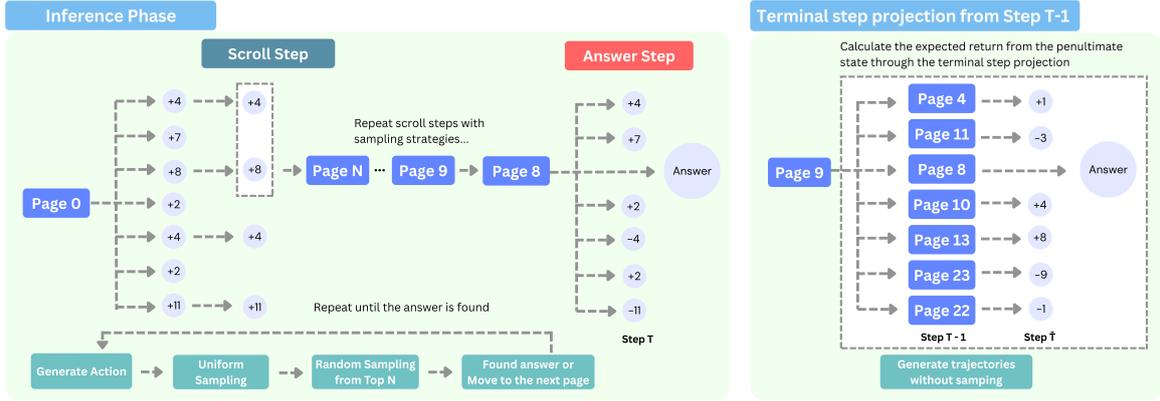


Figure 4: Overview of Episodic Group Relative Policy Optimization. The tailored framework for CoS employs group-based sampling with top-N selection during document navigation. Terminal step projection from the penultimate state enables return estimation that guides the model toward correct answers.

Since CoS is a novel task, even state-of-the-art proprietary VLMs lack the capabilities to generate high-quality CoS episodes without guidance. To address this limitation, we leverage datasets with explicit annotations of the evidence pages that must be visited. For M3DocVQA, we utilize the annotation model to pre-select query-relevant pages for each question.

Before generating CoS step responses, we sample page trajectories that models choose to read. Each trajectory designates page 0 as the initial state and selects one target page as the terminal state. To introduce natural variance, we randomly sample trajectory lengths, while incorporating both target pages and unrelated pages as intermediate steps to simulate realistic reading patterns.

During annotation, pseudo-labeling models receive three key inputs: the current page index, the question, and a defined scroll action transitioning to the next step along the predefined trajectory. This approach enables generations of appropriate scroll step responses. By explicitly providing the next page in advance, we induce high-quality reasoning that semantically aligns with the annotated action. For terminal steps, the annotation model formulates the final answer based on the current image, accumulated notes from previous steps, and the answer. This guided approach ensures that reasoning, memorization behavior, and decision-making are tightly aligned throughout the trajectory. Also, we annotated samples without answers to return “The answer cannot be found.” The response of our model to a query with no answer is presented in Section F.

3.3 Reinforcement Learning

With the SCoPE dataset, we perform Supervised Fine-Tuning (SFT) on Qwen2.5 VL 3B Instruct

due to its strong performance in high-resolution document understanding. While SFT effectively establishes basic decision-making capabilities, several challenges persist. First, the SFT model occasionally selects invalid actions that exceed the allowed image index range. Second, it often continues scrolling even after all document pages have been processed. See Section E for the full Chain of Scroll trajectory of SCoPE VLM SFT which fails to generate a valid answer after visiting all pages in the given document. These issues highlight a fundamental challenge: determining whether to continue scrolling or to provide an answer based on potentially incomplete information. Such decision-making under uncertainty cannot be adequately addressed by SFT alone, as it lacks mechanisms to encourage exploration beyond ground-truth trajectories. To overcome these limitations, we incorporate reinforcement learning (RL) to optimize CoS capabilities of SCoPE VLM.

Episodic Group Relative Policy Optimization:

To adapt GRPO (Shao et al., 2024) to the episodic structure of CoS, we introduce Episodic Group Relative Policy Optimization (EGRPO), a group-based approach tailored for multi-step document navigation tasks. As illustrated in Figure 4, EGRPO generates \tilde{G} candidate responses at each step and employs a two-stage sampling strategy: first, we uniformly sample G candidates from the original \tilde{G} to preserve reward diversity, then randomly select one response from the top- N candidates ranked by reward scores. This hybrid approach effectively balances exploration and exploitation—the uniform sampling maintains diversity across the reward spectrum while the top- N selection ensures high-quality action selection, allowing the model to efficiently explore diverse trajectories while systemati-

cally pursuing paths that maximize returns.

EGRPO extends GRPO by optimizing both the terminal step T (the answer step) and the penultimate step $T - 1$ (the final scroll step). To estimate future returns without exhaustive trajectory generation, we introduce a terminal step projection mechanism. After episode termination, we perform one additional generation to the projected terminal step (\hat{T}) from each penultimate state without sampling. This allows us to evaluate the quality of penultimate actions by observing where they lead, effectively propagating reward signals backward through the trajectory. The terminal state projection is under the assumption that the penultimate scroll step represents a state from which critical information needed to answer the user’s question is accessible. This assumption suggests that the probability of providing a correct answer given the penultimate state approximates the product of reaching the answer page and correctly extracting the answer once there. Also, we assume that following the dynamic programming principle, the training process optimizes not only the final transitions but the entire trajectories across diverse training sequences of varying lengths.

The EGRPO objective is provided below:

$$\mathcal{J}_{\text{EGRPO}}(\theta) = \gamma \mathcal{J}^{(T)}(\theta) + \mathcal{J}^{(T-1)}(\theta) \quad (1)$$

$$\mathcal{J}^{(t)}(\theta) = \mathbb{E}_{(q,a) \sim \mathcal{D}} \mathbb{E}_{\{o_i^{(t)}\}_{i=1}^G} \left[\frac{1}{G} \sum_{i=1}^G \frac{1}{|o_i^{(t)}|} \right. \quad (2)$$

$$\left. \sum_{k=1}^{|o_i^{(t)}|} \min\left(\rho_{i,k}^{(t)} \hat{A}_i^{(t)}, \text{clip}(\rho_{i,k}^{(t)}, 1 - \varepsilon, 1 + \varepsilon) \hat{A}_i^{(t)}\right) \right] \quad (3)$$

where $t \in \{T - 1, T\}$, $\{\hat{o}_j^{(t)}\}_{j=1}^{\hat{G}} \sim \pi_{\theta_{\text{old}}}(\cdot | q)$,

$$\{o_i^{(t)}\}_{i=1}^G \subseteq \{\hat{o}_j^{(t)}\}_{j=1}^{\hat{G}}$$

$$\rho_{i,k}^{(t)}(\theta) = \frac{\pi_{\theta}(o_{i,k}^{(t)} | q, o_{i,<k}^{(t)})}{\pi_{\theta_{\text{ref}}}(o_{i,k}^{(t)} | q, o_{i,<k}^{(t)})}, \quad \hat{A}_i^{(t)} = \frac{\hat{r}_i^{(t)} - \mu^{(t)}}{\sigma^{(t)}}, \quad (4)$$

$$\hat{r}_i^{(T)} = r_i^{(T)}, \quad \hat{r}_i^{(T-1)} = r_i^{(T-1)} + r_i^{(\hat{T})}. \quad (5)$$

where $\pi_{\theta_{\text{old}}}$ generates the initial candidates and $\{o_j^{(t)}\}_{i=1}^G$ represents the uniformly sampled subset. k indexes the token position. t refers to the time step. The advantage estimation incorporates both immediate rewards and projected future returns. γ balances the relative importance of terminal and penultimate objectives. By focusing on these critical transition steps and maintaining computational efficiency through selective sampling, EGRPO enables effective learning in the episodic CoS framework. See Section D in the appendix for formal derivations, assumptions, and the algorithm.

4 Experiments

4.1 Training setup for SFT and RL

During the SFT stage, we fine-tune Qwen 2.5 VL 3B on the full SCoPE dataset with a maximum of 1,003,520 pixels. For the RL stage, we deploy EGRPO to optimize both overall performance and navigation efficiency for Chain of Scroll. SCoPE VLM EGRPO is trained with LoRA (Hu et al., 2021) for 2,500 steps using the same maximum pixel setting as the SFT stage. Further details can be found in Section G of the appendix.

4.2 Evaluation

The experimental evaluation of CoS focuses on assessing its effectiveness and efficiency on six distinct multi-page document question answering benchmarks with ANLS score. We conducted these evaluations under consistent constraints of the maximum number of visual tokens processed per step (Table 2, and 5, and 6) and per image (Table 3, and 4) with CoS and Multi Image (MI) inference. For evaluation, we have extended LMMs-Eval (Zhang et al., 2024b) to support DUDE (Lan-deghem et al., 2023), M3DocVQA, SlideVQA, and MMLongBench-Doc. Additionally, the LoRA weights for SCoPE VLM EGRPO are merged prior to evaluation to ensure accurate VRAM usage measurement. See Section I from the appendix for further details on the experiment setup.

4.3 Enhancing decision-making capabilities under Chain of Scroll framework

Challenges of CoS: In Table 2, the baseline 3B model shows a drastic performance drop in CoS from the conventional method. Even the 72B model scores lower than the 3B model with multi-image inference on average. This shows that even the largest scale model in Qwen2.5 VL series fails to generalize decision-making performance, although it shows moderate capabilities in GUI control and document question answering. This highlights the difficulties in question answering under agentic exploration. Furthermore, CoS consistently underperforms on DocVQA, a single-image document question answering benchmark (Mathew et al., 2021). For single-page document question answering, models do not require complex agentic navigation but rather rely solely on visual understanding. This confirms the observation from MPO (Wang et al., 2025) that although Chain of Thought effectively enhances logical reasoning, the distribution

Models	Type	DocVQA	MP-DocVQA	DUDE	M3DocVQA (single doc.)	SlideVQA	MMLongBench-Doc
Average #Images per Question	–	1	5.57	5.61	13.18	20.00	51.72
Open Source General VLMs							
LLAVA One Vision 7B	MI	87.05	59.29	35.60	35.42	48.73	9.53
Qwen 2.5 VL 3B	MI	92.54	76.28	46.14	43.72	51.46	14.02
Qwen 2.5 VL 7B	MI	94.61	80.99	49.16	52.54	57.22	15.21
Chain of Scroll (Ours)							
Qwen 2.5 VL 3B	CoS	66.65	47.78	31.73	19.54	23.23	9.08
Qwen 2.5 VL 72B	CoS	90.60	80.83	48.42	35.83	66.73	–
SCoPE VLM 3B SFT (Ours)	CoS	85.41	74.49	42.82	46.13	59.88	16.89
SCoPE VLM 3B EGRPO (Ours)	CoS	85.39	73.07	42.29	48.27	57.31	17.90

Table 2: Performance comparison of VLMs on multi-page document benchmarks. SCoPE VLMs and Qwen series used the identical maximum image tokens per generation step for comparable memory usage.

Model	Method	Visit Ratio	M3DocVQA (single doc.)			SlideVQA			
			ANLS	Max VRAM (GB)	ANLS / VRAM	Visit Ratio	ANLS	Max VRAM (GB)	ANLS / VRAM
Qwen 2.5 VL 3B	MI	1.0	53.43	78.98	0.67	1.0	64.95	40.76	1.59
Qwen 2.5 VL 7B	MI	1.0	61.32	79.41	0.77	1.0	71.82	51.35	1.39
Qwen 2.5 VL 32B	MI	1.0	53.67	278.40	0.19	1.0	69.50	232.84	0.30
Qwen 2.5 VL 72B	MI	1.0	63.69	582.17	0.11	1.0	75.45	323.52	0.23
Qwen 2.5 VL 3B	CoS	0.52	20.66	13.35	1.55	0.30	23.23	13.41	1.73
Qwen 2.5 VL 72B	CoS	0.59	40.96	160.48	0.26	0.50	66.73	170.63	0.39
SCoPE VLM 3B SFT (Ours)	CoS	0.54	46.26	14.59	3.17	0.56	59.84	13.42	4.45
SCoPE VLM 3B EGRPO (Ours)	CoS	0.45	45.58	13.72	3.32	0.52	59.13	13.12	4.51

Table 3: ANLS scores, memory usage, and average visited-page ratios per question for SCoPE VLM and Qwen 2.5 VL, evaluated on M3DocVQA and SlideVQA. All models are constrained to use the same maximum number of visual tokens per image.

Model	General					Web Shopping					Overall				
	Click	Scroll	Stop	Total	Goal	Click	Scroll	Stop	Total	Goal	Click	Scroll	Stop	Total	Goal
Qwen 2.5-VL-3B	30.21	6.57	66.67	33.53	35.46	37.76	39.01	83.57	41.49	41.17	33.99	22.79	75.12	37.51	38.32
SCoPE VLM 3B SFT (Ours)	31.00	8.76	73.08	34.28	36.10	37.95	36.77	87.14	40.83	41.56	34.48	22.77	80.11	37.56	38.83
SCoPE VLM 3B EGRPO (Ours)	33.55	8.03	76.28	35.69	37.51	38.31	45.29	84.29	42.33	42.32	35.93	26.66	80.29	39.01	39.92

Table 4: Performance results on the AitZ (Zhang et al., 2024a) across General and Web Shopping test splits. The table reports exact-match accuracies for Click, Scroll, Stop and their average (Total). Goal progress (Goal) indicates average task completion rate per step. All models are fine-tuned on the benchmark’s training set prior to evaluation.

shift between SFT and inference degrades multi-modal understanding.

Effectiveness of SCoPE dataset: Despite the great difficulties, our models successfully leverage the inference time scaling strategy and consistently outperform the equivalent baseline both in CoS and multi-image inference in benchmarks with the longer pages under similar memory resource limitations. In Table 2, the gains are well demonstrated on M3DocVQA and SlideVQA, where SCoPE VLM 3B SFT scores more than two times higher in ANLS (19.54 \rightarrow 46.13) than the Qwen 2.5VL 3B with CoS and 5% higher than the multi-image inference. This clearly supports the effectiveness of the proposed SCoPE dataset, showing substantial performance gains throughout the experiments. This further highlights the supremacy of the CoS framework as it performs better than the multi-image inference under similar restrictions of VRAM allocation. It is also notable that both SCoPE VLMs achieve performance comparable to the 72B model on longer documents. This shows that our models have com-

parable exploration abilities to the largest model in the Qwen2.5 VL series.

Memory Efficiency of SCoPE VLM: CoS ultimately transforms the multi-page document question answering problem into a single image. These characteristics of the CoS framework bring a significant increase in the performance per memory usage. Table 3 highlights the memory efficiency of SCoPE VLMs, which shows more than two times higher ANLS / VRAM on both benchmarks. In addition, SCoPE VLMs show significantly lower VRAM usage to complete the benchmarks, which is about three to five times less than the conventional method requires. Compared to the baseline, SCoPE VLM SFT exhibits an increase in visit ratio due to the enhanced exploration capabilities, which results in an increase in overall performance. Moreover, the additional EGRPO stage successfully mitigates this issue by reducing the visit ratio. This is demonstrated on M3DocVQA, where EGRPO achieves approximately 16.67% improvement over SFT. Despite fewer page visits, SCoPE VLM EGRPO main-

Model	MP-DocVQA						M3DocVQA						Average	
	Serial		Random		CoS		Serial		Random		CoS		Visit Ratio	ANLS
	Visit Ratio	ANLS												
SFT	22.02	40.25	21.65	46.80	75.59	74.49	44.65	41.22	48.82	39.21	53.40	46.13	44.36	48.02
GRPO (Last, 2)	20.05	39.83	20.36	45.20	71.40	73.14	31.69	38.51	35.31	37.30	46.21	46.71	37.50	46.78
EGRPO	18.85	39.93	18.50	44.76	69.97	73.90	37.14	39.74	37.80	41.32	47.20	47.67	38.24	47.89

Table 5: Comparison of Visit Ratio and ANLS across inference strategies on MP-DocVQA and M3DocVQA. All models achieve higher ANLS under CoS, demonstrating effective learned search strategies. EGRPO further achieves superior navigation efficiency and task performance, validating its effectiveness in optimizing both objectives.

Model	ANLS (%)						Visit Ratio (%)					
	MP-DocVQA	SlideVQA	M3DocVQA	DUDE	MMLong	Avg.	MP-DocVQA	SlideVQA	M3DocVQA	DUDE	MMLong	Avg.
SFT	74.49	59.88	46.13	42.82	16.89	48.04	75.59	55.79	53.40	82.40	122.68	77.97
GRPO (Last)	72.40	56.53	46.71	42.40	17.99	47.20	70.31	49.86	45.94	74.39	101.47	68.39
GRPO (Last, 2)	73.14	56.78	46.71	42.05	17.72	47.28	71.40	53.45	46.21	75.02	108.84	70.98
EGRPO	73.90	58.78	47.67	42.87	17.53	48.15	69.97	54.20	47.20	76.84	113.04	72.25

Model	Action Success Ratio (%)						No Answer Ratio (%)					
	MP-DocVQA	SlideVQA	M3DocVQA	DUDE	MMLong	Avg.	MP-DocVQA	SlideVQA	M3DocVQA	DUDE	MMLong	Avg.
SFT	94.42	78.49	87.39	90.90	62.52	82.74	1.33	0.79	3.50	3.71	31.44	8.15
GRPO (Last)	95.78	80.10	91.02	92.69	65.55	85.02	0.46	0.42	1.40	2.95	18.06	4.66
GRPO (Last, 2)	96.02	79.97	91.33	93.11	63.98	84.88	0.64	0.30	1.40	3.47	24.47	6.06
EGRPO	95.44	79.33	89.69	92.77	65.84	84.61	0.62	0.48	2.33	5.73	27.04	7.24

Table 6: Performance-efficiency comparison of SFT, GRPO, and EGRPO on the Chain of Scroll framework across five document understanding benchmarks. EGRPO maintains task performance comparable to SFT while improving efficiency metrics, whereas GRPO methods exhibit performance-efficiency trade-offs with reduced ANLS scores.

tains performance in ANLS, indicating its improved reasoning and decision-making capabilities.

Effectiveness of Document Navigation Capabilities on GUI Control: To evaluate whether the multimodal agentic reasoning and decision-making capabilities learned from Chain of Scroll generalize to other downstream tasks, we conducted experiments on the AitZ dataset, a mobile GUI control dataset. We focused on two key actions that serve as direct proxies for core competencies: Scroll actions for assessing navigation proficiency and Stop actions for measuring task completion accuracy. These actions provide ideal benchmarks for validating the transferability of our model’s navigation capabilities to new domains. As shown in Table 4, the two SCoPE VLM variants demonstrate strong adaptation performance. Especially, the SCoPE VLM EGRPO exhibits better adaptation than the SFT variant, achieving an overall average exact-match accuracy of 39.01%. This performance gain stems from substantial improvements over Qwen2.5 VL 3B in both targeted actions: Stop accuracy increased from 75.12% to 80.29%, while Scroll accuracy improved from 22.79% to 26.66%. These results provide compelling evidence that our document navigation capabilities can be effectively generalized to GUI control across different contexts. Although the performance gains in Scroll and Stop accuracy between SFT and EGRPO are mixed, EGRPO demonstrates better adaptation to ground truth actions overall and achieves notably higher goal progress rates. This further confirms the

effectiveness of learning the transition from searching steps to the terminal step in multimodal agents. Detailed experimental setup and comprehensive results are provided in Sections I and L, respectively.

