

Thinking with Map: Reinforced Parallel Map-Augmented Agent for Geolocalization

Yuxiang Ji^{1,2*} Yong Wang^{2†‡} Ziyu Ma² Yiming Hu² Hailang Huang²

Xuecai Hu² Guanhua Chen³ Liaoni Wu¹ Xiangxiang Chu²

¹Xiamen University ²AMAP, Alibaba Group ³Southern University of Science and Technology

<https://amap-ml.github.io/Thinking-with-Map>


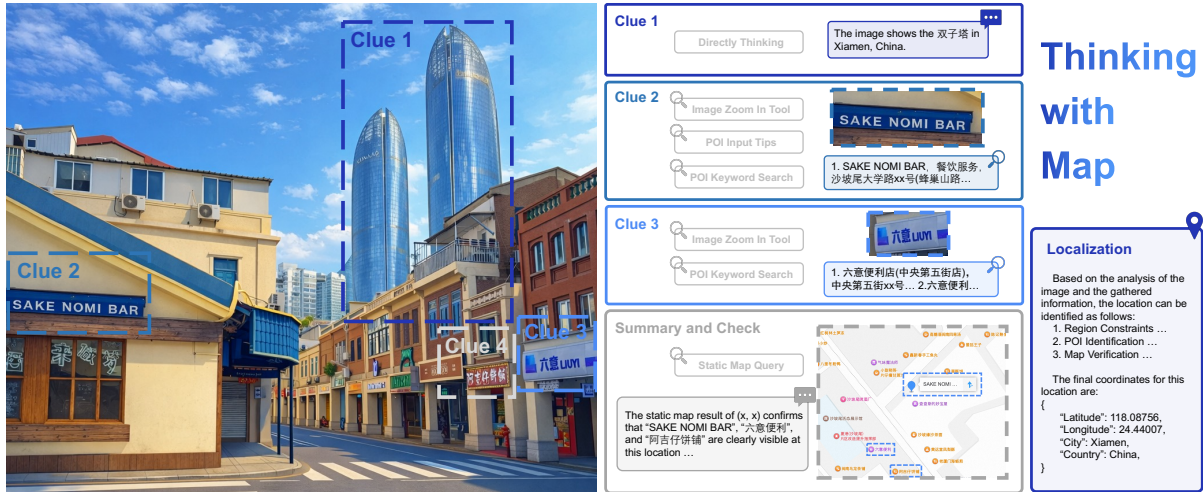


Figure 1: (Up) The illustration of a complete *Thinking with Map* process. (Bottom) Comparison with up-to-date open- and closed-source models on three geolocalization benchmarks. Our method is built upon the model *Qwen3-VL-30B-A3B*. POI represents Point of Interest.

Abstract

The image geolocalization task aims to predict the location where an image was taken anywhere on Earth using visual clues. Existing large vision-language model (LVL) approaches leverage world knowledge, chain-of-thought reasoning, and agentic capabilities, but overlook a common strategy used by humans — using maps. In this work, we first equip the model *Thinking with Map* ability and formulate it as an agent-in-the-map loop. We develop a two-stage optimization scheme for it, including agentic reinforcement learning (RL) fol-

lowed by parallel test-time scaling (TTS). The RL strengthens the agentic capability of model to improve sampling efficiency, and the parallel TTS enables the model to explore multiple candidate paths before making the final prediction, which is crucial for geolocalization. To evaluate our method on up-to-date and in-the-wild images, we further present MAPBench, a comprehensive geolocalization training and evaluation benchmark composed entirely of real-world images. Experimental results show that our method outperforms existing open- and closed-source models on most metrics, specifically improving Acc@500m from 8.0% to 22.1% compared to *Gemini-3-Pro* with Google Search/Map grounded mode.

*Work done during internship at AMAP, Alibaba Group.

†Project lead.

‡Corresponding author.

1 Introduction

Image geolocation is a challenging task to determine the latitude and longitude of an image as accurately as possible. Conventional vision research typically attributes this problem to a classification (Seo et al., 2018; Weyand et al., 2016; Müller-Budack et al., 2018; Clark et al., 2023) or retrieval (Ji et al., 2025a,b; Haas et al., 2024; Yang et al., 2021; Jia et al., 2024) task, achieving localization by predicting a region-level cell or retrieving the most similar image from a geo-tagged database. Although these methods are well established in applications such as indoor localization (Taira et al., 2018; Sarlin et al., 2019) and landmark recognition (Arandjelovic et al., 2016; Noh et al., 2017; Weyand et al., 2020; Zheng et al., 2009), they treat the entire image as a coupled feature for discrimination and fail to disentangle independent clues. This less interpretable paradigm is inherently constrained by the training data and is difficult to generalize to images in the wild.