4.4 Ablation study

To evaluate the effectiveness of the Chain of Scroll (CoS) framework and EGRPO, we conducted an ablation study as presented in Table 5 and Table 6. In Table 5, we compare CoS with two alternative strategies: serial and random processing. We implement random navigation by replacing the model’s scroll action output with a random scroll to assess whether our models learned effective search trajectories or simply visited pages without meaningful selection. In Table 6, we demonstrate the effectiveness of EGRPO by comparing it with SFT and GRPO variants. EGRPO and GRPO variant models in both tables are trained for a reduced step count of 1,000 with the same effective group size for gradient updates. Further details are provided in Section I in the appendix.

Inference Strategy Comparison: The evaluation results in Table 5 demonstrate that CoS consistently outperforms both Serial and Random approaches across the MP-DocVQA and M3DocVQA benchmarks. For MP-DocVQA, all compared models using CoS achieve significantly higher ANLS scores than the serial and random methods. For instance, the SFT model improves from 40.25 (Serial) and 46.80 (Random) to 74.49 (CoS). This substantial improvement demonstrates that the models have

learned effective search strategies, and this trend is consistently observed regardless of the post-training method employed. A similar trend is observed on M3DocVQA, where the SFT model improves from 41.22 (Serial) to 46.13 (CoS), with GRPO and EGRPO showing a similar tendency. This clearly demonstrates that all of the SCoPE VLM variants successfully learned effective context-aware, action-based reasoning for document navigation. Furthermore, EGRPO consistently shows lower page visit ratios compared to GRPO while maintaining high performance across all three inference strategies. This validates its effectiveness in optimizing both navigation efficiency and task performance.

Training Method Comparison: The extended analysis in Table 6 further confirms EGRPO’s effectiveness in addressing the limitations of SFT and vanilla GRPO variants. To validate the contributions of terminal step projection from EGRPO, we compared EGRPO with GRPO trained using only the terminal step (Last) and both the terminal and penultimate steps (Last, 2) as noted in the table.

In Table 6, the EGRPO model performs on par with SFT in ANLS scores (48.15 vs. 48.04 average) while significantly improving navigation efficiency. EGRPO reduces the average visit ratio from 77.97 to 72.25. The Action Success Ratio results show that EGRPO achieves consistent improvements in legal action selection. GRPO variants show improvements in the objectives but experience a clear performance drop. This is evidenced by MP-DocVQA and SlideVQA, where the GRPO (Last) model drops from 74.49 to 72.40 in MP-DocVQA and from 59.88 to 56.53 in SlideVQA.

Although the vanilla GRPO variants significantly sacrifice performance, they show better optimization in average page visits per question and answer return rate. Regarding the No Answer Ratio, vanilla GRPO variants optimize this metric more aggressively than EGRPO. GRPO (Last) achieved a No Answer Ratio of 4.66, whereas EGRPO achieved 7.24. However, this aggressive optimization comes at the cost of reduced ANLS performance. This performance gap suggests that GRPO variants are more prone to reward hacking and overfitting, leading to premature or speculative answers rather than thorough document exploration. In contrast, EGRPO delays returning an answer until it finds definitive evidence, resulting in more accurate responses overall. While the improvement in No Answer Ratio is more modest for EGRPO compared to GRPO variants, it still notably reduces this ratio compared to

the SFT baseline, from 8.15 to 7.24. This tendency is also observed in SCoPE VLM SFT and EGRPO. The performance-efficiency comparison of SCoPE VLMs and baseline models in Table 2 is detailed in Section K in the appendix.

These results validate that EGRPO effectively addresses the fundamental challenge of balancing exploration and exploitation in agentic document navigation. While vanilla GRPO variants achieve aggressive optimization of efficiency metrics, they do so at the expense of task performance, exhibiting patterns that prioritize quick termination over thorough evidence gathering. In contrast, EGRPO’s terminal step projection mechanism enables efficient backward propagation of reward signals through trajectories, delivering focused optimization toward answer-critical decisions while preserving the sample-efficient grouped updates of GRPO. This design enables EGRPO to achieve balanced optimization essential in real-world long document understanding scenarios.

5 Conclusion

In conclusion, this work addresses a critical gap in VLMs by introducing Chain of Scroll, the first framework to model agentic behaviors in multimodal document understanding. Through SCoPE dataset, we substantially improve the basic decision-making capabilities. With the Episodic Group Relative Policy Optimization, SCoPE VLM learns better Chain of Scroll trajectories, enabling effective long document question answering.

Through comprehensive evaluations, SCoPE VLM achieves high memory efficiency and competitive accuracy compared to state-of-the-art VLMs. Despite using fewer parameters and substantially less memory, our model demonstrates promising performance across benchmarks, underscoring the value of selective document exploration over exhaustive context processing. Furthermore, the strong adaptation to GUI control tasks validates the importance of document-based navigation pretraining for developing transferable agentic capabilities. Our ablation study confirms that our models learned effective reasoning and decision-making capabilities, while the additional EGRPO stage further optimizes navigation behavior by balancing efficiency with performance. By introducing an action-driven approach to document understanding, our work represents a significant step toward more capable and efficient multimodal agents.

6 Limitations

The major limitations of this work can be summarized in three key aspects: the scale of the training data, the variance in document lengths across the training and benchmark datasets, and the constrained demonstration of EGRPO’s potential due to the use of LoRA. First, the training data is notably limited, especially for final answer pages, which are restricted by the number of available questions (approximately 19K). When distributed across intermediate scroll steps, the number of answer-step samples becomes insufficient to fully train a 3B model from the SFT stage. In addition, due to the limited number of training document samples, there is a risk of amplifying the bias present in the training data. Using Gemini series alone for the annotation model may also have introduced bias in the generated reasoning. Second, both the training and benchmark datasets primarily consist of documents with fewer than 100 pages. This calls for further research to explore behavioral differences between short documents and longer, more contextually extensive documents. Furthermore, while LoRA helps reduce the training burden, it inherently restricts the potential of Episodic Group Relative Policy Optimization.

Another limitation concerns the trade-off between domain-specific adaptation and general capability. Recent works (Hu et al., 2024; Zhai et al., 2024) explore parameter updates primarily for document understanding or agentic tasks, often at the sacrifice of general ability for domain-specific adaptation. While SCoPE VLM may compromise generality just as those works, this can be mitigated through two strategies: (1) employing larger-scale base models that can generalize better even without additional training steps, as they already possess agentic decision-making capabilities, or (2) using blended SFT and RL that combine domain-specific and general-purpose corpora. Despite our method remaining suitable for local deployment due to its efficient parameter footprint, the potential trade-off between specialization and generalization in our domain-adapted sLLMs through the proposed Chain of Scroll framework requires further investigation.

Additionally, our current scope is limited to single-document QA. While SCoPE VLM supports multi-hop reasoning within a single document by actively performing memory-based aggregation during recursive scrolling, we do not address cross-

document reasoning. Cross-document reasoning represents an important capability gap and constitutes an exciting extension for future work.

References

- Hao Bai, Yifei Zhou, Mert Cemri, Jiayi Pan, Alane Suhr, Sergey Levine, and Aviral Kumar. 2024. [Digirl: Training in-the-wild device-control agents with autonomous reinforcement learning](#). *Preprint*, arXiv:2406.11896.
- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. 2025. [Qwen2.5-vl technical report](#). *Preprint*, arXiv:2502.13923.
- Ali Furkan Biten, Ruben Tito, Andres Mafla, Lluís Gomez, Marçal Rusiñol, Ernest Valveny, C. V. Jawahar, and Dimosthenis Karatzas. 2019. [Scene text visual question answering](#). *Preprint*, arXiv:1905.13648.
- Junbum Cha, Wooyoung Kang, Jonghwan Mun, and Byungseok Roh. 2023. Honeybee: Locality-enhanced projector for multimodal llm. *arXiv preprint arXiv:2312.06742*.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, Bin Li, Ping Luo, Tong Lu, Yu Qiao, and Jifeng Dai. 2023. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. *arXiv preprint arXiv:2312.14238*.
- Jaemin Cho, Debanjan Mahata, Ozan Irsoy, Yujie He, and Mohit Bansal. 2024. [M3docrag: Multi-modal retrieval is what you need for multi-page multi-document understanding](#). *Preprint*, arXiv:2411.04952.
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui

- Ding, Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Shengfeng Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanxia Zhao, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaosha Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanhong Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhen Xu, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *Preprint*, arXiv:2501.12948.
- Mostafa Dehghani, Basil Mustafa, Josip Djolonga, Jonathan Heek, Matthias Minderer, Mathilde Caron, Andreas Steiner, Joan Puigcerver, Robert Geirhos, Ibrahim Alabdulmohsin, Avital Oliver, Piotr Padlewski, Alexey Gritsenko, Mario Lučić, and Neil Houlsby. 2023. [Patch n’ pack: Navit, a vision transformer for any aspect ratio and resolution](#). *Preprint*, arXiv:2307.06304.
- Jarvis Guo, Tuney Zheng, Yuelin Bai, Bo Li, Yubo Wang, King Zhu, Yizhi Li, Graham Neubig, Wenhui Chen, and Xiang Yue. 2024. [Mammoth-vl: Eliciting multimodal reasoning with instruction tuning at scale](#). *Preprint*, arXiv:2412.05237.
- Anwen Hu, Haiyang Xu, Liang Zhang, Jiabo Ye, Ming Yan, Ji Zhang, Qin Jin, Fei Huang, and Jingren Zhou. 2024. [mplug-docowl2: High-resolution compressing for ocr-free multi-page document understanding](#). *Preprint*, arXiv:2409.03420.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. [Lora: Low-rank adaptation of large language models](#). *Preprint*, arXiv:2106.09685.
- Jing Yu Koh, Robert Lo, Lawrence Jang, Vikram Duvvur, Ming Chong Lim, Po-Yu Huang, Graham Neubig, Shuyan Zhou, Ruslan Salakhutdinov, and Daniel Fried. 2024. [Visualwebarena: Evaluating multimodal agents on realistic visual web tasks](#). *Preprint*, arXiv:2401.13649.
- Jordy Van Landeghem, Rubén Tito, Łukasz Borchmann, Michał Pietruszka, Paweł Józiać, Rafał Powalski, Dawid Jurkiewicz, Mickaël Coustaty, Bertrand Ackert, Ernest Valveny, Matthew Blaschko, Sien Moens, and Tomasz Stanisławek. 2023. [Document understanding dataset and evaluation \(dude\)](#). *Preprint*, arXiv:2305.08455.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2021. [Retrieval-augmented generation for knowledge-intensive nlp tasks](#). *Preprint*, arXiv:2005.11401.
- Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. 2024. [Llava-onevision: Easy visual task transfer](#). *Preprint*, arXiv:2408.03326.
- Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. 2024. [Llava-next: Improved reasoning, ocr, and world knowledge](#).
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. [Visual instruction tuning](#). *Preprint*, arXiv:2304.08485.
- Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. *Advances in Neural Information Processing Systems*, 35:2507–2521.
- Yubo Ma, Yuhang Zang, Liangyu Chen, Meiqi Chen, Yizhu Jiao, Xinze Li, Xinyuan Lu, Ziyu Liu, Yan Ma, Xiaoyi Dong, Pan Zhang, Liangming Pan, Yu-Gang Jiang, Jiaqi Wang, Yixin Cao, and Aixin Sun. 2024. [Mmlongbench-doc: Benchmarking long-context document understanding with visualizations](#). *Preprint*, arXiv:2407.01523.
- Minesh Mathew, Dimosthenis Karatzas, and C. V. Jawahar. 2021. [Docvqa: A dataset for vqa on document images](#). *Preprint*, arXiv:2007.00398.
- Fanqing Meng, Lingxiao Du, Zongkai Liu, Zhixiang Zhou, Quanfeng Lu, Daocheng Fu, Botian Shi, Wenhui Wang, Junjun He, Kaipeng Zhang, et al. 2025. [Mm-eureka: Exploring visual aha moment with rule-based large-scale reinforcement learning](#). *arXiv preprint arXiv:2503.07365*.

Yao Mu, Qinglong Zhang, Mengkang Hu, Wenhai Wang, Mingyu Ding, Jun Jin, Bin Wang, Jifeng Dai, Yu Qiao, and Ping Luo. 2023. Embodiedgpt: Vision-language pre-training via embodied chain of thought. *arXiv preprint arXiv:2305.15021*.

OpenAI, :, Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Mądry, Alex Baker-Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, Alex Nichol, Alex Paino, Alex Renzin, Alex Tachard Passos, Alexander Kirillov, Alexi Christakis, Alexis Conneau, Ali Kamali, Allan Jabri, Allison Moyer, Allison Tam, Amadou Crookes, Amin Tootoochian, Amin Tootoonchian, Ananya Kumar, Andrea Vallone, Andrej Karpathy, Andrew Braunstein, Andrew Cann, Andrew Codispoti, Andrew Galu, Andrew Kondrich, Andrew Tulloch, Andrey Mishchenko, Angela Baek, Angela Jiang, Antoine Pelisse, Antonia Woodford, Anuj Gosalia, Arka Dhar, Ashley Pantuliano, Avi Nayak, Avital Oliver, Barret Zoph, Behrooz Ghorbani, Ben Leimberger, Ben Rossen, Ben Sokolowsky, Ben Wang, Benjamin Zweig, Beth Hoover, Blake Samic, Bob McGrew, Bobby Spero, Bogo Giertler, Bowen Cheng, Brad Lightcap, Brandon Walkin, Brendan Quinn, Brian Guarraci, Brian Hsu, Bright Kellogg, Brydon Eastman, Camillo Lugaresi, Carroll Wainwright, Cary Bassin, Cary Hudson, Casey Chu, Chad Nelson, Chak Li, Chan Jun Shern, Channing Conger, Charlotte Barette, Chelsea Voss, Chen Ding, Cheng Lu, Chong Zhang, Chris Beaumont, Chris Hallacy, Chris Koch, Christian Gibson, Christina Kim, Christine Choi, Christine McLeavey, Christopher Hesse, Claudia Fischer, Clemens Winter, Coley Czarnecki, Colin Jarvis, Colin Wei, Constantin Koumouzelis, Dane Sherburn, Daniel Kappler, Daniel Levin, Daniel Levy, David Carr, David Farhi, David Mely, David Robinson, David Sasaki, Denny Jin, Dev Valladares, Dimitris Tsipras, Doug Li, Duc Phong Nguyen, Duncan Findlay, Edede Oiwoh, Edmund Wong, Ehsan Asdar, Elizabeth Proehl, Elizabeth Yang, Eric Antonow, Eric Kramer, Eric Peterson, Eric Sigler, Eric Wallace, Eugene Brevdo, Evan Mays, Farzad Khorasani, Felipe Petroski Such, Filippo Raso, Francis Zhang, Fred von Lohmann, Freddie Sulit, Gabriel Goh, Gene Oden, Geoff Salmon, Giulio Starace, Greg Brockman, Hadi Salman, Haiming Bao, Haitang Hu, Hannah Wong, Haoyu Wang, Heather Schmidt, Heather Whitney, Heewoo Jun, Hendrik Kirchner, Henrique Ponde de Oliveira Pinto, Hongyu Ren, Huiwen Chang, Hyung Won Chung, Ian Kivlichan, Ian O’Connell, Ian O’Connell, Ian Osband, Ian Silber, Ian Sohl, Ibrahim Okuyucu, Ikai Lan, Ilya Kostrikov, Ilya Sutskever, Ingmar Kanitscheider, Ishaan Gulrajani, Jacob Coxon, Jacob Menick, Jakub Pachocki, James Aung, James Betker, James Crooks, James Lennon, Jamie Kiros, Jan Leike, Jane Park, Jason Kwon, Jason Phang, Jason Teplitz, Jason Wei, Jason Wolfe, Jay Chen, Jeff Harris, Jenia Varavva, Jessica Gan Lee, Jessica Shieh, Ji Lin, Jiahui Yu, Jiayi Weng, Jie Tang, Jieqi Yu, Joanne Jang, Joaquin Quinonero Candela, Joe Beutler, Joe Landers, Joel Parish, Johannes Heidecke,

John Schulman, Jonathan Lachman, Jonathan McKay, Jonathan Uesato, Jonathan Ward, Jong Wook Kim, Joost Huizinga, Jordan Sitkin, Jos Kraaijeveld, Josh Gross, Josh Kaplan, Josh Snyder, Joshua Achiam, Joy Jiao, Joyce Lee, Juntang Zhuang, Justyn Harriman, Kai Fricke, Kai Hayashi, Karan Singhal, Katy Shi, Kavin Karthik, Kayla Wood, Kendra Rimbach, Kenny Hsu, Kenny Nguyen, Keren Gu-Lemberg, Kevin But-ton, Kevin Liu, Kiel Howe, Krithika Muthukumar, Kyle Luther, Lama Ahmad, Larry Kai, Lauren Itow, Lauren Workman, Leher Pathak, Leo Chen, Li Jing, Lia Guy, Liam Fedus, Liang Zhou, Lien Mamit-suka, Lilian Weng, Lindsay McCallum, Lindsey Held, Long Ouyang, Louis Feувrier, Lu Zhang, Lukas Kon-draciuk, Lukasz Kaiser, Luke Hewitt, Luke Metz, Lyric Doshi, Mada Aflak, Maddie Simens, Madelaine Boyd, Madeleine Thompson, Marat Dukhan, Mark Chen, Mark Gray, Mark Hudnall, Marvin Zhang, Mar-wan Aljubeih, Mateusz Litwin, Matthew Zeng, Max Johnson, Maya Shetty, Mayank Gupta, Meghan Shah, Mehmet Yatbaz, Meng Jia Yang, Mengchao Zhong, Mia Glaese, Mianna Chen, Michael Janner, Michael Lampe, Michael Petrov, Michael Wu, Michele Wang, Michelle Fradin, Michelle Pokrass, Miguel Castro, Miguel Oom Temudo de Castro, Mikhail Pavlov, Miles Brundage, Miles Wang, Minal Khan, Mira Mu-rati, Mo Bavarian, Molly Lin, Murat Yesildal, Nacho Soto, Natalia Gimelshein, Natalie Cone, Natalie Stau-dacher, Natalie Summers, Natan LaFontaine, Neil Chowdhury, Nick Ryder, Nick Stathas, Nick Tur-ley, Nik Tezak, Niko Felix, Nithanth Kudige, Nitish Keskar, Noah Deutsch, Noel Bundick, Nora Puck-ett, Ofir Nachum, Ola Okelola, Oleg Boiko, Oleg Murk, Oliver Jaffe, Olivia Watkins, Olivier Gode-ment, Owen Campbell-Moore, Patrick Chao, Paul McMillan, Pavel Belov, Peng Su, Peter Bak, Peter Bakkum, Peter Deng, Peter Dolan, Peter Hoeschele, Peter Welinder, Phil Tillet, Philip Pronin, Philippe Tillet, Prafulla Dhariwal, Qiming Yuan, Rachel Dias, Rachel Lim, Rahul Arora, Rajan Troll, Randall Lin, Rapha Gontijo Lopes, Raul Puri, Reah Miyara, Reimar Leike, Renaud Gaubert, Reza Zamani, Ricky Wang, Rob Donnelly, Rob Honsby, Rocky Smith, Ro-han Sahai, Rohit Ramchandani, Romain Huet, Rory Carmichael, Rowan Zellers, Roy Chen, Ruby Chen, Ruslan Nigmatullin, Ryan Cheu, Saachi Jain, Sam Altman, Sam Schoenholz, Sam Toizer, Samuel Mis-erendino, Sandhini Agarwal, Sara Culver, Scott Ether-smith, Scott Gray, Sean Grove, Sean Metzger, Shamez Hermani, Shantanu Jain, Shengjia Zhao, Sherwin Wu, Shino Jomoto, Shirong Wu, Shuaiqi, Xia, Sonia Phene, Spencer Papay, Srinivas Narayanan, Steve Cof-fey, Steve Lee, Stewart Hall, Suchir Balaji, Tal Broda, Tal Stramer, Tao Xu, Tarun Gogineni, Taya Christian-son, Ted Sanders, Tejal Patwardhan, Thomas Cun-ninghman, Thomas Degry, Thomas Dimson, Thomas Raoux, Thomas Shadwell, Tianhao Zheng, Todd Un-derwood, Todor Markov, Toki Sherbakov, Tom Ru-bin, Tom Stasi, Tomer Kaftan, Tristan Heywood, Troy Peterson, Tyce Walters, Tyna Eloundou, Valerie Qi, Veit Moeller, Vinnie Monaco, Vishal Kuo, Vlad Fomenko, Wayne Chang, Weiye Zheng, Wenda Zhou, Wesam Manassra, Will Sheu, Wojciech Zaremba, Yash Patil, Yilei Qian, Yongjik Kim, Youlong Cheng,

- Yu Zhang, Yuchen He, Yuchen Zhang, Yujia Jin, Yunxing Dai, and Yury Malkov. 2024. [Gpt-4o system card](#). *Preprint*, arXiv:2410.21276.
- OpenAI. 2023. [Gpt-4v](#).
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024. [Deepseekmath: Pushing the limits of mathematical reasoning in open language models](#). *Preprint*, arXiv:2402.03300.
- Xiaoqian Shen, Yunyang Xiong, Changsheng Zhao, Lemeng Wu, Jun Chen, Chenchen Zhu, Zechun Liu, Fanyi Xiao, Balakrishnan Varadarajan, Florian Bordes, Zhuang Liu, Hu Xu, Hyunwoo J. Kim, Bilge Soran, Raghuraman Krishnamoorthi, Mohamed Elhoseiny, and Vikas Chandra. 2024. [Longvu: Spatiotemporal adaptive compression for long video-language understanding](#). *Preprint*, arXiv:2410.17434.
- Qiushi Sun, Kanzhi Cheng, Zichen Ding, Chuanyang Jin, Yian Wang, Fangzhi Xu, Zhenyu Wu, Chengyou Jia, Liheng Chen, Zhounianze Liu, Ben Kao, Guohao Li, Junxian He, Yu Qiao, and Zhiyong Wu. 2025. [Os-genesis: Automating gui agent trajectory construction via reverse task synthesis](#). *Preprint*, arXiv:2412.19723.
- Ryota Tanaka, Kyosuke Nishida, Kosuke Nishida, Taku Hasegawa, Itsumi Saito, and Kuniko Saito. 2023. [Slidevqa: A dataset for document visual question answering on multiple images](#). *Preprint*, arXiv:2301.04883.
- Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, Soroosh Mariooryad, Yifan Ding, Xinyang Geng, Fred Alcober, Roy Frostig, Mark Omernick, Lexi Walker, Cosmin Paduraru, Christina Sorokin, Andrea Tacchetti, Colin Gaffney, Samira Daruki, Olcan Serincinoglu, Zach Gleicher, Juliette Love, Paul Voigtlaender, Rohan Jain, Gabriela Surita, Kareem Mohamed, Rory Blevins, Junwhan Ahn, Tao Zhu, Kornaphop Kawintiranon, Orhan Firat, Yiming Gu, Yujing Zhang, Matthew Rahtz, Manaal Faruqui, Natalie Clay, Justin Gilmer, JD Co-Reyes, Ivo Penchev, Rui Zhu, Nobuyuki Morioka, Kevin Hui, Krishna Haridasan, Victor Campos, Mahdis Mahdieh, Mandy Guo, Samer Hassan, Kevin Kilgour, Arpi Vezer, Heng-Tze Cheng, Raoul de Liedekerke, Siddharth Goyal, Paul Barham, DJ Strouse, Seb Noury, Jonas Adler, Mukund Sundararajan, Sharad Vikram, Dmitry Lepikhin, Michela Paganini, Xavier Garcia, Fan Yang, Dasha Valter, Maja Trebacz, Kiran Vodrahalli, Chulayuth Asawaroengchai, Roman Ring, Norbert Kalb, Livio Baldini Soares, Siddhartha Brahma, David Steiner, Tianhe Yu, Fabian Mentzer, Antoine He, Lucas Gonzalez, Bibo Xu, Raphael Lopez Kaufman, Laurent El Shafey, Junhyuk Oh, Tom Hennigan, George van den Driessche, Seth Odoom, Mario Lucic, Becca Roelofs, Sid Lall, Amit Marathe, Betty Chan, Santiago Ontanon, Luheng He, Denis Teplyashin, Jonathan Lai, Phil Crone, Bogdan Damoc, Lewis Ho, Sebastian Riedel, Karel Lenc, Chih-Kuan Yeh, Aakanksha Chowdhery, Yang Xu, Mehran Kazemi, Ehsan Amid, Anastasia Petrushkina, Kevin Swersky, Ali Khodaei, Gowoon Chen, Chris Larkin, Mario Pinto, Geng Yan, Adria Puigdomenech Badia, Piyush Patil, Steven Hansen, Dave Orr, Sebastien M. R. Arnold, Jordan Grimstad, Andrew Dai, Sholto Douglas, Rishika Sinha, Vikas Yadav, Xi Chen, Elena Gribovskaya, Jacob Austin, Jeffrey Zhao, Kaushal Patel, Paul Komarek, Sophia Austin, Sebastian Borgeaud, Linda Friso, Abhimanyu Goyal, Ben Caine, Kris Cao, Da-Woon Chung, Matthew Lamm, Gabe Barth-Maron, Thais Kagohara, Kate Olszewska, Mia Chen, Kaushik Shivakumar, Rishabh Agarwal, Harshal Godhia, Ravi Rajwar, Javier Snaider, Xerxes Dotiwalla, Yuan Liu, Aditya Barua, Victor Ungureanu, Yuan Zhang, Bat-Orgil Batsaikhan, Mateo Wirth, James Qin, Ivo Danihelka, Tulsee Doshi, Martin Chadwick, Jilin Chen, Sanil Jain, Quoc Le, Arjun Kar, Madhu Gurumurthy, Cheng Li, Ruoxin Sang, Fangyu Liu, Lampros Lamprou, Rich Munoz, Nathan Lintz, Harsh Mehta, Heidi Howard, Malcolm Reynolds, Lora Aroyo, Quan Wang, Lorenzo Blanco, Albin Cassirer, Jordan Griffith, Dipanjan Das, Stephan Lee, Jakub Sygnowski, Zach Fisher, James Besley, Richard Powell, Zafarali Ahmed, Dominik Paulus, David Reiter, Zalan Borsos, Rishabh Joshi, Aedan Pope, Steven Hand, Vittorio Selo, Vihan Jain, Nikhil Sethi, Megha Goel, Takaki Makino, Rhys May, Zhen Yang, Johan Schalkwyk, Christina Butterfield, Anja Hauth, Alex Goldin, Will Hawkins, Evan Senter, Sergey Brin, Oliver Woodman, Marvin Ritter, Eric Noland, Minh Giang, Vijay Bolina, Lisa Lee, Tim Blyth, Ian Mackinnon, Machel Reid, Obaid Sarvana, David Silver, Alexander Chen, Lily Wang, Loren Maggiore, Oscar Chang, Nithya Attaluri, Gregory Thornton, Chung-Cheng Chiu, Oskar Bunyan, Nir Levine, Timothy Chung, Evgenii Eltyshev, Xiance Si, Timothy Lillicrap, Demetra Brady, Vaibhav Aggarwal, Boxi Wu, Yuanzhong Xu, Ross McIlroy, Kartikeya Badola, Paramjit Sandhu, Erica Moreira, Wojciech Stokowiec, Ross Hemsley, Dong Li, Alex Tudor, Pranav Shyam, Elahe Rahimtoroghi, Salem Haykal, Pablo Sprechmann, Xiang Zhou, Diana Mincu, Yujia Li, Ravi Addanki, Kalpesh Krishna, Xiao Wu, Alexandre Frechette, Matan Eyal, Allan Dafoe, Dave Lacey, Jay Whang, Thi Avrahami, Ye Zhang, Emanuel Taropa, Hanzhao Lin, Daniel Toyama, Eliza Rutherford, Motoki Sano, HyunJeong Choe, Alex Tomala, Chalance Safranek-Shrader, Nora Kassner, Mantas Pajarskas, Matt Harvey, Sean Sechrist, Meire Fortunato, Christina Lyu, Gamaleldin Elsayed, Chenkai Kuang, James Lottes, Eric Chu, Chao Jia, Chih-Wei Chen, Peter Humphreys, Kate Baumli, Connie Tao, Rajkumar Samuel, Cicero Nogueira dos Santos, Anders Andreassen, Nemanja Rakićević, Dominik Grewe, Aviral Kumar, Stephanie Winkler, Jonathan Caton, Andrew Brock, Sid Dalmia, Hannah Sheahan, Iain Barr, Yingjie Miao, Paul Natsev, Jacob Devlin, Fer- yal Behbahani, Flavien Prost, Yanhua Sun, Artiom Myaskovsky, Thanumalayan Sankaranarayanan Pillai, Dan Hurt, Angeliki Lazaridou, Xi Xiong, Ce Zheng, Fabio Pardo, Xiaowei Li, Dan Horgan, Joe Stanton,

Moran Ambar, Fei Xia, Alejandro Lince, Mingqiu Wang, Basil Mustafa, Albert Webson, Hyo Lee, Rohan Anil, Martin Wicke, Timothy Dozat, Abhishek Sinha, Enrique Piqueras, Elahe Dabir, Shyam Upadhyay, Anudhyan Boral, Lisa Anne Hendricks, Corey Fry, Josip Djolonga, Yi Su, Jake Walker, Jane Labanowski, Ronny Huang, Vedant Misra, Jeremy Chen, RJ Skerry-Ryan, Avi Singh, Shruti Rijhwani, Dian Yu, Alex Castro-Ros, Beer Changpinyo, Romina Datta, Sumit Bagri, Arnar Mar Hrafnkelsson, Marcello Maggioni, Daniel Zheng, Yury Sulsky, Shaobo Hou, Tom Le Paine, Antoine Yang, Jason Riesa, Dominika Rogozinska, Dror Marcus, Dalia El Badawy, Qiao Zhang, Luyu Wang, Helen Miller, Jeremy Greer, Lars Lowe Sjos, Azade Nova, Heiga Zen, Rahma Chaabouni, Mihaela Rosca, Jiepu Jiang, Charlie Chen, Ruibo Liu, Tara Sainath, Maxim Krikun, Alex Polozov, Jean-Baptiste Lespiau, Josh Newlan, Zeynep Cankara, Soo Kwak, Yunhan Xu, Phil Chen, Andy Coenen, Clemens Meyer, Katerina Tsihla, Ada Ma, Juraj Gottweis, Jinwei Xing, Chenjie Gu, Jin Miao, Christian Frank, Zeynep Cankara, Sanjay Ganapathy, Ishita Dasgupta, Steph Hughes-Fitt, Heng Chen, David Reid, Keran Rong, Hongmin Fan, Joost van Amersfoort, Vincent Zhuang, Aaron Cohen, Shixiang Shane Gu, Anhad Mohananey, Anastasija Ilic, Taylor Tobin, John Wieting, Anna Bortsova, Phoebe Thacker, Emma Wang, Emily Caveness, Justin Chiu, Eren Sezener, Alex Kaskasoli, Steven Baker, Katie Millican, Mohamed Elhawaty, Kostas Aisopos, Carl Lebsack, Nathan Byrd, Hanjun Dai, Wenhao Jia, Matthew Wiethoff, Elnaz Davoodi, Albert Weston, Lakshman Yagati, Arun Ahuja, Isabel Gao, Golan Pundak, Susan Zhang, Michael Azzam, Khe Chai Sim, Sergi Caelles, James Keeling, Abhanshu Sharma, Andy Swing, YaGuang Li, Chenxi Liu, Carrie Grimes Bostock, Yamini Bansal, Zachary Nado, Ankesh Anand, Josh Lipschultz, Abhijit Karmarkar, Lev Proleev, Abe Ittycheriah, Soheil Hassas Yeganeh, George Polovets, Aleksandra Faust, Jiao Sun, Alban Rrustemi, Pen Li, Rakesh Shivanna, Jeremiah Liu, Chris Welty, Federico Lebron, Anirudh Baddepudi, Sebastian Krause, Emilio Parisotto, Radu Soricut, Zheng Xu, Dawn Bloxwich, Melvin Johnson, Behnam Neyshabur, Justin Mao-Jones, Renshen Wang, Vinay Ramasesh, Zaheer Abbas, Arthur Guez, Constant Segal, Duc Dung Nguyen, James Svensson, Le Hou, Sarah York, Kieran Milan, Sophie Bridgers, Wiktor Gworek, Marco Tagliasacchi, James Lee-Thorp, Michael Chang, Alexey Guseynov, Ale Jakse Hartman, Michael Kwong, Ruizhe Zhao, Sheleem Kashem, Elizabeth Cole, Antoine Miech, Richard Tanburn, Mary Phuong, Filip Pavetic, Sebastien Cevey, Ramona Comanescu, Richard Ives, Sherry Yang, Cosmo Du, Bo Li, Zizhao Zhang, Mariko Iinuma, Clara Huiyi Hu, Aurko Roy, Shaan Bijwadia, Zhenkai Zhu, Danilo Martins, Rachel Saputro, Anita Gergely, Steven Zheng, Dawei Jia, Ioannis Antonoglou, Adam Sadovskiy, Shane Gu, Yingying Bi, Alek Andreev, Sina Samangoeei, Mina Khan, Tomas Kocisky, Angelos Filos, Chintu Kumar, Colton Bishop, Adams Yu, Sarah Hodkinson, Sid Mittal, Premal Shah, Alexandre Moufarek, Yong Cheng, Adam Bloniarz, Jaehoon Lee, Pedram Pejman, Paul Michel,

Stephen Spencer, Vladimir Feinberg, Xuehan Xiong, Nikolay Savinov, Charlotte Smith, Siamak Shakeri, Dustin Tran, Mary Chesus, Bernd Bohnet, George Tucker, Tamara von Glehn, Carrie Muir, Yiran Mao, Hideto Kazawa, Ambrose Slone, Kedar Soparkar, Disha Shrivastava, James Cobon-Kerr, Michael Sherman, Jay Pavagadhi, Carlos Araya, Karolis Misiunas, Nimesh Ghelani, Michael Laskin, David Barker, Qiuji Li, Anton Briukhov, Neil Houlsby, Mia Glaese, Balaji Lakshminarayanan, Nathan Schucher, Yunhao Tang, Eli Collins, Hyeontaek Lim, Fangxiaoyu Feng, Adria Recasens, Guangda Lai, Alberto Magni, Nicola De Cao, Aditya Siddhant, Zoe Ashwood, Jordi Orbay, Mostafa Dehghani, Jenny Brennan, Yifan He, Kelvin Xu, Yang Gao, Carl Saroufim, James Mollo, Xinyi Wu, Seb Arnold, Solomon Chang, Julian Schrittwieser, Elena Buchatskaya, Soroush Radpour, Martin Polacek, Skye Giordano, Ankur Bapna, Simon Tokumine, Vincent Hellendoorn, Thibault Sottiaux, Sarah Cogan, Aliaksei Severyn, Mohammad Saleh, Shantanu Thakoor, Laurent Shefey, Siyuan Qiao, Meenu Gaba, Shuo yiin Chang, Craig Swanson, Biao Zhang, Benjamin Lee, Paul Kishan Rubenstein, Gan Song, Tom Kwiatkowski, Anna Koop, Ajay Kannan, David Kao, Parker Schuh, Axel Stjerngren, Golnaz Ghiasi, Gena Gibson, Luke Vilnis, Ye Yuan, Felipe Tiengo Ferreira, Aishwarya Kamath, Ted Klimentko, Ken Franko, Kefan Xiao, Indro Bhattacharya, Miteyan Patel, Rui Wang, Alex Morris, Robin Strudel, Vivek Sharma, Peter Choy, Sayed Hashemi, Jessica Landon, Mara Finkelstein, Priya Jhakra, Justin Frye, Megan Barnes, Matthew Mauger, Dennis Daun, Khuslen Baatarsukh, Matthew Tung, Wael Farhan, Henryk Michalewski, Fabio Viola, Felix de Chaumont Quitry, Charline Le Lan, Tom Hudson, Qingze Wang, Felix Fischer, Ivy Zheng, Elspeth White, Anca Dragan, Jean baptiste Alayrac, Eric Ni, Alexander Pritzel, Adam Iwanicki, Michael Isard, Anna Bulanova, Lukas Zilka, Ethan Dyer, Deendra Sachan, Srivatsan Srinivasan, Hannah Muckenhirn, Honglong Cai, Amol Mandhane, Mukarram Tariq, Jack W. Rae, Gary Wang, Kareem Ayoub, Nicholas FitzGerald, Yao Zhao, Woohyun Han, Chris Alberti, Dan Garrette, Kashyap Krishnakumar, Mai Gimenez, Anselm Levskaya, Daniel Sohn, Josip Matak, Inaki Iturrate, Michael B. Chang, Jackie Xi-ang, Yuan Cao, Nishant Ranka, Geoff Brown, Adrian Hutter, Vahab Mirrokni, Nanxin Chen, Kaisheng Yao, Zoltan Egyed, Francois Galilee, Tyler Liechty, Praveen Kallakuri, Evan Palmer, Sanjay Ghemawat, Jasmine Liu, David Tao, Chloe Thornton, Tim Green, Mimi Jasarevic, Sharon Lin, Victor Cotruta, Yi-Xuan Tan, Noah Fiedel, Hongkun Yu, Ed Chi, Alexander Neitz, Jens Heitkaemper, Anu Sinha, Denny Zhou, Yi Sun, Charbel Kaed, Brice Hulse, Swaroop Mishra, Maria Georgaki, Sneha Kudugunta, Clement Farabet, Izhak Shafran, Daniel Vlasic, Anton Tsitsulin, Rajagopal Ananthanarayanan, Alen Carin, Guolong Su, Pei Sun, Shashank V, Gabriel Carvajal, Josef Broder, Iulia Comsa, Alena Repina, William Wong, Warren Weilun Chen, Peter Hawkins, Egor Filonov, Lucia Loher, Christoph Hirsenschall, Weiyi Wang, Jingchen Ye, Andrea Burns, Hardie Cate, Diana Gage Wright, Federico Piccinini, Lei