In the era of large vision-language models (LVLM), geolocation can be viewed as a natural testbed for vision, understanding and reasoning. Beyond single-image discriminative paradigm, it requires LVLMs to inspect visual clues (e.g., climate, architecture, and cultural context) in detail, and reason over the complex intersection of evidence to make the final prediction. This process is closer to how human beings behave when inferring image locations. Recent studies follow frontier models (DeepSeek-AI et al., 2025; OpenAI, 2025; Google DeepMind, 2025b; Bai et al., 2025; Seed, 2025; Wang et al., 2025a) to further enhance such behavior by using chain-of-thought (CoT) reasoning (Li et al., 2024, 2025a; Jia et al., 2025) and incorporating external tools within the reasoning chain (Lai et al., 2025; Su et al., 2025; Qian et al., 2025; Wang et al., 2025b). However, despite their increased reasoning capability, these methods still depend on the model internal reasoning ability over knowledge.

In contrast, human beings rarely rely on internal reasoning alone for geolocation. When identifying visual clues, humans typically propose multiple location hypotheses and then verify them in turn using map tools. By querying points of interest (POIs), examining road topology, and checking spatial consistency, maps provide an essential mechanism for validating visual clues against real-world geography. Surprisingly, despite being the

most fundamental tool for geolocation, maps are almost absent from existing LVLM-based methods. To bridge this gap, we equip the LVLM with map tools for the first time, enabling the model to *Think with Map*. Specifically, we expose map interfaces such as keyword search, POI details lookup, and static map query as callable tools, allowing the model to retrieve information and verify visual clues in the structured map environment during reasoning. As illustrated in Figure 1, the process of *Thinking with Map* is a multi-turn agentic behavior. The model invokes tools based on multiple visual clues, then cross-validates the gathered evidence to produce the final prediction. We further formulate this localization process as an agent-in-the-map loop, in which the agent iteratively proposes and verifies location hypotheses.

Similar to human beings, when the model encounters an ambiguous image, it needs to go through an iterative process of repeated hypothesis generation and verification. However, simply increasing the reasoning budget to let the model explore sequentially not only leads to context explosion, but has also been found to yield marginal gains (Wen et al., 2025; Zheng et al., 2025a). Inspired by the success of Google Gemini in parallel thinking (Google DeepMind, 2025a), we also enable the model to explore multiple hypotheses in a parallel paradigm. Unlike conventional reasoning tasks, *Thinking with Map* inherently leaves a large number of map-API results in the reasoning trace. These factual outputs make the reasoning trajectory largely self-verifying. As Figures 2 and 4, we find that the LVLM can easily identify the better trajectory among multiple parallel *Thinking with Map* trajectories by causal relationships. Based on this observation, we introduce a simple parallel sampling with verifier framework for test-time scaling (TTS) in *Thinking with Map*. To further improve the model’s pass@K performance and enable more effective parallel sampling, we conduct agentic Reinforcement Learning (RL) training for *Thinking with Map*.

To evaluate our method, we propose MAPBench, which consists of up-to-date and broadly covered Chinese urban street-view and POI images. We categorize the data into two difficulty levels for further analysis of the model’s performance: *easy* cases are those that the model can localize at a glance, while *hard* cases contain less distinctive clues and are unlikely to be encountered during pre-training. We also conduct rigorous evaluations on recently re-




Figure 2: The *Thinking with Map* trajectories from parallel sampling. The abundant map-API results make the trajectories easily verified based on their causal relationships.

leased benchmarks, including IMAGEO-Bench (Li et al., 2025b) and GeoBench (Wang et al., 2025b). The results show that our method consistently outperforms all open-source models by a large margin and even surpasses *Gemini-3-Pro* (with Google Search/Map grounded mode) on most metrics. Our contributions are summarized as follows:

- We propose a map-augmented agent for the world-wide image geolocation, equipped with the model *Thinking with Map* ability.
- Building on the *Thinking with Map* capability, we propose a parallel-and-verifier TTS method and further enhance it with agentic RL.
- We evaluate our method on the proposed MAP-Bench and other geolocation benchmarks. The results show that our method outperforms all open- and closed-source models on most metrics.