Zhang, Chu-Cheng Lin, Ionel Gog, Yana Kulizhskaya, Ashwin Sreevatsa, Shuang Song, Luis C. Cobo, Anand Iyer, Chetan Tekur, Guillermo Garrido, Zhuyun Xiao, Rupert Kemp, Huaixiu Steven Zheng, Hui Li, Ananth Agarwal, Christel Ngani, Kati Goshvadi, Rebeca Santamaria-Fernandez, Wojciech Fica, Xinyun Chen, Chris Gorgolewski, Sean Sun, Roopal Garg, Xinyu Ye, S. M. Ali Eslami, Nan Hua, Jon Simon, Pratik Joshi, Yelin Kim, Ian Tenney, Sahitya Potluri, Lam Nguyen Thiet, Quan Yuan, Florian Luisier, Alexandra Chronopoulou, Salvatore Scellato, Praveen Srinivasan, Minmin Chen, Vinod Koverkathu, Valentin Dalibard, Yaming Xu, Brennan Saeta, Keith Anderson, Thibault Sellam, Nick Fernando, Fantine Huot, Junehyuk Jung, Mani Varadarajan, Michael Quinn, Amit Raul, Maigo Le, Ruslan Habalov, Jon Clark, Komal Jalan, Kalesha Bullard, Achintya Singhal, Thang Luong, Boyu Wang, Sujeevan Rajayogam, Julian Eisenschlos, Johnson Jia, Daniel Finchelstein, Alex Yakubovich, Daniel Balle, Michael Fink, Sameer Agarwal, Jing Li, Dj Dvijotham, Shalini Pal, Kai Kang, Jaclyn Konzelmann, Jennifer Beattie, Olivier Dousse, Diane Wu, Remi Crocker, Chen Elkind, Siddhartha Reddy Jonnalagadda, Jong Lee, Dan Holtmann-Rice, Krystal Kallarackal, Rosanne Liu, Denis Vnukov, Neera Vats, Luca Invernizzi, Mohsen Jafari, Huanjie Zhou, Lilly Taylor, Jennifer Prendki, Marcus Wu, Tom Eccles, Tianqi Liu, Kavaya Koppurapu, Francoise Beaufays, Christof Angermueller, Andreea Marzoca, Shourya Sarcar, Hilal Dib, Jeff Stanway, Frank Perbet, Nejc Trdin, Rachel Sterneck, Andrey Khorlin, Dinghua Li, Xihui Wu, Sonam Goenka, David Madras, Sasha Goldshtein, Willi Gierke, Tong Zhou, Yaxin Liu, Yannie Liang, Anais White, Yunjie Li, Shreya Singh, Sanaz Bahargam, Mark Epstein, Sujoy Basu, Li Lao, Adnan Ozturel, Carl Crous, Alex Zhai, Han Lu, Zora Tung, Neeraj Gaur, Alanna Walton, Lucas Dixon, Ming Zhang, Amir Globerson, Grant Uy, Andrew Bolt, Olivia Wiles, Milad Nasr, Ilia Shumailov, Marco Selvi, Francesco Piccinno, Ricardo Aguilar, Sara McCarthy, Misha Khalman, Mrinal Shukla, Vlado Galic, John Carpenter, Kevin Vellela, Haibin Zhang, Harry Richardson, James Martens, Matko Bosnjak, Shreyas Rammohan Belle, Jeff Seibert, Mahmoud Alnahlawi, Brian McWilliams, Sankalp Singh, Annie Louis, Wen Ding, Dan Popovici, Lenin Simicich, Laura Knight, Pulkit Mehta, Nishesh Gupta, Chongyang Shi, Saaber Fatehi, Jovana Mitrovic, Alex Grills, Joseph Pagadora, Tsendsuren Munkhdalai, Dessie Petrova, Danielle Eisenbud, Zhishuai Zhang, Damion Yates, Bhavishya Mittal, Nilesh Tripuraneni, Yannis Assael, Thomas Brovelli, Prateek Jain, Mihajlo Velimirovic, Canfer Akbulut, Jiaqi Mu, Wolfgang Macherey, Ravin Kumar, Jun Xu, Haroon Qureshi, Gheorghe Comanici, Jeremy Wiesner, Zhitao Gong, Anton Ruddock, Matthias Bauer, Nick Felt, Anirudh GP, Anurag Arnab, Dustin Zelle, Jonas Rothfuss, Bill Rosgen, Ashish Shenoy, Bryan Seybold, Xinjian Li, Jayaram Mudigonda, Goker Erdogan, Jiawei Xia, Jiri Simsa, Andrea Michi, Yi Yao, Christopher Yew, Steven Kan, Isaac Caswell, Carey Radebaugh, Andre Elisseeff, Pedro Valenzuela, Kay McKinney, Kim Paterson, Albert Cui, Eri Latorre-Chimoto, Solomon

Kim, William Zeng, Ken Durden, Priya Ponnappalli, Tiberiu Sosea, Christopher A. Choquette-Choo, James Manyika, Brona Robenek, Harsha Vashisht, Sebastien Pereira, Hoi Lam, Marko Velic, Denese Owusu-Afriyie, Katherine Lee, Tolga Bolukbasi, Alicia Parrish, Shawn Lu, Jane Park, Balaji Venkatraman, Alice Talbert, Lambert Rosique, Yuchung Cheng, Andrei Sozanschi, Adam Paszke, Praveen Kumar, Jessica Austin, Lu Li, Khalid Salama, Bartek Perz, Wooyeol Kim, Nandita Dukkupati, Anthony Baryshnikov, Christos Kaplanis, XiangHai Sheng, Yuri Chervonyi, Caglar Unlu, Diego de Las Casas, Harry Askham, Kathryn Tunyasuvunakool, Felix Gimeno, Siim Poder, Chester Kwak, Matt Miecnikowski, Vahab Mirrokni, Alek Dimitriev, Aaron Parisi, Dangyi Liu, Tomy Tsai, Toby Shevlane, Christina Kouridi, Drew Garmon, Adrian Goedeckemeyer, Adam R. Brown, Anitha Vijayakumar, Ali Elqursh, Sadegh Jazayeri, Jin Huang, Sara Mc Carthy, Jay Hoover, Lucy Kim, Sandeep Kumar, Wei Chen, Courtney Biles, Garrett Bingham, Evan Rosen, Lisa Wang, Qijun Tan, David Engel, Francesco Pongetti, Dario de Cesare, Dongseong Hwang, Lily Yu, Jennifer Pullman, Srinu Narayanan, Kyle Levin, Siddharth Gopal, Megan Li, Asaf Aharoni, Trieu Trinh, Jessica Lo, Norman Casagrande, Roopali Vij, Loic Matthey, Bramandia Ramadhana, Austin Matthews, CJ Carey, Matthew Johnson, Kremena Goranova, Rohin Shah, Shereen Ashraf, Kingshuk Dasgupta, Rasmus Larsen, Yicheng Wang, Manish Reddy Vuyyuru, Chong Jiang, Joana Ijazi, Kazuki Osawa, Celine Smith, Ramya Sree Boppana, Taylan Bilal, Yuma Koizumi, Ying Xu, Yasemin Altun, Nir Shabat, Ben Bariach, Alex Korchemniy, Kiam Choo, Olaf Ronneberger, Chimezie Iwuanyanwu, Shubin Zhao, David Soergel, Cho-Jui Hsieh, Irene Cai, Shariq Iqbal, Martin Sundermeyer, Zhe Chen, Elie Bursztein, Chaitanya Malaviya, Fadi Biadsy, Prakash Shroff, Inderjit Dhillon, Tejasi Latkar, Chris Dyer, Hannah Forbes, Massimo Nicosia, Vitaly Nikolaev, Somer Greene, Marin Georgiev, Piding Wang, Nina Martin, Hanie Sedghi, John Zhang, Praseem Banzal, Doug Fritz, Vikram Rao, Xuezhi Wang, Jiageng Zhang, Viorica Patraucean, Dayou Du, Igor Mordatch, Ivan Jurin, Lewis Liu, Ayush Dubey, Abhi Mohan, Janek Nowakowski, Vlad-Doru Ion, Nan Wei, Reiko Tojo, Maria Abi Raad, Drew A. Hudson, Vaishakh Keshava, Shubham Agrawal, Kevin Ramirez, Zhichun Wu, Hoang Nguyen, Ji Liu, Madhavi Sewak, Bryce Petrini, DongHyun Choi, Ivan Philips, Ziyue Wang, Ioana Bica, Ankush Garg, Jarek Wilkiewicz, Priyanka Agrawal, Xiaowei Li, Danhao Guo, Emily Xue, Naseer Shaik, Andrew Leach, Sath MNM Khan, Julia Wiesinger, Sammy Jerome, Abhishek Chakladar, Alek Wenjiao Wang, Tina Ornduff, Folake Abu, Alireza Ghaffarkhah, Marcus Wainwright, Mario Cortes, Frederick Liu, Joshua Maynez, Andreas Terzis, Pouya Samangouei, Riham Mansour, Tomasz Kępa, François-Xavier Aubet, Anton Algymer, Dan Banica, Agoston Weisz, Andras Orban, Alexandre Senges, Ewa Andrejczuk, Mark Geller, Niccolo Dal Santo, Valentin Anklin, Majd Al Merey, Martin Baeuml, Trevor Strohmman, Junwen Bai, Slav Petrov, Yonghui Wu, Demis Hassabis, Koray Kavukcuoglu, Jeff Dean, and Oriol Vinyals.

2024. [Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context](#). *Preprint*, arXiv:2403.05530.
- Rubèn Tito, Dimosthenis Karatzas, and Ernest Valveny. 2023. [Hierarchical multimodal transformers for multi-page docvqa](#). *Preprint*, arXiv:2212.05935.
- Weiyun Wang, Zhe Chen, Wenhai Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Jinguo Zhu, Xizhou Zhu, Lewei Lu, Yu Qiao, and Jifeng Dai. 2025. [Enhancing the reasoning ability of multimodal large language models via mixed preference optimization](#). *Preprint*, arXiv:2411.10442.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. [Self-consistency improves chain of thought reasoning in language models](#). In *The Eleventh International Conference on Learning Representations*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). *Advances in Neural Information Processing Systems*, 35:24824–24837.
- Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh Jing Hua, Zhoujun Cheng, Dongchan Shin, Fangyu Lei, Yitao Liu, Yiheng Xu, Shuyan Zhou, Silvio Savarese, Caiming Xiong, Victor Zhong, and Tao Yu. 2024. [Osworld: Benchmarking multimodal agents for open-ended tasks in real computer environments](#). *Preprint*, arXiv:2404.07972.
- Chenyu Yang, Xuan Dong, Xizhou Zhu, Weijie Su, Jiahao Wang, Hao Tian, Zhe Chen, Wenhai Wang, Lewei Lu, and Jifeng Dai. 2024. [Pvc: Progressive visual token compression for unified image and video processing in large vision-language models](#). *Preprint*, arXiv:2412.09613.
- Michihiro Yasunaga, Armen Aghajanyan, Weijia Shi, Rich James, Jure Leskovec, Percy Liang, Mike Lewis, Luke Zettlemoyer, and Wen tau Yih. 2023. [Retrieval-augmented multimodal language modeling](#). *Preprint*, arXiv:2211.12561.
- Jiabo Ye, Haiyang Xu, Haowei Liu, Anwen Hu, Ming Yan, Qi Qian, Ji Zhang, Fei Huang, and Jingren Zhou. 2024. [mplug-owl3: Towards long image-sequence understanding in multi-modal large language models](#). *Preprint*, arXiv:2408.04840.
- Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Tiantian Fan, Gaohong Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi Zhang, Mofan Zhang, Wang Zhang, Hang Zhu, Jinhua Zhu, Jiase Chen, Jiangjie Chen, Chengyi Wang, Hongli Yu, Weinan Dai, Yuxuan Song, Xiangpeng Wei, Hao Zhou, Jingjing Liu, Wei-Ying Ma, Ya-Qin Zhang, Lin Yan, Mu Qiao, Yonghui Wu, and Mingxuan Wang. 2025. [Dapo: An open-source llm reinforcement learning system at scale](#). *Preprint*, arXiv:2503.14476.
- Yuexiang Zhai, Hao Bai, Zipeng Lin, Jiayi Pan, Shengbang Tong, Yifei Zhou, Alane Suhr, Saining Xie, Yann LeCun, Yi Ma, and Sergey Levine. 2024. [Fine-tuning large vision-language models as decision-making agents via reinforcement learning](#). *Preprint*, arXiv:2405.10292.
- Jiwen Zhang, Jihao Wu, Yihua Teng, Minghui Liao, Nuo Xu, Xiao Xiao, Zhongyu Wei, and Duyu Tang. 2024a. [Android in the zoo: Chain-of-action-thought for gui agents](#). *Preprint*, arXiv:2403.02713.
- Kaichen Zhang, Bo Li, Peiyuan Zhang, Fanyi Pu, Joshua Adrian Cahyono, Kairui Hu, Shuai Liu, Yuanhan Zhang, Jingkang Yang, Chunyuan Li, and Ziwei Liu. 2024b. [Lmms-eval: Reality check on the evaluation of large multimodal models](#). *Preprint*, arXiv:2407.12772.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena](#). *Preprint*, arXiv:2306.05685.

A Prompt for Chain of Scroll

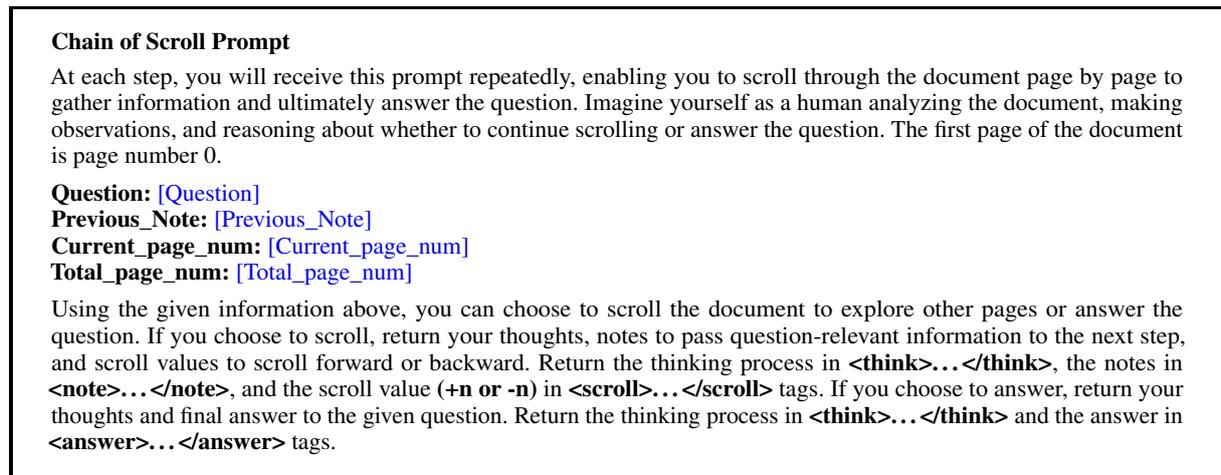


Figure 5: Input prompt of CoS framework. At every step, the blue part is replaced with the input query, accumulated previous notes, current page number, and total page number until the model chooses to return the answer.

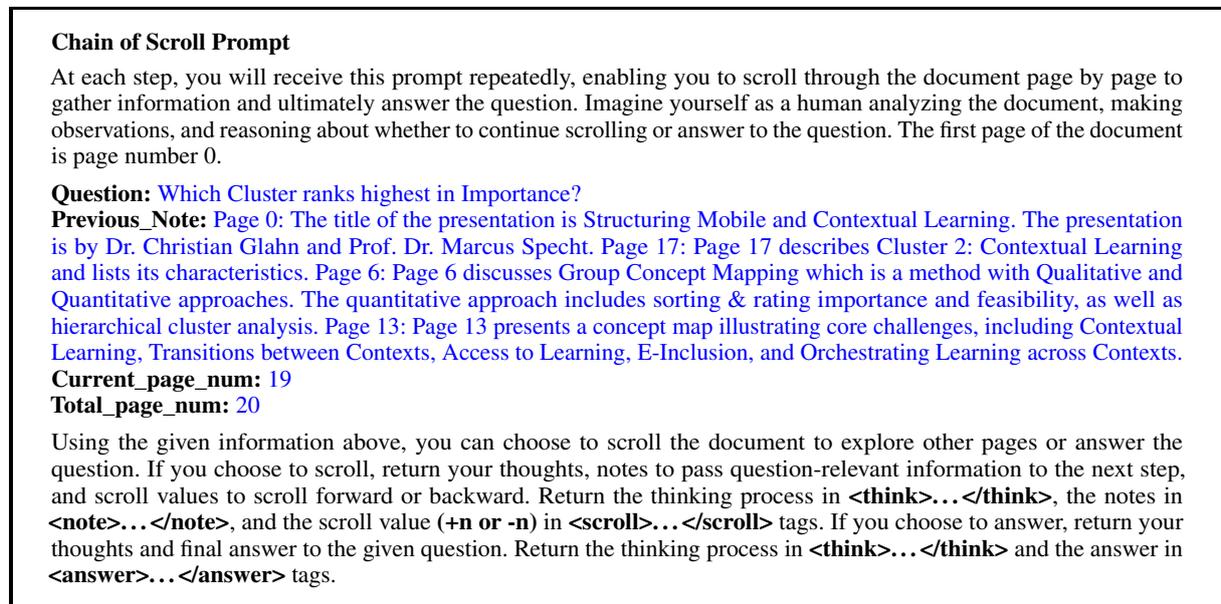


Figure 6: Example shows the input prompt of CoS framework for slidevqa in SCoPE dataset with the dataset id of "00mlearn2011glahn-111020092730-phpapp02_95_28_4".

Figure 5 and 6 illustrates the input prompt template used in the Chain of Scroll (CoS) framework. At each decision step, the model is presented with a structured prompt that includes the user's question, the current page number being viewed, the total number of pages in the document, and any previously accumulated notes. The model is asked to simulate human-like reasoning by either scrolling through additional pages or answering the question based on the information seen so far. It must explicitly output its reasoning process using '`<think>...</think>`', relevant information to carry forward in '`<note>...</note>`', and a scroll decision using '`<scroll>...</scroll>`'. If it opts to answer the question, it returns its final reasoning and conclusion using '`<think>...</think>`' and '`<answer>...</answer>`' tags. This recursive prompting mechanism allows the model to incrementally explore long documents and build up a chain of contextual understanding before answering. Since this prompt contains all query-related information from the conversation history, the CoS framework treats each step as a single turn. Previous conversation history is not provided at each step. This design allows CoS to retain global information while navigating the document, enabling the incorporation of holistic information in the final answer.

B SCoPE dataset generation

B.1 Overall Annotation Process

The annotation generation pipeline for SCoPE dataset creates synthetic human-like document navigation traces for multi-page document question answering. The system employs a four-stage approach: (1) first identifying evidence pages containing answer-relevant information when ground truth is unavailable, (2) randomly selecting intermediate pages to create realistic exploration trajectories, (3) annotating scroll steps with human-like reasoning for each page visited, and (4) generating the final answer step with reasoning. The pipeline leverages large language models to simulate human cognitive processes during document exploration, producing training data that captures both successful information-seeking strategies and natural exploration behaviors. The resulting annotations include page-level observations, accumulated notes across pages, and human-like reasoning chains that justify navigation decisions and answer derivation.

B.2 Prompt for Evidence Page Identification

Evidence Page Identification Prompt

You are given multiple pages of a document, along with a question and an answer that were provided for a query about the document.

Question: [Question]

Answer: [Answer]

Your task is to finding all pages of the document are necessary to answer the query accurately. You need to return the image number which starts from 0. Return your findings as a list of page numbers in the following format. If all of the images seem to be necessary, put all page indices in the list.

However, if there is no required pages in the given images return an empty list.

Each image is a page, even if the image is a crop of a bigger image.

Make sure to look back your choice and there must be the **given answer** in the selected pages.

You can only speak json and put all of your thoughts under the thoughts.

Use this JSON schema:

Return: {'thoughts': str, 'output': int}

Figure 7: Target page identification prompt used to determine which pages contain information necessary to answer a given question. This prompt is only deployed when ground truth answer pages are not available in the dataset.

Figure 7 presents the prompt template for identifying evidence pages within a document. This prompt is employed only when the dataset lacks pre-annotated answer page information. Given a question-answer pair and multiple document pages, the model must analyze each page to determine which ones contain essential information for answering the question. The model outputs its reasoning process and a list of relevant page indices in JSON format, ensuring that the identified pages collectively contain the information needed to derive the given answer.

B.3 Prompt for Scroll Step Generation

Scroll Step Generation Prompt

You are given:

- A single-page image of a document
- A question
- The number of pages to be skipped and scrolled in the next step
- Current page number
- The total number of pages in the document
- Notes containing relevant information from other pages

At each step, you receive this prompt repeatedly, enabling you to scroll through the document page by page to gather information and ultimately answer the question. Imagine yourself as a human analyzing the document, making observations, and reasoning about whether to continue scrolling. Your task is to generate realistic, human-like reasoning for decision-making. Think as if you have the choice to either answer or continue exploring based on your notes and findings, while also determining the appropriate scroll value—though you are not allowed to answer at this step. The first page of the document is page number 0.

Question: [Question]
Previous Note: [Previous_note]
Scroll_value: [Scroll_value]
Current_page_num: [Current_page_num]
Total_page_num: [Total_page_num]

Your job is to:

1. Identify any information on the **current page** that can be useful to answer the question. Do not repeat the previous note and its information. Only return the new information. Also, note a brief summary of the current page. Do not format the note at all. Put a string simply.
2. Write out an in-depth thinking process about how you find this relevant information and reasoning to conclude to scroll by the scroll value (The thoughts should not reveal that it is instructed and the scroll value is provided.. Answer as if you are not given the scroll value. You still need to provide profound reasoning that you need to scroll [Scroll_value]).

Return your response in **JSON format**.

Figure 8: Scroll step generation prompt for creating intermediate navigation annotations. This prompt generates human-like reasoning for each page visited during document exploration.

Figure 8 shows the prompt template for generating intermediate navigation steps. For each page in the exploration trajectory except the final page, this prompt elicits reasoning about the current page’s content and justification for the navigation decision. The model must identify new relevant information on the current page, maintain accumulated notes from previous pages, and generate plausible reasoning for why a human would choose to scroll to the next page. Crucially, the model must generate reasoning as if it independently decided to navigate, even though the scroll direction is predetermined, ensuring the annotations reflect natural exploration patterns.

B.4 Prompt for Answer Step Generation

Answer Step Generation Prompt

You are given:

- A single-page image of a document
- A question
- Current page number
- The total number of pages in the document
- Notes containing relevant information from other pages

At each step, you receive this prompt repeatedly, enabling you to scroll through the document page by page to gather information and ultimately answer the question. Imagine yourself as a human analyzing the document, making observations, and reasoning about whether to continue scrolling. Your task is to generate realistic, human-like reasoning for decision-making. Think as if you have the choice to either answer or continue exploring based on your notes and findings without the answer given, though you must provide the final answer to the question in the end. The first page of the document is page number 0. Any pages including the first and last page may have the enough information to answer the question.

Question: [Question]
Answer: [Answer]
Previous Note: [Previous_Note]
Current_page_num: [Current_page_num]

Try to examine each step in depth and as if it is a realistic thinking process.

Your task is to:

1. Identify any information on the **current page** that can be useful to answer the question.
2. Write out an in-depth thinking process step by step to identify the answer to the question on the current page. There must be the answer in the current page.
3. Explain why you now have enough information to provide the answer
4. You **MUST DERIVE THE FINAL ANSWER** based on the previous note and the current page. *You cannot say that the answer is given.*

Return your response in the exact following **JSON format**.

Figure 9: Answer step generation prompt for creating answer derivation annotations at the terminal page of the exploration trajectory.

Figure 9 illustrates the prompt template for generating the answer step. Applied to the last page in the exploration trajectory, this prompt requires the model to synthesize information from all previously visited pages from accumulated notes with content from the current page to derive the answer. The model must demonstrate explicit reasoning about why sufficient information has been gathered, how the current page contributes to the answer, and provide a step-by-step derivation process. The prompt ensures that the model generates reasoning as if discovering the answer organically, rather than simply restating a provided answer, creating annotations that reflect genuine comprehension and synthesis processes.

C Formal representation of Chain of Scroll

Let $\tau = (s_0, a_0, \dots, s_T, a_T)$ be a trajectory generated by policy π_θ . In the Chain of Scroll framework, each state s_t depends on the input query q , all input images $imgs$, current page index, and accumulated context (notes).

Let \mathcal{V} be the vocabulary of tokens, and \mathcal{V}^* denote the set of all finite sequences over \mathcal{V} .

C.1 State Space

We formalize the state at time t as $s_t = (page_t, notes_t, visited_t)$ where:

- $page_t \in \{0, 1, \dots, N\}$ is the current page index (0-based)
- $notes_t \subseteq \mathcal{V}^*$ is the set of accumulated notes containing extracted information
- $visited_t \in \{0, 1\}^{N+1}$ is a boolean array tracking which pages have been visited

C.2 Context and Transition Function

At each step, the context is assembled from s_t and a_t as:

$$c_t = (q, imgs, page_t, scroll_t, notes_t, visited_t) \quad (6)$$

where $imgs = \{Image_0, \dots, Image_N\}$.

The transition function TRAN_FN processes this context:

$$\text{TRAN_FN}(c_t) \rightarrow (prompt_t, cur_img_t, page_t, visited_t) \quad (7)$$

where:

- $prompt_t \in \mathcal{V}^*$ is the generated prompt containing task instructions and accumulated context
- $cur_img_t = imgs[page_t]$ is the current page image

C.3 Policy

The policy π_θ takes the image and prompt to generate a response:

$$\pi_\theta(cur_img_t, prompt_t) \rightarrow response_t \in \mathcal{V}^* \quad (8)$$

The policy never directly receives the user query q , full image set $imgs$, or raw state components. These are processed by TRAN_FN into the prompt and $prompt_t$.

C.4 Parse Function

The PARSE function extracts structured output from the model's response:

$$\text{PARSE}(response_t) \rightarrow (cur_note_t, scroll_t, answer_t) \quad (9)$$

where $answer_t \in \mathcal{V}^*$.

C.5 Action Space

The action space is implicitly defined by the parse output:

- Scroll action: $(cur_note_t, scroll_t, \emptyset)$ where $scroll_t$ must satisfy:

$$-page_t \leq scroll_t \leq N - page_t \quad (10)$$

- Answer action: $(cur_note_t, scroll_t, answer_t)$ where $answer_t \neq \emptyset$

C.6 State Updates and Termination

C.6.1 Initial State

The initial state is defined as:

$$s_0 = (page_0 = 0, notes_0 = \emptyset, visited_0 = [0]^{N+1}) \quad (11)$$

with $scroll_0 = 0$.

C.6.2 Notes Accumulation

$$notes_{t+1} = notes_t \cup \{cur_note_t\} \quad (12)$$

C.6.3 Page Update

The page index is updated based on the scroll value:

$$page_{t+1} = \min(\max(0, page_t + scroll_t), N) \quad (13)$$

C.6.4 Visit History Update

$$visited_{t+1}[i] = \begin{cases} 1 & \text{if } i = page_{t+1} \\ visited_t[i] & \text{otherwise} \end{cases} \quad (14)$$

C.6.5 Episode Termination and Final Output

An episode terminates when one of the following conditions is met:

- An answer is found: $answer_t \neq \emptyset$
- Maximum steps reached: $t = T_{\max}$, where $T_{\max} = \min(maxSteps, |imgs|)$

The algorithm returns:

$$answer = \begin{cases} answer_t & \text{if } answer_t \neq \emptyset \text{ for some } t \leq T_{\max} \\ \emptyset & \text{otherwise} \end{cases} \quad (15)$$

where the first non-empty answer encountered during the episode is returned.

D Episodic Group Relative Policy Optimization (EGRPO)¹

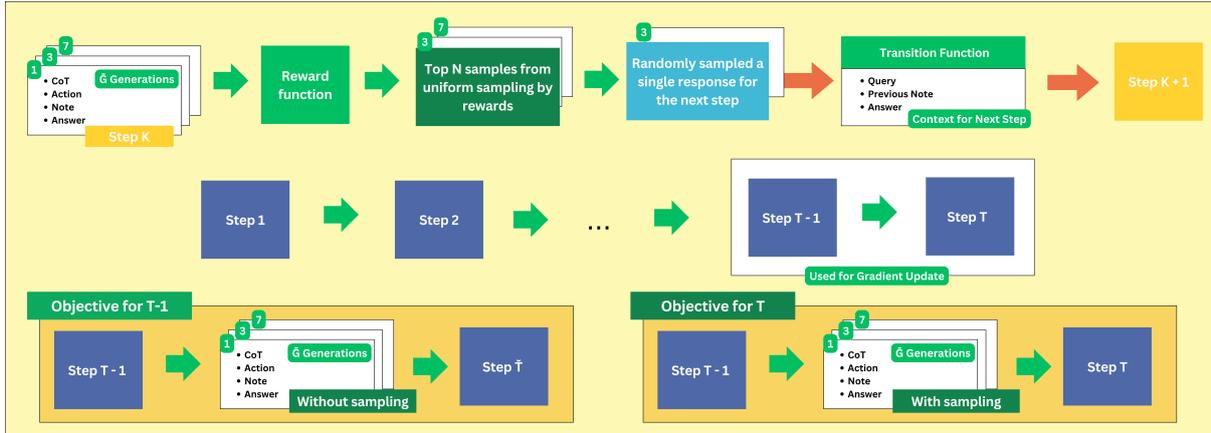


Figure 10: Overview of Episodic Group Relative Policy Optimization (EGRPO). The algorithm generates \tilde{G} candidate actions at each step, applies uniform sampling followed by random selection from top- N to balance exploration-exploitation.

While GRPO (DeepSeek-AI et al., 2025) has proven effective for single-step Chain-of-Thought reasoning, it does not account for multi-step frameworks like Chain of Scroll (CoS) where agents must navigate documents, accumulate information, and generate answers. Two critical problems emerge: (1) computing fine-grained rewards and values for every step is computationally expensive, especially with LLM-as-judge methods (Zheng et al., 2023), and (2) group-based value computation becomes intractable due to exponential growth in candidate trajectories. EGRPO addresses these by focusing gradient updates on penultimate and terminal steps, using projected states \tilde{T} to estimate returns without exhaustive generation.

EGRPO bridges this gap by introducing a two-step objective that back-propagates rewards to the penultimate scroll step, exploiting that this state contains sufficient information to reach answer pages, coupling Random Sampling from Top- N with Uniform Sampling for balanced exploration-exploitation, and removing KL penalties while using $\pi_{\theta_{\text{ref}}}$ and LoRA regularization. As illustrated in Figure 10, the algorithm generates \tilde{G} candidates per step, applies uniform sampling of G candidates followed by random selection from top- N (by rewards), ensuring efficient exploration. EGRPO employs a two-step objective that focuses gradient updates on the penultimate step ($T - 1$) with future return estimation through projected terminal step \tilde{T} —which is generated without sampling strategies from $T - 1$ to estimate the expected reward if the agent continues from that state—and the terminal step (T) with direct reward signals, where \tilde{T} serves as a projection of what the terminal state would be if the agent took one more step from $T - 1$ without exploration noise.

D.1 Overview of EGRPO objectives

The EGRPO objective function combines terminal and penultimate timestep objectives:

$$\mathcal{J}_{\text{EGRPO}}(\theta) = \gamma \mathcal{J}^{(T)}(\theta) + \mathcal{J}^{(T-1)}(\theta) \quad (16)$$

where θ represents the policy parameters, γ is a weighting factor that balances the importance of the terminal step objective, T denotes the terminal (final) timestep, and $T - 1$ is the penultimate (second-to-last, the final scroll step) timestep.

D.2 Group Sampling: Uniform Sampling and Random Sampling from Top- N candidates

Following the inference phase illustrated in Figure 4 and 10, the model initially generates \tilde{G} candidates and samples to G candidates, evaluates them with the composite reward, and then proceeds in two stages:

¹The training code is based on MM-Eureka (Meng et al., 2025).

1. **Uniform Sampling.** G candidates are drawn uniformly from \tilde{G} to form a sorted mini-batch by rewards, preserving a wide range of diverse reward signals from generated samples.
2. **Random Sampling from Top- N candidates.** Among the G candidates, we identify the N highest-reward samples and randomly select the next environment action from this subset.

While a high initial number of generations is essential to increase exploration and discover diverse high-reward trajectories, this approach incurs substantial computational costs and often produces repetitive samples with limited unique solutions. Uniform sampling alone optimizes exploration by decreasing the size of the group while preserving the diversity; top- N selection maximizes exploitation but risks premature convergence to local optima. Our two-stage scheme inherits the best of both approaches while addressing the uniqueness problem: Top- N selection guarantees that each sampled action is within high-reward trajectories and filters out repetitive low-quality samples, while random selection encourages exploration.

D.3 Regularization Strategy

To further optimize the computational cost to train present long horizon of the CoS, we optimize the regularization strategy used in GRPO. To address these challenges, we adopt a streamlined approach inspired by DAPO (Yu et al., 2025). Specifically, we remove the explicit KL divergence penalty term from the objective function. Additionally, to promote exploration and prevent the policy from prematurely converging to the initial behavior, we modify the importance ratio $\rho_{i,k}^{(t)}(\theta)$ by replacing $\pi_{\theta_{\text{old}}}$ with the fixed reference policy $\pi_{\theta_{\text{ref}}}$. This design choice is motivated by our use of LoRA (Hu et al., 2021) adaptation, which provides implicit regularization through its low-rank parameterization. By constraining updates to a low-dimensional subspace, LoRA naturally prevents excessive policy drift while significantly reducing memory requirements and computational burden compared to full model fine-tuning. This approach achieves effective regularization without the computational overhead of maintaining and updating a separate old policy network.

D.4 Formal Derivation of the EGRPO Objective

D.4.1 Underlying Assumptions in EGRPO

In order to define the information content function $I : \mathcal{S} \rightarrow \mathcal{I}$ that maps states to their information content relevant for answer generation, we make the following key assumptions:

Assumption 1: Representativeness of Penultimate State The penultimate scroll step (Step T-1) is a representative state that contains or can access the critical information needed to reach the final page containing the answer. This assumption is valid for both single-hop retrieval and multi-hop reasoning where the last pieces of information needed to answer the user’s question are accessible from the penultimate state:

$$P(\text{correct answer} | s_{T-1}) \approx P(\text{reaching answer page} | s_{T-1}) \cdot P(\text{correct answer} | \text{answer page reached}) \quad (17)$$

Assumption 2: Context-Dependent Valid States. For single-hop reasoning, there exist multiple independent valid penultimate states $s_{T-1} \in \mathcal{S}_{T-1}^*$ and answer states $s_T \in \mathcal{S}_T^*$. For multi-hop reasoning, the set of valid states may be constrained by information dependencies, where $\mathcal{S}_{T-1}^*(n_t) \subseteq \mathcal{S}_{T-1}^*$ depends on the accumulated information in the note buffer n_t . Formally:

$$\mathcal{S}_{T-1}^*(n_t) = \{s_{T-1} \in \mathcal{S} : I(n_t) \supseteq I_{\text{prereq}}\} \quad (18)$$

where I_{prereq} represents prerequisite information needed for the final reasoning step.

Assumption 3: Dynamic Programming with Diverse Training. By optimizing transitions from diverse penultimate to answer states across varied trajectory lengths during training, the model implicitly learns to increase probabilities of productive actions throughout entire trajectories—specifically navigation actions

toward relevant pages, note-taking actions for critical information extraction, and appropriately-timed answer generation. This occurs through two mechanisms:

1. Value propagates backward through the trajectory via the Bellman equation:

$$V^\pi(s_t) = \mathbb{E}_{a_t \sim \pi} [r_t + \gamma V^\pi(s_{t+1})] \quad (19)$$

Improving $V^\pi(s_{T-1})$ creates pressure to improve $V^\pi(s_t)$ for all $t < T - 1$, which increases $\pi(a|s_t)$ for actions a that lead to high-value successor states.

2. During training, the same state can appear at different positions across trajectories of varying lengths. A state that serves as penultimate (s_{T-1}) in a short trajectory may appear much earlier (s_t for $t \ll T$) in a longer trajectory. This teaches the model to recognize and execute critical transitions at any position.

D.4.2 Validity of Assumptions in Multi-Hop Reasoning

The three assumptions naturally extend to multi-hop reasoning where agents must integrate information from multiple sources:

Assumption 1: The penultimate state s_{T-1} in multi-hop reasoning contains all prerequisite information in the note buffer n_{T-1} , with only the final navigation remaining. The representative state property holds with the extended condition.

Assumption 2: Information dependencies constrain valid states based on accumulated facts. For a query requiring facts $\{f_1, f_2, \dots, f_k\}$, valid penultimate states must satisfy:

$$\mathcal{S}_{T-1}^* = \{s : \{f_1, \dots, f_k\} \subseteq n_t(s) \wedge \text{can_reach_answer}(s)\} \quad (20)$$

Assumption 3: Multi-hop reasoning creates naturally diverse trajectory lengths. A state containing facts $\{f_1, f_2\}$ serves as penultimate in a 2-hop query but appears early in a 3-hop query requiring $\{f_1, f_2, f_3\}$. This diversity ensures the model learns to execute productive transitions from any position. Thus, learning to handle final-hop transitions improves navigation and note-taking decisions throughout the entire multi-hop chain.