2 Related Work

Worldwide Geolocation. Predicting the geographic location of a given image over the world is quite a challenging task (Haas et al., 2024; Qian et al., 2025). Over the past decades, computer vision research primarily treats this task as a retrieval or classification problem. The former relies on an enormous geo-tagged reference database as the retrieval gallery and introduces several large-scale benchmarks (Hays and Efros, 2008; Berton et al., 2022; Berton and Masone, 2025). The latter partitions the Earth into structured “geocells” and predicts geographic coordinates either directly or hierarchically (Müller-Budack et al., 2018; Clark et al., 2023; Haas et al., 2024). Recent LVLm-based methods leverage the visual understanding and reasoning capabilities of frontier models to directly infer a location from an image, without any database or map partitioning (Jia et al., 2025; Li et al., 2025a; Wang et al., 2024a; Li et al., 2025b; Huang et al., 2025). Although explicit reasoning reduces the black-box nature of the model, it cannot prevent hallucinations and biases of LVLms.

Tool Name	Parameter	Output
image_zoom_tool	Zoom in bounding box	Zoomed region image
poi_input_tips	Query text	Search Suggestions
poi_keyword_search	POI keyword	POI list
poi_detail_query	POI id	POI details
static_map_query	Location center	Static map image
satellite_map_query	Location center	Satellite map image

Table 1: The involved tools for *Thinking with Map*.

LVLm Powered Agent. As foundation models advance, researchers begin to focus on agentic capabilities and apply LVLm-powered agents to tasks that require interaction with open environments (Team et al., 2025; Li et al., 2025d; Gur et al., 2023; Yao et al., 2023; Hu et al., 2026; Zhu et al., 2026). Recent works employ an end-to-end agentic RL (Feng et al., 2025; Wang et al., 2025c; Ji et al., 2025c; Dong et al., 2025; Chu et al., 2025; Yuan et al., 2025; Li et al., 2025c; Xiong et al., 2025; Dai et al., 2026) to improve tool use and long-horizon decision-making abilities of the base model in specific task environments, demonstrating a broad vision. GeoVista (Wang et al., 2025b) applies this paradigm to geolocation by optimizing models to use vision and search tools for localization. Some studies (Qian et al., 2025) also argue that general search tools offer very limited benefits for localization. Beyond RL, some works also try to improve agent performance via test-time scaling methods such as parallel sampling (Wen et al., 2025), sequential revision (Zhu et al., 2025), and multi-agent exploration (Qiao et al., 2025).

3 Method

In this section, we present *Thinking with Map*, a map-augmented agent for improved LVLm-based geolocation. The overview of our method can be viewed in Figure 3. We first present the definition and implementation of *Thinking with Map* (§ 3.1). Then we use agentic RL to improve sampling efficiency by optimizing performance from pass@N to pass@K (§ 3.2). Finally, we apply parallel TTS to explore multiple candidate hypotheses