D.4.3 Terminal Step Projection

The traditional advantage function is:

$$A^\pi(s, a) = Q^\pi(s, a) - V^\pi(s) \quad (21)$$

For the penultimate step in a trajectory, this becomes:

$$A^\pi(s_{T-1}, a_{T-1}) = Q^\pi(s_{T-1}, a_{T-1}) - V^\pi(s_{T-1}) \quad (22)$$

In EGRPO, we make the following estimations: $Q^\pi(s_{T-1}, a_{T-1})$ is the expected return from taking action a_{T-1} at state s_{T-1} :

$$Q^\pi(s_{T-1}, a_{T-1}) = \mathbb{E}_\pi \left[\sum_{t=T-1}^T r^{(t)} \mid s_{T-1}, a_{T-1} \right] \quad (23)$$

For trajectory i , we estimate this as:

$$Q^\pi(s_{T-1}, a_{T-1}) \approx \hat{r}_i^{(T-1)} = \underbrace{r_i^{(T-1)}}_{\text{current reward}} + \underbrace{r_i^{(\hat{T})}}_{\text{future return estimate}} \quad (24)$$

where:

- $r_i^{(T-1)}$ is the immediate reward at step $T - 1$

- $r_i^{(\tilde{T})}$ is the estimated future return obtained through terminal step projection

The estimation $r_i^{(\tilde{T})}$ is critical for converting GRPO to the episodic setting:

1. Unlike traditional RL that requires complete trajectories, EGRPO can compute returns and values using only the current group of samples by leveraging the group-based advantage estimation from GRPO. By projecting incomplete trajectories to terminal states, we obtain return estimates without waiting for full episode completion.
2. Incomplete trajectories receive lower return estimates compared to successfully completed ones. This creates a natural reward gradient where:
 - Successfully completed trajectories receive full terminal rewards $r_i^{(T)}$
 - Incomplete trajectories receive projected rewards $r_i^{(\tilde{T})}$ that are typically lower
 - This difference provides stronger learning signals for actions leading to completion

D.4.4 The EGRPO Objective

Upon the assumptions and derivations, the complete EGRPO objective extends GRPO and optimizes both terminal and penultimate steps jointly:

$$\mathcal{J}_{\text{EGRPO}}(\theta) = \gamma \mathcal{J}^{(T)}(\theta) + \mathcal{J}^{(T-1)}(\theta) \quad (25)$$

Where each component $\mathcal{J}^{(t)}(\theta)$ for $t \in \{T-1, T\}$ is:

$$\mathcal{J}^{(t)}(\theta) = \mathbb{E}_{(q,a) \sim \mathcal{D}} \mathbb{E}_{\{o_i^{(t)}\}_{i=1}^G} \left[\frac{1}{G} \sum_{i=1}^G \frac{1}{|o_i^{(t)}|} \sum_{k=1}^{|o_i^{(t)}|} \min\left(\rho_{i,k}^{(t)} \hat{A}_i^{(t)}, \text{clip}(\rho_{i,k}^{(t)}, 1 - \varepsilon, 1 + \varepsilon) \hat{A}_i^{(t)}\right) \right]$$

with $\{\tilde{o}_j^{(t)}\}_{j=1}^{\tilde{G}} \sim \pi_{\theta_{\text{old}}}(\cdot | q)$ (Total generated responses prior to sampling)

and $\{o_i^{(t)}\}_{i=1}^G \subseteq \{\tilde{o}_j^{(t)}\}_{j=1}^{\tilde{G}}$ (Sampled G outputs by uniform sampling) (26)

The probability ratio is:

$$\rho_{i,k}^{(t)}(\theta) = \frac{\pi_{\theta}(o_{i,k}^{(t)} | q, o_{i,<k}^{(t)})}{\pi_{\theta_{\text{ref}}}(o_{i,k}^{(t)} | q, o_{i,<k}^{(t)})} \quad (27)$$

Where k indexes the token position within each output sequence.

The normalized advantage is:

$$\hat{A}_i^{(t)} = \frac{\hat{r}_i^{(t)} - \mu^{(t)}}{\sigma^{(t)}} \quad (28)$$

As derived in Section D.4.3, rewards are defined as:

$$\hat{r}_i^{(T)} = r_i^{(T)}, \quad \hat{r}_i^{(T-1)} = r_i^{(T-1)} + r_i^{(\tilde{T})} \quad (29)$$

EGRPO introduces several key differences over standard GRPO to handle episodic tasks. The terminal step projection $r_i^{(\tilde{T})}$ enables efficient learning from variable-length trajectories and partial episodes, crucial for CoS where the SFT model suffers with not always reaching terminal states during exploration. By focusing optimization on penultimate-to-terminal transitions across diverse trajectory lengths, the model implicitly learns productive actions throughout entire trajectories, leveraging the dynamic programming principle stated in Assumption 3. The two-stage group sampling combines uniform sampling from \tilde{G} to G candidates to preserve diversity with random selection from top-N candidates to ensure high-quality actions while maintaining exploration. Together, these modifications extend GRPO's single-step optimization framework to episodic settings while maintaining computational efficiency through selective focus on critical state transitions.

D.4.5 Reward Function

Reward func.	Step	Condition	Return
<i>Accuracy reward</i>		Exception	-1
	Scroll	Valid scroll	+2
		Invalid scroll	-2
		Valid scroll after $> \frac{2}{3}$ pages read	$2 \times \frac{\text{pages_read}}{\text{max_page_num}}$
		Scroll when all pages visited	-4
	Answer	Valid Answer	$w \times \text{ANLS Score}$
Answer $\geq 4 \times \text{GT length}$		-1	
<i>Format reward</i>	Answer	Base score	1
		Valid <answer> tag	+4
		Valid <think> tag	+2
	Scroll	Valid <scroll> tag	+2
		Valid <think> / <note> / Scroll Value	+1/ +1/ +2

Table 7: Accuracy and format rewards for EGRPO. Maximum reward w is set to 7 for each.

For EGRPO, we design two reward functions to evaluate both scroll and answer step, as shown in Table 7. For answer step, Accuracy reward utilizes the ANLS metric (Biten et al., 2019) to evaluate the correctness of the final answer provided at the episode’s conclusion. To encourage exploration in the early steps and to penalize reading entire documents, we have set a decaying step accuracy and a stronger penalty for reading entire documents. Format reward monitors the adherence to structural conventions of every step.

D.5 Pseudocode of EGRPO

Algorithm 2 Episodic Group Relative Policy Optimization (EGRPO)

Input:

q ▷ user’s query
 $imgs = \{Image_1, \dots, Image_P\}$ ▷ ordered pages
 $maxSteps$ ▷ maximum episode length
 $\mathcal{R}(\cdot)$ ▷ reward function
 $TRANSITIONFN(\cdot)$ ▷ transition function
 \tilde{G} ▷ original candidate groups
 G ▷ uniformly sampled subset size
 N ▷ top-N for ranking
 γ ▷ terminal step weight
 ε ▷ clipping parameter

Output: \mathcal{L}_{EGPRO} ▷ policy gradient loss

Initialization:

$step \leftarrow 0$ ▷ Episode step counter
 $page \leftarrow [0]^{\times G}$ ▷ Start at first page for all groups
 $scroll \leftarrow [0]^{\times G}$ ▷ Initial scroll values
 $notes \leftarrow [\emptyset]^{\times G}$ ▷ Empty note buffers
 $visited \leftarrow False^{G \times P}$ ▷ Visitation matrix
 $rewards \leftarrow [\emptyset]^{\times G}$
 $trajectory \leftarrow []^G$ ▷ Page trajectories
 $maxSteps \leftarrow \min(P, maxSteps)$

// Initialize with transition function (scroll to first page)

$c \leftarrow (q, imgs, page, scroll, notes, visited)$
 $(prompts, imgInputs, page, visited, trajectory, done, validScroll) \leftarrow TRANSITIONFN(c)$
 $buffer \leftarrow []$ ▷ Initialize empty buffer

// Generate CoS trajectories

while $step < maxSteps$ **and** $\neg done$ **do**
 $\{\tilde{o}_j\}_{j=1}^{\tilde{G}} \sim \pi_{\theta_{old}}(\cdot | imgInputs, prompts)$ ▷ Generate \tilde{G} candidates
 $rewards_{\tilde{G}} \leftarrow \mathcal{R}(\{\tilde{o}_j\}_{j=1}^{\tilde{G}})$
 $\{o_i\}_{i=1}^G \leftarrow ORDEREDUNIFORMSAMPLE(\{\tilde{o}_{indices[i]}\}_{i=1}^G, rewards_{\tilde{G}}, G)$ ▷ Ordered uniform sampling by rewards

// Parse all outputs in the uniformly sampled group, G

for $i \in \{1, \dots, G\}$ **do**
 $(cur_note_i, scroll_i, answer_i) \leftarrow PARSE(o_i)$
 $notes[i] \leftarrow notes[i] \cup \{cur_note_i\}$
 $rewards[i] \leftarrow rewards_{\tilde{G}}[indices[i]]$

end for
 $buffer \leftarrow buffer \cup \{(\{o_i\}_{i=1}^G, rewards, prompts, imgInputs, page, scroll, notes, visited)\}$

// Select best from top-N for next state

$top_N \leftarrow TOPN(rewards, N)$
 $i^* \leftarrow RANDOMCHOICE(top_N)$ ▷ Select the next state index

// Broadcast selected action to all groups

$page \leftarrow [page[i^*]]^{\times G}$ ▷ All groups move to selected page
 $scroll \leftarrow [scroll[i^*]]^{\times G}$ ▷ All use selected scroll
 $notes \leftarrow [notes[i^*]]^{\times G}$ ▷ All use selected notes
 $visited \leftarrow [visited[i^*]]^{\times G}$ ▷ Copy visitation state

// Transition function with batched input

$c \leftarrow (q, imgs, page, scroll, notes, visited)$
 $(prompts, imgInputs, page, visited, trajectory, done, validScroll) \leftarrow TRANSITION_FN(c)$
 $step \leftarrow step + 1$

end while

Algorithm 2 Episodic Group Relative Policy Optimization (EGPRO) (continued)

$T \leftarrow \text{step}$ ▷ Assign terminal step

// Terminal step loss

$(\{o_i^{(T)}\}_{i=1}^G, \text{rewards}^{(T)}, \text{prompts}^{(T)}, \text{imgInputs}^{(T)}, \dots) \leftarrow \text{buffer}[T]$

$\mu^{(T)} \leftarrow \text{mean}(\text{rewards}^{(T)})$

$\sigma^{(T)} \leftarrow \text{std}(\text{rewards}^{(T)})$

for $i \in \{1, \dots, G\}$ **do**

$\hat{A}_i^{(T)} \leftarrow \frac{\text{rewards}_i^{(T)} - \mu^{(T)}}{\sigma^{(T)} + \epsilon}$ ▷ Normalize advantages

$\rho_i^{(T)} \leftarrow \frac{\pi_{\theta}(o_i^{(T)} | \text{imgInputs}^{(T)}, \text{prompts}^{(T)})}{\pi_{\theta_{\text{ref}}}(o_i^{(T)} | \text{imgInputs}^{(T)}, \text{prompts}^{(T)})}$ ▷ Probability ratio

end for

$\mathcal{L}^{(T)} \leftarrow -\frac{1}{G} \sum_{i=1}^G \min(\rho_i^{(T)} \hat{A}_i^{(T)}, \text{clip}(\rho_i^{(T)}, 1 - \epsilon, 1 + \epsilon) \hat{A}_i^{(T)})$

// Terminal step projection from Penultimate step to estimate the future reward

$(\{o_i^{(T-1)}\}_{i=1}^G, \text{rewards}^{(T-1)}, \text{prompts}^{(T-1)}, \text{imgInputs}^{(T-1)}, \dots) \leftarrow \text{buffer}[T - 1]$ ▷ Get unsampled penultimate step

for $i \in \{1, \dots, G\}$ **do**

$o_i^{(\tilde{T})} \sim \pi_{\theta_{\text{old}}}(\cdot | \text{imgInputs}^{(T-1)}, \text{prompts}^{(T-1)})$ ▷ Generate terminal state from each penultimate step

$\text{rewards}_i^{(\tilde{T})} \leftarrow \mathcal{R}(o_i^{(\tilde{T})})$

$\text{rewards}_i^{(T-1)} \leftarrow \text{rewards}_i^{(T-1)} + \text{rewards}_i^{(\tilde{T})}$ ▷ Sum terminal state reward with the penultimate step to model return

end for

$\mu^{(T-1)} \leftarrow \text{mean}(\text{rewards}^{(T-1)})$

$\sigma^{(T-1)} \leftarrow \text{std}(\text{rewards}^{(T-1)})$

for $i \in \{1, \dots, G\}$ **do**

$\hat{A}_i^{(T-1)} \leftarrow \frac{\text{rewards}_i^{(T-1)} - \mu^{(T-1)}}{\sigma^{(T-1)} + \epsilon}$

$\rho_i^{(T-1)} \leftarrow \frac{\pi_{\theta}(o_i^{(T-1)} | \text{imgInputs}^{(T-1)}, \text{prompts}^{(T-1)})}{\pi_{\theta_{\text{ref}}}(o_i^{(T-1)} | \text{imgInputs}^{(T-1)}, \text{prompts}^{(T-1)})}$

end for

$\mathcal{L}^{(T-1)} \leftarrow -\frac{1}{G} \sum_{i=1}^G \min(\rho_i^{(T-1)} \hat{A}_i^{(T-1)}, \text{clip}(\rho_i^{(T-1)}, 1 - \epsilon, 1 + \epsilon) \hat{A}_i^{(T-1)})$

$\mathcal{L}_{\text{EGPRO}} \leftarrow \gamma \cdot \mathcal{L}^{(T)} + \mathcal{L}^{(T-1)}$ ▷ $\gamma > 1$ emphasizes terminal step loss

return $\mathcal{L}_{\text{EGPRO}}$

As shown in Algorithm 2, EGPRO is an extension of GRPO designed for episodic tasks tailored to sequential navigation through documents to answer queries. The algorithm initializes G groups at the first page with empty note buffers and tracks visitation states across P pages. During each episode step (up to maxSteps), EGPRO generates \tilde{G} candidate outputs from the current policy $\pi_{\theta_{\text{old}}}$, evaluates them with reward function \mathcal{R} , and performs ordered uniform sampling to select G diverse candidates. Each candidate output is parsed to extract notes, scroll positions, and potential answers, with rewards stored in a trajectory buffer. The algorithm then randomly selects one action from the top- N highest-rewarded candidates and broadcasts this choice to all groups. The transition function updates the prompts, image inputs, page positions, and visitation matrix based on the selected action. After the training policy returns the answer, EGPRO computes a two-step loss: for the terminal step T , the loss is the conventional GRPO loss for the generated answer; for the penultimate step $T - 1$, it estimates future returns by projecting terminal states from each candidate without sampling and incorporates these projected rewards into the advantage calculation.

E Chain of Scroll trajectory examples of SCoPE VLM

As shown in Figure 11, 12, 13, 14, 15, 16, we compare the Chain of Scroll trajectories of SCoPE VLM SFT and SCoPE VLM EGRPO for the question: "What is the telephone number of H.Stern branch at Sao Paulo?" (MP-docVQA Val-split QuestionID: 56580). SCoPE VLM SFT fails to find the answer, whereas SCoPE VLM EGRPO successfully finds the correct answer, "224-7727".

E.1 SCoPE VLM SFT

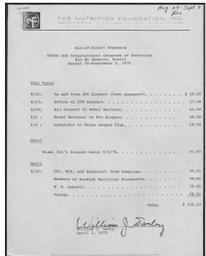
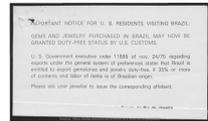
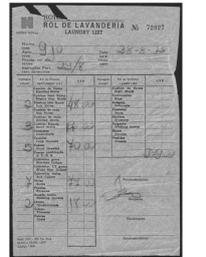
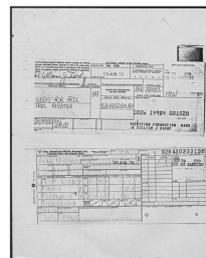
	<p><think>The current page is an expense report from The Nutrition Foundation, Inc. It lists out-of-pocket expenses for the International Congress of Nutrition held in Rio de Janeiro, Brazil. There's no mention of H.Stern or Sao Paulo on this page. Since the question asks for the telephone number of H.Stern branch at Sao Paulo, I need to look for documents that might contain contact information for businesses or organizations located there. Given that this document is related to a conference, it's possible that later pages might contain contact details for attendees or sponsors. Therefore, I should continue scrolling to see if any subsequent pages contain the information I need.</think></p>		<p><think>The current page is an important notice regarding duty-free status for gems and jewelry purchased in Brazil. It mentions U.S. customs and government executive orders related to exports under the general system of preferences. There's no mention of H.Stern or Sao Paulo on this page. Since the question asks for a telephone number associated with H.Stern in Sao Paulo, I need to look for documents that might contain contact information or business listings. Given that the previous page was also irrelevant, I'll go back one page to see if there's any useful information there.</think></p>
<p>Step 0 Page 0</p> <p>Scroll Step</p>	<p><note>This page contains an expense report from The Nutrition Foundation, Inc. for the International Congress of Nutrition held in Rio de Janeiro, Brazil. No mention of H.Stern or Sao Paulo.</note></p> <p><scroll>+12</scroll></p>	<p>Step 1 Page 12</p> <p>Scroll Step</p>	<p><note>This page is an important notice for U.S. residents visiting Brazil regarding duty-free status for gems and jewelry.</note></p> <p><scroll>-1</scroll></p>
	<p><think>The current page is a laundry list from Hotel ROL DE LAVANDERIA. It includes items like shirts, drawers, socks, and dresses, along with their prices. There's no mention of H.Stern or Sao Paulo on this page. Since the question asks for the telephone number of H.Stern branch at Sao Paulo, I need to find a page that contains contact information or a directory listing. I'll scroll back to see if there's any relevant information in the previous pages.</think></p> <p><note>This page is a laundry list from Hotel ROL DE LAVANDERIA, detailing items of clothing and their corresponding costs.</note></p> <p><scroll>-3</scroll></p>		<p><think>The current page appears to be a transportation ticket from Transportes Paranapan S/A. It includes the address, phone number (396-1565), and some codes related to the transportation service. However, it does not contain any information about H.Stern or Sao Paulo. I need to keep looking for the information about H.Stern's branch in Sao Paulo. Since the previous pages didn't have the answer either, I'll skip ahead a few pages to see if there's anything relevant later in the document.</think></p> <p><note>This page is a transportation ticket from Transportes Paranapan S/A with the phone number 396-1565.</note></p> <p><scroll>+4</scroll></p>
<p>Step 2 Page 11</p> <p>Scroll Step</p>	<p><think>The current page shows two airline tickets. One is issued by Pan American World Airways, Inc. and the other is issued by Universal Credit Card Charge Form. Neither of these tickets mentions H.Stern or Sao Paulo. I need to keep looking for the address or contact information for H.Stern's branch in Sao Paulo. Since the previous pages didn't have any useful information either, I'll skip ahead a few pages to see if there's anything related to business addresses or contact details.</think></p> <p><note>This page contains two airline tickets. One is issued by Pan American World Airways, Inc. and the other is issued by Universal Credit Card Charge Form. No mention of H.Stern or Sao Paulo.</note></p> <p><scroll>+4</scroll></p>	<p>Step 3 Page 8</p> <p>Scroll Step</p>	<p><think>The current page appears to be a ticket issued by ARSA Aeroportos do Rio de Janeiro S.A. It includes information such as the airport name, flight details, and a number. However, it does not contain any information related to H.Stern or Sao Paulo. I need to continue searching for the telephone number of H.Stern's branch in Sao Paulo. Since the previous pages didn't have the answer either, I'll skip ahead a few pages to see if there's any relevant information later in the document.</think></p> <p><note>This page is a ticket issued by ARSA Aeroportos do Rio de Janeiro S.A. for a flight from Rio de Janeiro to Sao Paulo. It doesn't mention H.Stern.</note></p> <p><scroll>+4</scroll></p>
	<p>Step 4 Page 3</p> <p>Scroll Step</p>	 <p>Step 5 Page 7</p> <p>Scroll Step</p>	<p><note>This page is a ticket issued by ARSA Aeroportos do Rio de Janeiro S.A. for a flight from Rio de Janeiro to Sao Paulo. It doesn't mention H.Stern.</note></p> <p><scroll>+4</scroll></p>

Figure 11: Model responses from Step 0 to Step 5 of SCoPE VLM SFT to the question: What is the telephone number of H.Stern branch at Sao Paulo?



<think>The current page is a stock control form from H. Stern. It includes the name 'H. Stern' at the top left corner. However, it does not contain any information about a telephone number or Sao Paulo. I need to check if there's a connection between H. Stern and Sao Paulo on previous pages. Since the previous pages did not have the answer, I will go back one page.</think>

<note>This page is a stock control form from H. Stern, but it does not contain a telephone number or mention Sao Paulo.</note>

<scroll>-1</scroll>

Step 6 Page 6
Scroll Step



<think>The current page is a hotel bill from Hotel Nacional in Rio de Janeiro, Brazil. It includes charges for room, meals, and other services. I don't see any mention of H. Stern or Sao Paulo on this page. Since the previous note mentioned that page 6 had information about H. Stern, I should go back to that page to check if there's a phone number listed there.</think>

<note>This page is a hotel bill from Hotel Nacional in Rio de Janeiro, Brazil, listing charges for room, meals, and other services. No mention of H. Stern or Sao Paulo.</note>

<scroll>-4</scroll>

Step 7 Page 5
Scroll Step

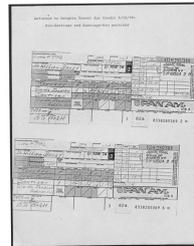


<think>The current page is a hotel bill from Ouro Verde Hotel in Rio de Janeiro, Brazil. It includes the name William J Darby and the address of the establishment. I need to find the telephone number of H. Stern's branch in Sao Paulo. Since this page is related to Rio de Janeiro, it is unlikely to contain the answer. I have already checked several pages including the ones that mentioned H. Stern, but none of them contained the telephone number. I will skip ahead to see if there are any pages that might contain contact information for H. Stern or Sao Paulo.</think>

<note>This page is a hotel bill from Ouro Verde Hotel in Rio de Janeiro, Brazil, for William J Darby. It does not contain the telephone number of H. Stern or mention Sao Paulo.</note>

<scroll>+14</scroll>

Step 8 Page 1
Scroll Step



<think>The current page shows two airline tickets issued by PAN AM. The tickets are for Dr. William Darby and include flights from Rio de Janeiro to Santiago and back to Rio de Janeiro. There's no mention of H. Stern or Sao Paulo on this page. I need to go back to find the information about H. Stern's telephone number in Sao Paulo. Based on my previous notes, page 7 had some information about H. Stern, so I should scroll back to that page.</think>

<note>This page contains two airline tickets issued by PAN AM for Dr. William Darby, one for Rio de Janeiro to Santiago and another for Santiago to Rio de Janeiro. No mention of H. Stern or Sao Paulo.</note>

<scroll>-14</scroll>

Step 9 Page 15
Scroll Step

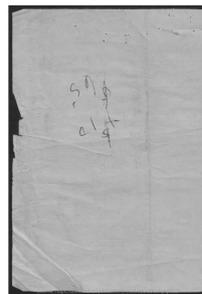


<think>The current page is a hotel bill from Miami International Airport Hotel for William J Darby. It includes charges for room and distance. I need to find the telephone number of H. Stern's branch in Sao Paulo. Since this page is irrelevant to the question, I should continue searching. I'll skip ahead a few pages to see if there's any relevant information.</think>

<note>This page is a hotel bill from Miami International Airport Hotel for William J Darby. It does not contain the telephone number of H. Stern or mention Sao Paulo.</note>

<scroll>+4</scroll>

Step 10 Page 2
Scroll Step



<think>The current page appears to be blank with some handwriting on it. I don't see any information related to H. Stern or Sao Paulo. Based on my previous notes, I have already checked several pages including pages that mentioned H. Stern, but did not contain the telephone number. I need to go back further in the document to check if there's any other mention of H. Stern or Sao Paulo. Therefore, I will scroll back.</think>

<note>This page is blank with some handwriting.</note>

<scroll>-4</scroll>

Step 11 Page 10
Scroll Step

Figure 12: Model responses from Step 6 to Step 11 of SCoPE VLM SFT to the question: What is the telephone number of H. Stern branch at Sao Paulo?



Step 12 Page 17

Scroll Step

<think>The current page shows two airline tickets issued by PAN AM for Dr. William Darby. The destinations include Rio de Janeiro, Santiago, and Nashville. There's no mention of H.Stern or Sao Paulo on this page. I need to go back to find the information about H.Stern's telephone number in Sao Paulo. Based on my previous notes, page 7 had a ticket issued by ARSA Aeroportos do Rio de Janeiro S.A. for a flight from Rio de Janeiro to Sao Paulo. I should scroll back to that page.</think>

<note>This page contains two airline tickets issued by PAN AM for Dr. William Darby, one for Rio de Janeiro to Santiago and another for Santiago to Rio de Janeiro. No mention of H.Stern or Sao Paulo.</note>

<scroll>-14</scroll>



Step 13 Page 14

Scroll Step

<think>The current page is a check from Heights Travel Service Inc. to The Nutrition Foundation, Inc. for \$194.10. The check is dated November 15, 1978. There's no mention of H.Stern or Sao Paulo on this page. I need to go back and check previous pages to see if I missed anything related to H.Stern or Sao Paulo. I'll scroll back to page 13.</think>

<note>This page is a check from Heights Travel Service Inc. to The Nutrition Foundation, Inc. for \$194.10, dated Nov. 15, 1978. No mention of H.Stern or Sao Paulo.</note>

<scroll>-1</scroll>



Step 14 Page 13

Scroll Step

<think>The current page (page 13) is an expense report from H.Stern. I see the name 'H.Stern' at the top left corner of the page. Also, I see the address 'Sao Paulo: Praça da Republica, 242'. This confirms that H.Stern has a branch in Sao Paulo. However, there is no telephone number on this page. I need to go back to find the telephone number of H.Stern's branch in Sao Paulo.</think>

<note>This page is an expense report from H.Stern, including the address of their Sao Paulo branch. It also includes the name of the person who received the report, Dr. William J. Darby. There is no telephone number on this page.</note>

<scroll>-1</scroll>



Step 15 Page 4

Scroll Step

<think>The current page is a hotel bill from Hotel Nacional in Rio de Janeiro, Brazil, for William J Darby. It lists charges for room, meals, and other services. I need to find the telephone number of H.Stern's branch in Sao Paulo. I have already checked several pages, including the Sao Paulo address on page 13, but there was no phone number. Since this page is also a hotel bill, it's unlikely to contain the phone number I'm looking for. I'll skip ahead a few pages to see if I can find something related to H.Stern or Sao Paulo.</think>

<note>This page is a hotel bill from Hotel Nacional in Rio de Janeiro, Brazil, for William J Darby. It does not contain the telephone number of H.Stern or mention Sao Paulo.</note>

<scroll>+4</scroll>



Step 16 Page 18

Scroll Step

<think>The current page is an itinerary for Dr. William Darby's trip. It lists dates, locations (Rio de Janeiro, Miami, Nashville), airlines, and times. I need to find the telephone number of H.Stern's branch in Sao Paulo. Since this page is about Dr. Darby's itinerary and doesn't mention H.Stern or Sao Paulo, it's unlikely to contain the answer. I have already checked several pages that mentioned H.Stern, but none contained a phone number. I should go back to the page where H.Stern was mentioned to see if there is a phone number listed there.</think>

<note>This page is an itinerary for Dr. William Darby's trip, listing dates, locations, airlines, and times. It does not contain the telephone number of H.Stern or mention Sao Paulo.</note>

<scroll>-1</scroll>



Step 17 Page 9

Scroll Step

<think>The current page is a Citibank traveler's check form. It shows that the currency is Brazilian Real (BRL), the amount is 4000, and the rate is 0.52. The form is addressed to someone at Citibank, N.A., JFK International Airport, Jamaica, New York. I need to find the telephone number of H.Stern's branch in Sao Paulo. I have already checked several pages, including pages related to travel expenses, hotel bills, and airline tickets. I also found a page with the address of H.Stern's Sao Paulo branch, but no phone number. Since this page is a bank check, it is unlikely to contain the phone number I'm looking for. I should go back to the page where the address of H.Stern was mentioned to see if there is a phone number listed there.</think>

<note>This page is a Citibank traveler's check form for Brazilian Real, issued at JFK International Airport, New York. It does not contain the telephone number of H.Stern or mention Sao Paulo.</note>

<scroll>-1</scroll>

Figure 13: Model responses from Step 12 to Step 17 of SCoPE VLM SFT to the question: What is the telephone number of H.Stern branch at Sao Paulo?



Step 18 Page 16

Scroll Step

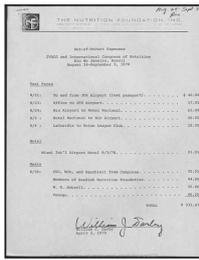
<think>The current page is an invoice from Heights Travel Service Inc. to The Nutrition Foundation, Inc. for transportation for Dr. W. Darby. It lists the dates of travel and the locations. I need to find the telephone number of H.Stern's branch in Sao Paulo. Since this page is an invoice related to Dr. Darby's travel, it's unlikely to contain the information I'm looking for. I have already checked several pages that were related to Dr. Darby's travel expenses and didn't find the answer. I should go back to the page where H.Stern was mentioned to see if there is any contact information listed there.</think>

<note>This page is an invoice from Heights Travel Service Inc. to The Nutrition Foundation, Inc. for transportation for Dr. W. Darby. It does not contain the telephone number of H.Stern or mention Sao Paulo.</note>

<scroll>-1</scroll>

Figure 14: Model responses at Step 18 of SCoPE VLM SFT to the question: What is the telephone number of H.Stern branch at Sao Paulo? Ultimately, the SFT model returns the scroll action, failing to return the correct answer.

E.2 SCoPE VLM EGRPO



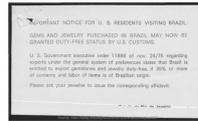
Step 0 Page 0

Scroll Step

<think>The current page is an expense report from The Nutrition Foundation, Inc. It lists out-of-pocket expenses for the International Congress of Nutrition held in Rio de Janeiro, Brazil. There's no mention of H.Stern or Sao Paulo on this page. Since the question asks for the telephone number of H.Stern branch at Sao Paulo, I need to look for documents related to H.Stern or Sao Paulo. I will skip ahead to see if there are any pages that contain the information I need.</think>

<note>This page contains an expense report from The Nutrition Foundation, Inc. for the International Congress of Nutrition held in Rio de Janeiro, Brazil. It includes taxi fares, hotel costs, and meals expenses.</note>

<scroll>+12</scroll>



Step 1 Page 12

Scroll Step

<think>The current page is an important notice regarding duty-free status for gems and jewelry purchased in Brazil. It mentions U.S. customs and government executive orders related to exports. This page doesn't seem to contain any information about the H.Stern branch or its telephone number in Sao Paulo. I need to go back to find the information.</think>

<note>This page is an important notice for U.S. residents visiting Brazil regarding duty-free status for gems and jewelry.</note>

<scroll>-1</scroll>



Step 2 Page 11

Scroll Step

<think>The current page is a laundry list from Hotel ROL DE LAVANDERIA. It lists items being washed and their corresponding costs. There's no mention of H.Stern or any phone numbers on this page. Since the question asks for the telephone number of H.Stern branch at Sao Paulo, I need to find a page that contains contact information for H.Stern. I'll go back a few pages to see if I missed anything related to business contacts or addresses.</think>

<note>This page is a laundry list from Hotel ROL DE LAVANDERIA, listing items of clothing and their costs.</note>

<scroll>-3</scroll>



Step 3 Page 8

Scroll Step

<think>The current page appears to be a transportation ticket from Transportes Paranapanuan S/A. It includes the address, phone number (396-1565), and fare information. However, it doesn't mention anything about H.Stern or Sao Paulo. I need to find the telephone number of H.Stern branch at Sao Paulo. Since this page does not contain any relevant information, I should go back to previous pages to see if I missed anything.</think>

<note>This page is a transportation ticket from Transportes Paranapanuan S/A, including the address, phone number (396-1565), and fare information.</note>

<scroll>-7</scroll>



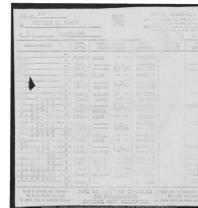
Step 4 Page 1

Scroll Step

<think>The current page is a hotel bill from Ouro Verde Hotel in Rio de Janeiro, Brazil. It includes charges for room, meals, and other services. I need to find the telephone number of H.Stern branch at Sao Paulo. Since this page doesn't contain any information related to H.Stern or Sao Paulo, I should continue searching. I'll skip ahead a few pages to see if there's any relevant information.</think>

<note>This page is a hotel bill from Ouro Verde Hotel in Rio de Janeiro, Brazil, dated August 7, 1978.</note>

<scroll>+4</scroll>



Step 5 Page 5

Scroll Step

<think>The current page is a hotel bill from Hotel Nacional in Rio de Janeiro, Brazil. It lists charges for various services like room, restaurant, and bar. I need to find the telephone number of H.Stern branch at Sao Paulo. Since this page is related to hotel expenses in Rio de Janeiro, it's unlikely to contain the answer. I should check earlier pages in case there was any mention of H.Stern or Sao Paulo. Therefore, I will scroll back one page.</think>

<note>This page is a hotel bill from Hotel Nacional in Rio de Janeiro, Brazil, dated August 7, 1978, with charges for room, restaurant, and bar.</note>

<scroll>-1</scroll>

Figure 15: Model responses from Step 0 to Step 5 of SCoPE VLM EGRPO to the question: What is the telephone number of H.Stern branch at Sao Paulo?



Step 6 Page 4

Scroll Step

<think>The current page is a hotel bill from Hotel Nacional in Rio de Janeiro, Brazil. I need to find the telephone number of H.Stern branch at Sao Paulo. This page does not contain any information about H.Stern or Sao Paulo. I have already checked several pages, including hotel bills and transportation tickets, but none of them contained the information I need. I will continue scrolling to see if there's any mention of H.Stern or Sao Paulo on subsequent pages. Since I haven't found anything yet, I'll scroll further ahead to increase my chances of finding the answer.</think>

<note>This page is a hotel bill from Hotel Nacional in Rio de Janeiro, Brazil, dated August 7, 1978, with charges for room, restaurant, and bar.</note>

<scroll>+12</scroll>



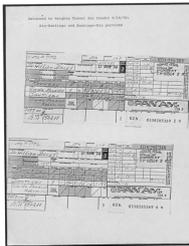
Step 7 Page 16

Scroll Step

<think>The current page is an invoice from Heights Travel Service Inc. to The Nutrition Foundation, Inc. for transportation for Dr. W. Darby. It includes the date, location, and payment details. I need to find the telephone number of H.Stern branch at Sao Paulo. This page does not contain any information related to H.Stern or Sao Paulo. I should go back to previous pages where I might have missed something.</think>

<note>This page is an invoice from Heights Travel Service Inc. to The Nutrition Foundation, Inc. for transportation for Dr. W. Darby.</note>

<scroll>-1</scroll>



Step 8 Page 15

Scroll Step

<think>The current page shows two passenger tickets issued by Pan Am for Dr. William Darby. The tickets indicate flights from Rio de Janeiro to Santiago and back to Rio. There's no mention of H.Stern or Sao Paulo on this page. I need to find the telephone number of H.Stern branch at Sao Paulo. Since this page doesn't seem to have any relevant information, I should go back to previous pages where I might have missed something related to travel or contact information. I'll scroll back to page 13.</think>

<note>This page contains two passenger tickets from Rio de Janeiro to Santiago and back to Rio, issued by Pan Am for Dr. William Darby.</note>

<scroll>-2</scroll>



Step 9 Page 13

Answer Step

<think>The question asks for the telephone number of H.Stern's branch in Sao Paulo. I am currently on page 13. Upon examining page 13, I see the text 'H.Stern' prominently displayed at the top left corner. Directly below it, there are two lines indicating locations and phone numbers: 'RIO DE JANEIRO: AVENIDA RIO BRANCO, 173 - TEL: 224-7727' and 'SAO PAULO: PRAÇA DA REPÚBLICA, 242 - TEL: 258-1222'. The second line clearly states 'SAO PAULO' followed by a city name and a phone number. Therefore, the telephone number for H.Stern's branch in Sao Paulo is explicitly provided on this page. I have found the answer directly on the current page and do not need to look at any further pages.</think><answer>224-7727</answer>

Figure 16: Model responses from Step 6 to Step 9 of SCoPE VLM EGRPO to the question: What is the telephone number of H.Stern branch in Sao Paulo? In Step 9, SCoPE VLM EGRPO correctly answers the question, while the SFT model fails.

F Responses of SCoPE VLM SFT When No Answer Exists in the Context

SCoPE-VLM is trained to answer "The answer cannot be found." when it concludes that the given document does not have the relevant information to answer the question.

Total pages #:27, Question: What do the red squares denote in Fig 1?



Step 0 Page 0

Scroll Step

<think>The question asks about what the red squares denote in Figure 1. This page introduces the Chameleon model and mentions Figures 2-4. It also mentions Figure 1, but it doesn't describe what the red squares represent. I need to find Figure 1 to answer the question. Since the figures are likely presented later in the paper, I should scroll further down to find them.</think>

<note>This page introduces the Chameleon model and mentions Figure 1, but does not explain what the red squares denote.</note>

<scroll>+13</scroll>



Step 1 Page 13

Scroll Step



Step 2 Page 12

Scroll Step

<think>The current page discusses relative evaluation and inter-annotator agreement. It presents Figure 10 which shows the inter-annotator agreement on questions in the absolute evaluation. However, it doesn't mention anything about red squares in Figure 1. I need to go back to find the figure that contains red squares.</think>

<note>This page discusses relative evaluation and mentions Figure 9, which compares Chameleon to other models. It does not mention red squares.</note>

<scroll>-1</scroll>



Step 21 Page 8

Scroll Step



Step 22 Page 26

Scroll Step

<think>The current page (page 26) presents two tables, Table 13 and Table 14, displaying win rates for the Chameleon model against Gemini and GPT-4V, respectively. These tables provide quantitative comparisons across different categories and modalities. However, neither table nor the text directly addresses the meaning of red squares in Figure 1. I need to go back to find the figure and the description of the red squares.</think>

<note>This page contains Table 13 and Table 14, which show complete win rates for Chameleon vs Gemini and Chameleon vs GPT-4V, respectively. There is no mention of Figure 1 or red squares.</note>

<scroll>-1</scroll>



Step 23 Page 9

Answer Step

<think>The current page (page 8) presents Table 3, which provides statistics on the supervised fine-tuning dataset. It lists categories such as Text, Code, Visual Chat, Image Generation, Interleaved Generation, and Safety, along with their respective numbers of samples, tokens, and images. While this table gives an overview of the dataset composition, it doesn't directly address the meaning of red squares in Figure 1. I need to find the page that actually contains Figure 1 to understand what the red squares represent. Based on previous notes, Figure 1 was mentioned on page 1. Therefore, I should scroll back to page 1 to examine Figure 1 more closely.</think>

<note>This page contains Table 3, which shows the supervised fine-tuning dataset statistics. It also mentions Figure 7, but does not explain what the red squares denote.</note>

<scroll>-7</scroll>

<think>The current page (page 9) displays Figure 7, which presents example alignment data categorized into text, code, image generation, visual chat, and interleaved text/image generation. I carefully examined the figure and its descriptions, focusing on identifying any elements that might be referred to as "red squares." However, after reviewing the text snippets associated with each category, I did not find any explicit mention or visual representation of red squares within these examples. The previous notes indicate that Figure 1, which is mentioned in the question, is located on page 1 and depicts the Chameleon model architecture. Since Figure 1 is described as showing the Chameleon model architecture and the question asks about red squares in Figure 1, and I have not found any information about red squares in Figure 1 or any other figures on the current page, I conclude that the question cannot be answered based on the information available on this page and the accumulated notes. Therefore, the answer to the question "What do the red squares denote in Fig 1?" is not present in the provided document extract.</think>

<answer>The answer cannot be found.</answer>

Figure 17: SCoPE VLM SFT's answer to MMLongBench Valid 773 Question: What do the red squares denote in Fig 1?