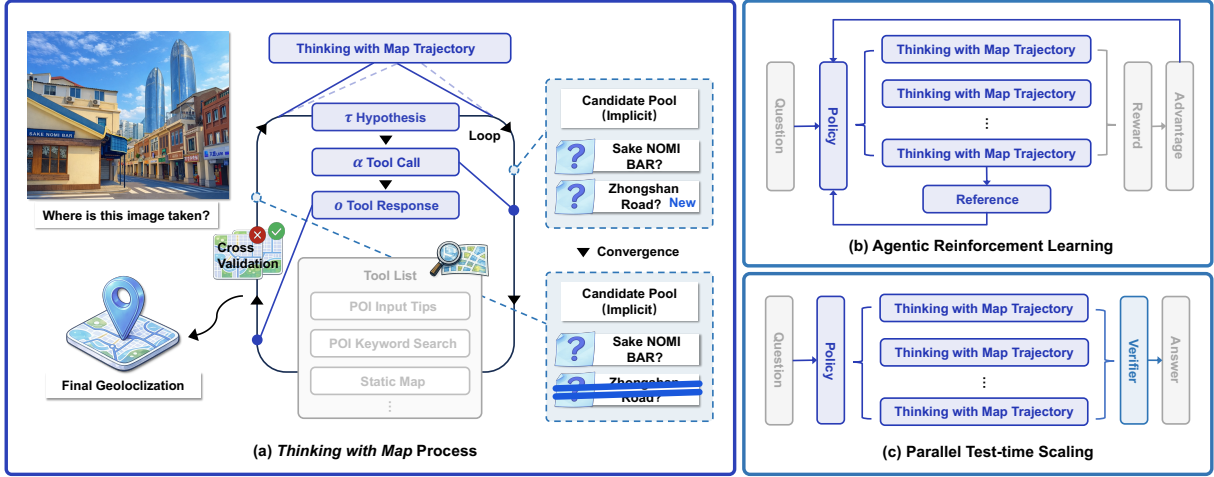


Figure 3: **(a)** The process of *Thinking with Map*, consists of an agent-in-the-map loop. During the loop, the agent implicitly maintains a candidate pool of hypotheses. **(b)** The agentic reinforcement learning for *Thinking with Map*. **(c)** The parallel test-time scaling with verifier pipeline for *Thinking with Map*.

during geolocalization, to gain performance from pass@K to pass@1 (§ 3.3).

3.1 Thinking with Map

Unlike direct discrimination or internal knowledge reasoning, we reformulate geolocalization as a *Thinking with Map* process. As Figure 3 (a), it follows an agent-in-the-map iterative loop of proposing location hypotheses, map retrieval, cross-validation and decision convergence. Formally, we model *Thinking with Map* as an iterative interaction process between a policy model π_θ and a structured map environment P_{env} . Given a geolocalization query $q_{\text{image, text}}$, at each iteration t the policy model can either propose a hypothesis τ_t (optional) explicitly/implicitly or verify existing hypotheses $\tau_{<t}$ through tool-call actions α_t to retrieve candidates within the map environment P_{env} . Then the map tool responses are treated as an observation o_t , and together with all previous interaction history, form an evidence chain s_t for cross-validation over the structured information:

$$s_t = \{(\tau_0, \alpha_0, o_0), \dots, (\tau_t, \alpha_t, o_t)\}, \quad (1)$$

$$p_\theta(\tau, \alpha, o|s_0) = \prod_{t=0}^{T-1} \left[\pi_\theta(\tau_t|s_t) \pi_\theta(\alpha_t|s_t, \tau_t) P_{\text{env}}(o_{t+1}|\alpha_t) \right]. \quad (2)$$

Let there be an implicit candidate pool \mathcal{C}_t in this iterative process. Then the evidence chain s_t composed by propositions and map observation at each step t can be regarded as a maintenance update to the candidate pool:

$$\mathcal{C}_{t+1} \triangleq \text{Update}(\mathcal{C}_t, s_t) \subseteq \mathcal{L}, \quad (3)$$

where \mathcal{L} is the overall location set. The policy model keeps maintaining this pool until it becomes sufficiently confident or the interaction budget is exhausted, and then selects the final answer from the candidate pool.

Here we provide a suite of map tools that human beings commonly use when looking for a location in Table 1. Among these tools, POI search serves as the primary information source from the map engine, helping the model obtain location details for specific places. Static and satellite maps then enable the model to verify and cross-check the surrounding scene and places around a candidate location. Due to the region-specific availability of map services, we employ two types of map API providers¹² to enable global geolocalization. In addition, we provide an `image_zoom_tool`, which helps the model progressively inspect visual clues in large-scene images.

3.2 RL for Map-augmented Agent

To enhance the model *Thinking with Map* capability, we adopt a widely explored RL paradigm to improve agentic performance from pass@N to pass@K. Instead of some recent *Qwen2.5-VL*-based works (Wang et al., 2025b; Lai et al., 2025; Zheng et al., 2025b) that adopt a two-stage SFT-then-RL training pipeline, we find that the *Qwen3-VL* model already shows basic tool-use ability after equipping it with map tools via the unified tool interface. Therefore, we directly apply agentic RL from this base model.

¹AMAP: <https://lbs.amap.com/>

²Google Map: <https://developers.google.com/maps>

Benchmark	Im2GPS3K	YFC100M	OSV-5M	IMAGEO-Bench	GeoBench	MAPBench
Reference	Vo et al. (2017)	Thomee et al. (2016)	Astruc et al. (2024)	Li et al. (2025b)	Wang et al. (2025b)	-
Number	3,000	100M	5M	6,152 / 2,929 / 220	512 / 512 / 108	2,500 / 2,500
Image Source	Flickr	Flickr	Mapillary	Mapillary/KartaView/Google Map	Web/Mapillary/Planetary Computer	AMAP
Up-to-date	✗	✗	✗	✗	✗	✓
Difficulty Tiering	✗	✗	✗	✗	✗	✓

Table 2: The comparison of MAPBench and existing geolocalization benchmarks.