G Details of training hyperparameters

G.1 SCoPE VLM SFT

Hyperparameter	Value
model	Qwen2.5 VL 3B Instruct
learning rate	2e-5
batch size	8
per device train batch size	2
gradient accumulation steps	4
max length	8192
max pixels	1003520
warmup ratio	0.03
num train epochs	1
optimizer	AdamW
data type	bf16
gradient checkpointing	enabled
deepspeed	zero3

Table 8: Training Hyperparameters for SFT of SCoPE VLM SFT

During the SFT stage, training is performed on 8*A100 GPU for one epoch. The full SCoPE dataset is fine-tuned with the max pixel of 1,003,520 pixels, 1280 tokens per image. Further EGRPO training hyperparameters are shown in Table 8.

G.2 SCoPE VLM EGRPO

Hyperparameter	Value
General	
model	SCOPE VLM 3B SFT
learning rate	1e-4
batch size	1
per device train batch size	1
gradient accumulation steps	1
max prompt length	8096
max completion length	4096
max pixels	1003520
warmup ratio	0.01
num train epochs	1 (early terminated at 2500 step)
num iterations	1
optimizer	AdamW
data type	bf16
gradient checkpointing	enabled
lr scheduler	cosine with min lr
min learning rate	1e-9
LoRA	
LoRA rank	128
LoRA alpha	16
LoRA dropout	0.05
EGRPO	
reward functions	accuracy, format
num generations	8
num samples per group	4
max steps	24
max window length	2
β (KL penalty)	0
ϵ (clipping)	0.2
γ (weighing coefficient)	3
temperature	0.9
top p	0.9
top k	50

Table 9: Training hyperparameters for EGRPO of SCoPE VLM EGRPO

For RL stage, we deploy EGRPO to enhance CoS performance through improved reasoning and memorization. We train our model on two H100 GPUs with LoRA configuration of $r=128$, $\alpha=16$, $\text{dropout}=0.05$, and with the same max pixels of SFT stage. For training dataset, we set up training with SlideVQA and 50% of MP-DocVQA but early terminated at 2500 steps with the maximum CoS step of 24. For ablation, the training step is further reduced to 1000 step. Also, we restrict the page visit to once in the transition function to limit extensive exploration for training. Further EGRPO training hyperparameters are shown in Table 9.

G.3 GRPO for ablation study

Hyperparameter	Value
General	
model	SCOPE VLM 3B SFT
learning rate	1e-4
batch size	1
per device train batch size	1
gradient accumulation steps	1
max prompt length	8096
max completion length	4096
max pixels	1003520
warmup ratio	0.01
num train epochs	1 (early terminated at 1000 step)
num iterations	1
optimizer	AdamW
data type	bf16
gradient checkpointing	enabled
lr scheduler	cosine with min lr
min learning rate	1e-9
LoRA	
LoRA rank	128
LoRA alpha	16
LoRA dropout	0.05
EGRPO	
reward functions	accuracy, format
num generations	4
max steps	24
max window length	1 or 2
β (KL penalty)	0
ϵ (clipping)	0.2
γ (weighing coefficient)	3
temperature	0.9
top p	0.9
top k	50

Table 10: Training hyperparameters for GRPO

For the ablation study, we train SCoPE VLM SFT with GRPO objectives. We train the model in a similar manner as SCoPE VLM EGRPO for 1,000 steps with a maximum CoS step of 24. Further GRPO training hyperparameters are shown in Table 10.

H Details of inference hyperparameters

H.1 Multi page document question answering experiments: Table 2, 3, 5, 6

Hyperparameter	Value
temperature	0
top_p	1.0(default)
num_beams	1
max_new_tokens	1280
Table 2, 5, 6 max_visit_count	2
Table 3 max_visit_count	1

Table 11: Hyperparameters for multi page document question answering experiments: Table 2, 3, 5, 6

H.2 GUI control experiments: Table 4

Hyperparameter	Value
temperature	0.1
top_p	0.9
do_sample	True
max_new_tokens	1024
repetition_penalty	1.1
seed	2020

Table 12: Hyperparameters for GUI control experiments

I Details of the Evaluation Setup in Section 4 and Examples

I.1 Experimental Setup for Table 2 and 6

To ensure fair comparison under identical hardware constraints, we standardized VRAM usage across all evaluated models. In this experiment, we applied a fixed image token budget of 2,560 tokens per inference step (corresponding to a maximum of 2,007,040 pixels). The two inference methods in Table 2 implement VRAM standardization as follows:

Chain of Scroll (CoS): Processes one image per inference step through sequential inference to reach the final conclusion. Therefore, each image in CoS inference receives the full allocation of 2,560 tokens.

Multi-Image (MI): Processes all images simultaneously in a single inference pass. The 2,560 token budget is distributed across all provided images. Consequently, the total token count for the entire image set cannot exceed this limit, resulting in lower per-image token allocation compared to CoS.

In addition, for all methods in this table, models were allowed to visit each image up to twice.

I.2 Experimental Setup for Table 3

In contrast to the setup in Table 2, we did not enforce uniform VRAM usage in Table 3. Instead, our primary objective was to evaluate the trade-off between model performance and VRAM consumption. To achieve this, we imposed a limit of 1,280 tokens per individual image. In addition, the maximum visit is restricted to once for efficiency and the LoRA for SCoPE VLM EGRPO has been merged in the inference time.

For multi-image models, each image within a single concurrent input can be allocated up to 1,280 tokens. For CoS models, input images are similarly limited to a maximum of 1,280 tokens. This setup was designed to analyze how efficiently each model utilizes a fixed per-image token budget and its impact on VRAM usage and task performance.

I.3 Experimental Setup for Table 4

This experiment evaluates model performance on practical GUI navigation tasks using the Android in the Zoo (AitZ) benchmark. The evaluation protocol follows standard AitZ procedures, measuring task

Hyperparameter	Value
learning rate	1e-5
batch size	8
per device train batch size	2
gradient accumulation steps	4
max length	8192
max pixels	1003520
warmup ratio	0.03
num train epochs	1
optimizer	AdamW
data type	bf16
gradient checkpointing	enabled
deepspeed	zero3

Table 13: Hyperparameters for SFT Training of SCoPE VLM and Qwen 2.5 VL for Table 4 in Section 4

completion capability by comparing model-generated actions against ground truth for achieving target goals on given UI screens.

SCoPE VLM SFT, SCoPE VLM EGRPO, and Qwen2.5-3B-VL are fine-tuned with the hyperparameters shown in the Table 13. Also, the max image token is set to 1,280 tokens per image. The LoRA for SCoPE VLM EGRPO has been merged in the inference time as well.

I.4 Experimental Setup for Table 5: Performance Comparison of Serial, Random, and CoS Methods

The ablation study is designed to isolate the effects of different information-seeking strategies. We standardized input conditions by allocating 2,560 tokens per image and limiting the maximum page visits to 2.

For this evaluation, three different inference strategies were compared:

- **Serial:** Pages are presented to the model sequentially in natural order, starting from page 0.
- **Random:** After initial presentation of page 0, subsequent pages are provided in random order.
- **CoS:** Following initial input of page 0, the model autonomously determines which page to explore next based on its reasoning.

J Detailed Examples of Per-Image Token Limits

To manipulate the maximum image token, we set the maximum pixel limit which is used to resize images to limit the number of image tokens. The image processor for this work follows the default processor provided by the baseline Qwen 2.5 VL series (Bai et al., 2025).

When an image is input, we first calculate its initial pixel count using the original height and width. If the calculated pixels exceed the maximum pixel limit, we compute a scaling factor β as follows:

$$\beta = \sqrt{\frac{\text{height} \times \text{width}}{\text{max_pixels}}}$$

We then divide both height and width by β , round down to the nearest multiple of 28 (the patch size), and multiply by 28 to obtain the new dimensions. The product of new height and new width becomes the final pixel count.

To illustrate this process, we provide examples for three cases under both experimental setups:

1. Single image input with dimensions 5120×2880
2. Ten images input, each with dimensions 1980×1080
3. Ten images input, each with dimensions 720×144

J.1 Pixel Resizing for Table 2, 5, and 6: limiting image token usage per single inference step

With $\text{max_pixel} = 2,007,040$:

Case 1: Single 5120×2880 image (Multi-image and CoS apply the same method for single images)

$$\text{Initial Pixels} = 5120 \times 2880 = 14,745,600$$

$$\beta = \sqrt{\frac{14,745,600}{2,007,040}} \approx 2.71$$

$$\text{New Height} = \left\lfloor \frac{2880/2.71}{28} \right\rfloor \times 28 = 1036$$

$$\text{New Width} = \left\lfloor \frac{5120/2.71}{28} \right\rfloor \times 28 = 1876$$

$$\text{Final Pixels} = 1036 \times 1876 = 1,943,536$$

$$\text{Final Tokens} = \frac{1,943,536}{784} = 2479$$

Case 2: Ten 1980×1080 images

For Multi-Image inference, the total pixels across all 10 images cannot exceed 2,007,040. Thus, the per-image pixel limit is 200,704:

$$\text{Initial Pixels} = 1980 \times 1080 = 2,138,400$$

$$\text{Max Pixels per Image} = \frac{2,007,040}{10} = 200,704$$

$$\beta = \sqrt{\frac{2,138,400}{200,704}} \approx 3.26$$

$$\text{New Height} = \left\lfloor \frac{1080/3.26}{28} \right\rfloor \times 28 = 308$$

$$\text{New Width} = \left\lfloor \frac{1980/3.26}{28} \right\rfloor \times 28 = 588$$

$$\text{Final Pixels} = 308 \times 588 = 181,104$$

$$\text{Final Tokens} = \frac{181,104}{784} = 231$$

For CoS inference (processing one image at a time), each image has a 2,007,040 pixel limit:

$$\text{Initial Pixels} = 1980 \times 1080 = 2,138,400$$

$$\beta = \sqrt{\frac{2,138,400}{2,007,040}} \approx 1.032$$

$$\text{New Height} = \left\lfloor \frac{1080/1.032}{28} \right\rfloor \times 28 = 1036$$

$$\text{New Width} = \left\lfloor \frac{1980/1.032}{28} \right\rfloor \times 28 = 1904$$

$$\text{Final Pixels} = 1036 \times 1904 = 1,972,544$$

$$\text{Final Tokens} = \frac{1,972,544}{784} = 2516$$

Case 3: Ten 720×144 images

For Multi-Image inference, the per-image pixel limit is 200,704. Since each image's initial pixel count

is below this limit, no downscaling is applied:

$$\begin{aligned} \text{Image Pixels} &= 720 \times 144 = 103,680 \\ \text{Per-Image Limit} &= 200,704 \\ 103,680 &< 200,704 \quad (\text{Below the limit}) \\ \therefore &\text{ No downscaling required} \end{aligned}$$

Therefore, the original image dimensions are used for tokenization:

$$\begin{aligned} \text{Final Pixels per image} &= 103,680 \\ \text{Final Tokens per image} &= \frac{103,680}{784} \approx 132 \end{aligned}$$

For CoS inference, the per-image limit is 2,007,040. The initial pixel count is also well below this threshold, so no downscaling is required:

$$\begin{aligned} \text{Image Pixels} &= 720 \times 144 = 103,680 \\ \text{Limit} &= 2,007,040 \\ 103,680 &< 2,007,040 \quad (\text{Below the limit}) \\ \therefore &\text{ No downscaling required} \end{aligned}$$

The final token count is therefore identical to the Multi-Image case:

$$\begin{aligned} \text{Final Pixels per image} &= 103,680 \\ \text{Final Tokens per image} &= \frac{103,680}{784} \approx 132 \end{aligned}$$

J.2 Pixel Resizing for Table 3: limiting image token usage per image

In Table 3, tokens are limited per image. All cases apply a 1,004,520 pixel limit per image:

Case 1: Single 5120 × 2880 image

$$\begin{aligned} \text{Initial Pixels} &= 5120 \times 2880 = 14,745,600 \\ \beta &= \sqrt{\frac{14,745,600}{1,004,520}} \approx 3.83 \\ \text{New Height} &= \left\lfloor \frac{2880/3.83}{28} \right\rfloor \times 28 = 728 \\ \text{New Width} &= \left\lfloor \frac{5120/3.83}{28} \right\rfloor \times 28 = 1316 \\ \text{Final Pixels} &= 728 \times 1316 = 957,808 \\ \text{Final Tokens} &= \frac{957,808}{784} = 1222 \end{aligned}$$

Case 2: Ten 1980 × 1080 images

All inference methods apply the same resizing:

$$\begin{aligned} \text{Initial Pixels} &= 1980 \times 1080 = 2,138,400 \\ \beta &= \sqrt{\frac{2,138,400}{1,004,520}} \approx 1.459 \\ \text{New Height} &= \left\lfloor \frac{1080/1.459}{28} \right\rfloor \times 28 = 728 \\ \text{New Width} &= \left\lfloor \frac{1980/1.459}{28} \right\rfloor \times 28 = 1344 \\ \text{Final Pixels} &= 728 \times 1344 = 978,432 \\ \text{Final Tokens} &= \frac{978,432}{784} = 1248 \end{aligned}$$

Case 3: Ten 720 × 144 images

Since each image's pixels are below the limit, no resizing is applied:

$$\begin{aligned}\text{Pixels} &= 720 \times 144 = 103,680 \\ \text{Limit} &= 1,004,520 \\ 103,680 &< 1,004,520 \quad (\text{below the limit}) \\ &\therefore \text{No resizing required}\end{aligned}$$

Each image maintains its original dimensions:

$$\text{Final Pixels per image} = 720 \times 144 = 103,680 \quad \text{Final Tokens per image} = \frac{103,680}{784} = 132.24 \approx 132$$

K Efficiency metrics of baseline models and SCoPE VLMs in Table 2

Model	ANLS (%)					Visit Ratio (%)				
	MP-DocVQA	SlideVQA	M3DocVQA	DUDE	MMLong	MP-DocVQA	SlideVQA	M3DocVQA	DUDE	MMLong
Qwen2.5 VL 3B	47.78	23.23	19.54	31.73	9.08	59.17	30.19	56.81	64.52	23.41
Qwen2.5 VL 72B	80.83	66.73	35.83	48.43	–	67.77	50.28	57.47	84.32	–
SCoPE VLM SFT	74.49	59.88	46.13	42.82	16.89	75.59	55.79	53.40	82.40	122.68
SCoPE VLM SFT EGRPO	73.07	57.31	48.27	42.29	17.90	68.98	81.89	89.33	92.86	108.57

Model	Action Success Ratio (%)					No Answer Ratio (%)				
	MP-DocVQA	SlideVQA	M3DocVQA	DUDE	MMLong	MP-DocVQA	SlideVQA	M3DocVQA	DUDE	MMLong
Qwen2.5 VL 3B	89.59	65.73	74.39	85.05	67.09	2.27	1.76	10.72	7.40	1.28
Qwen2.5 VL 72B	99.33	99.54	96.85	95.31	–	0.98	3.09	12.70	9.52	–
SCoPE VLM SFT	94.42	78.49	87.39	90.90	62.52	1.33	0.79	3.50	3.71	31.44
SCoPE VLM EGRPO	96.03	81.89	89.33	92.86	66.16	0.42	0.61	1.75	3.40	23.37

Table 14: Performance-efficiency comparison with baseline Qwen2.5 VL models and SCoPE VLMs on document understanding benchmarks.

Table 14 presents the performance-efficiency comparison between baseline Qwen2.5 VL models and SCoPE VLMs. The baseline models demonstrate limited performance in the Chain of Scroll framework, with even Qwen2.5 VL 72B struggling despite its larger capacity. While SCoPE VLM SFT successfully learns the navigation task and achieves competitive ANLS scores, it fails to optimize efficiency metrics and action validity. In contrast, SCoPE VLM EGRPO successfully optimizes across all dimensions, achieving improved action success ratios and substantially reduced no-answer rates while maintaining comparable task performance to SCoPE VLM SFT.

L GUI control experiment in Table 4: Input prompt and overall results.

L.1 GUI experiments using the AitZ dataset: Full evaluation results for the General and Web-shopping test set splits.

Table 15 presents extended experimental results from Table 4 for the General and Web Shopping domains of the AitZ benchmark. Performance across five actions is evaluated using two metrics: Type Accuracy (Type Acc.) and Exact Match. Type Accuracy measures whether the predicted action type matches the ground truth, while Exact Match requires both the action type and its parameters to match completely. For instance, if the ground truth is “Scroll Up” but the model predicts “Scroll Down,” Type Accuracy is satisfied (both are scroll actions), but Exact Match is not. Goal Progress, an episode-level metric, represents the average progress toward the goal across all episodes.

Model	Accuracy Type	General						Web Shopping					
		Action					Episode	Action					Episode
		Click	Scroll	Type	Press	Stop		Click	Scroll	Type	Press	Stop	
Qwen 2.5-VL-3B	Type Acc.	73.77	9.49	83.95	44.07	66.67	35.46	81.67	43.05	80.30	42.55	83.57	41.17
	Exact Match	30.21	6.57	61.73	0.00	66.67		37.76	39.01	64.65	0.00	83.57	
SCoPE VLM 3B SFT (Ours)	Type Acc.	70.11	10.22	79.01	32.20	73.08	36.10	82.03	39.91	73.74	35.46	87.14	41.56
	Exact Match	31.00	8.76	56.17	0.00	73.08		37.95	36.77	57.58	0.00	87.14	
SCoPE VLM 3B EGRPO (Ours)	Type Acc.	72.66	8.03	82.10	35.59	76.28	37.51	82.31	46.64	80.81	40.43	84.29	42.32
	Exact Match	33.55	8.03	54.32	0.00	76.28		38.31	45.29	61.62	0.00	84.29	

Table 15: Detailed performance breakdown on the AitZ benchmark. This table reports Type Accuracy and Exact Match scores for each action type, along with the episode-level Goal Progress, across the General and Web Shopping test splits.

L.2 Example of input prompt

Figure 18 shows a prompt converted by merging the system prompts from the AitZ dataset and CoAT.

Prompt for GUI Agent Action Prediction

You are a smart and helpful visual assistant that is well trained to manipulate mobile phones. Your task is to navigate on the current screen to complete the user request. You are provided with:

- Two screenshots of the current mobile phone (one raw, one with UI element annotations).
- A brief summarization of the screen content.
- A history of actions attempting to accomplish the user request.
- The result of the previous action that led to the current screen.

You are required to decide on the next single-step valid action to be conducted on the current screen so as to fulfill the user request.

Valid Action Spaces

CLICK_ELEMENT(*idx*) Clicks on a visible UI element on the screen.

SCROLL(*direction*) Scrolls the content in the given direction ('up', 'down', 'left', 'right').

INPUT(*text*) Inputs the given text into an editable area.

PRESS_ENTER() Confirms or submits an input.

PRESS_HOME() Returns to the home screen.

PRESS_BACK() Returns to the previously visited screen.

STOP(*task_status*) Stops the task and sets its status ('success' or 'failure').

Output Format

Your response must be strictly structured in an XML-like format with the following tags:

<think> Your reasoning and analysis behind the chosen action.

<next> A brief, human-readable description of the next action.

<action> The precise API function call for the action.

User Request: [User Request]
History Actions: [History Actions]
Previous Action Result: [Previous Action Result]
Screen Content: [Screen Content Summary]
UI Elements: [List of UI Elements on the current screen]

You should return your thoughts, a brief description of the output action, and the final action. Return the thinking process in <think>...</think>, a brief description in <next>...</next>, and the action in <action>...</action> tags.

Figure 18: The modified AitZ prompt used for training and testing comprises four components: (1) general system instructions, (2) definitions of valid actions, (3) output format requirements, and (4) context-specific information for the current step, including user request, screen summary, and UI elements.