As shown in Figure 3 (b), we adopt the Group Relative Policy Optimization (GRPO) (Shao et al., 2024) as the agentic RL algorithm. Specifically, for each geolocalization query q , the LVLM-based agent generates a group of agent trajectories $\{\mathcal{H}_i = (\tau_0, \alpha_0, o_0, \dots, \tau_T, \alpha_T)\}_{i=1}^G$ based on the previous policy $\pi_{\theta_{\text{old}}}$. The policy π_{θ} is then optimized by maximizing the advantages:

$$J_{\text{GRPO}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, \mathcal{H} \sim \pi_{\theta_{\text{old}}}(\cdot|q)} \left[\frac{1}{G} \sum_{i=1}^G \frac{1}{|\mathcal{H}^i|} \sum_{t=1}^{|\mathcal{H}^i|} \hat{r}_{i,t}(\theta) \hat{A}(\mathcal{H}^i) - \beta \mathbb{D}_{\text{KL}}(\pi_{\theta}(\mathcal{H}|q) \parallel \pi_{\text{ref}}(\mathcal{H}|q)) \right], \quad (4)$$

where $\hat{r}_{i,t}(\theta)$ is the importance sampling ratio, and clipping is applied in practice to stabilize RL training. We prompt the model to output answers in a fixed JSON format for each query, enabling structured parsing for the verifiable reward function. For geolocalization tasks evaluated by continuous distance, we simply use a piecewise discrete scheme that assigns different rewards to different distance ranges:

$$r = \begin{cases} 1, & \text{dis} \in [0, 500m) \\ 0.8, & \text{dis} \in [500m, 2km) \\ 0.6, & \text{dis} \in [2km, 10km) \\ 0.4, & \text{dis} \in [10km, 25km) \\ 0.2, & \text{dis} \in [25km, 200km) \\ 0.1, & \text{dis} \in [200km, 750km) \\ 0, & \text{dis} \in [750km, +\infty) \end{cases} \quad (5)$$

This hierarchical reward reflects different localization granularity, e.g., 500m for fine-level and 25km for city-level. In our experiments, this simple design works well with group-based RL and provides a discriminative learning signal.

3.3 Parallel Test-time Scaling

After RL training, the reinforced model can perform image localization reasoning while interacting with map tools. However, as with how human beings guess locations, images with limited clues

often require a sequence of hypotheses and verification steps. Due to the limited memory and reflection capabilities (Li et al., 2025d; Liu et al., 2025), such long-horizon sequential reasoning is a challenging task for LVLM-based agents.

Fortunately, we find that *Thinking with Map* trajectories naturally contain many self-verifiable factual information from map APIs, as shown in Figure 2. Therefore, we adopt a parallel-sampling pipeline with a verifier, where the model explores multiple paths through lightweight independent samples and a verifier aggregates the results. Formally, given a geolocalization query q and reinforced model π_{θ} , we first sample a set of N *Thinking with Map* trajectories in parallel as:

$$\left\{ \mathcal{H} | \mathcal{H}_i = \prod_{t=0}^{T-1} [\pi_{\theta}(\tau_t | s_t) \pi_{\theta}(\alpha_t | s_t, \tau_t) P_{\text{env}}(o_{t+1} | \alpha_t)] \right\}_{i=1}^N. \quad (6)$$

Then we feed the set of *Thinking with Map* trajectories, together with the original image and a simple instruction I into a LVLM-based verifier π_{verifier} , which summarizes the evidence and selects the most plausible prediction as:

$$\text{Answer} = \pi_{\text{verifier}}(q, \{\mathcal{H}\}_{i=1}^N, I). \quad (7)$$

As Figure 4, when we use *Qwen3-VL-30B-A3B* to perform parallel sampling with different numbers, verifier@N closely matches oracle best@N. In particular, when $N = 2$ or 4, the performance loss introduced by the verifier is almost negligible. With this parallel test-time scaling, we enable the model to explore multiple *Thinking with Map* hypotheses and aggregate self-verifiable trajectories to produce the final answer. This approach transfers performance gains from pass@K to pass@1.

4 Dataset

As Table 2, most existing geolocalization benchmarks use images collected earlier from Google Street View (Vo et al., 2017; Wang et al., 2024b, 2025b), Mapillary (Astruc et al., 2024), and Flickr (Thomee et al., 2016). In our early experiments, we identified several major issues with these datasets:

Method	MAPBench-test-easy ($Acc@Dis$, %)						MAPBench-test-hard ($Acc@Dis$, %)					
	Fine 500m	Local 2km	District 10km	City 25km	Region 200km	Country 750km	Fine 500m	Local 2km	District 10km	City 25km	Region 200km	Country 750km
<i>Closed Source Model</i>												
GPT-o3	7.68	35.23	86.64	88.98	89.82	92.32	0.05	0.74	4.53	9.10	20.73	44.13
GPT-5	9.02	34.39	87.48	90.32	92.99	95.49	0.05	0.79	4.10	8.94	22.30	47.45
Gemini-3-Pro (w/. Google Search/Map)	20.86	48.28	74.31	80.69	86.90	93.79	4.02	11.73	23.45	29.64	41.86	67.48
<i>Open Source Model</i>												
Qwen3-VL-235B-A22B	9.35	34.06	86.14	88.48	90.82	93.66	0.63	3.42	13.41	19.31	32.88	57.18
GLOBE-7B	0.17	6.53	42.21	58.29	73.70	82.91	0.05	0.85	6.34	11.35	27.68	52.29
GeoVista-7B (w/. Google Search)	0.33	4.17	28.21	39.39	47.74	51.08	0.00	0.53	4.16	6.52	10.94	18.99
Qwen3-VL-30B-A3B	4.01	21.87	68.61	71.95	75.63	83.31	0.21	1.89	10.36	14.20	28.56	52.76
+ Thinking with Map	33.10	40.28	53.68	56.89	59.94	64.73	10.83	12.05	16.08	19.06	25.58	38.28
+ Reinforcement Learning	41.51	50.88	76.88	79.35	83.07	89.67	12.33	14.67	26.89	31.62	42.58	67.17
+ Parallel×2 & Verifier	43.65	54.38	79.93	82.27	85.12	90.64	13.70	16.45	28.98	33.79	44.32	68.85
+ Parallel×4 & Verifier	44.98	55.02	80.27	82.27	85.79	91.30	14.86	17.40	29.88	34.37	45.21	68.85

Table 3: Comparison of *Thinking with Map* with open- and closed-source models on MAPBench. Results are reported as accuracy at multiple granularities ($Acc@Dis$). The **bold** indicates the best.

- **Timeliness.** Most of existing datasets are not up to date, and POIs shown in the images may no longer exist. As a result, they fail to assess an LLM-based geolocation method that leverage current, real-world knowledge. Moreover, obsolete POIs can contradict information from map APIs or the web, which can mislead the agent and impact localization performance.
- **Difficulty tiering.** Because LLMs are pre-trained on massive amount of world knowledge and images, many landmark-style images can be easily recognized and even memorized coordinates. Such images mainly measure memorization, but fail to evaluate the reasoning ability and capability to acquire and use external knowledge.
- **Global coverage.** Although existing datasets appear geographically diverse, their image sources bias them toward Europe and North America, with no coverage of China.

Based on these issues, we propose MAPBench, an up-to-date geolocation benchmark with broad coverage across China. The MAPBench consists of 5,000 nearby street-view images centered on POIs, with no POI repeated across samples. We randomly split the dataset into 2,500 training samples and 2,500 test samples. Furthermore, we categorize test samples based on the zero-shot predictions of three base models *GPT-5*, *GPT-o3*, and *Qwen3-VL-235B-A22B*. The sample is labeled as easy if at least two models predict locations within $10km$ of the ground truth, and labeled as hard otherwise. The easy split evaluates the memorization and world knowledge of base model, while the hard split specifically assesses agentic capabilities. As a result, 599 test samples are labeled as easy, while

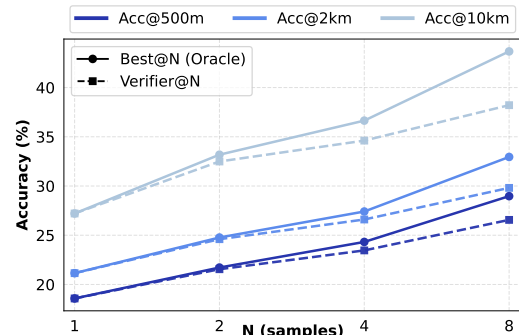


Figure 4: The comparison on parallel sampling.

the remaining 1,901 test samples are labeled as hard.

5 Experiment

5.1 Experimental Setup

Models. We compare the proposed *Thinking with Map* against multiple series of state-of-the-art closed-source models, including *GPT-o3* and *GPT-5* from OpenAI, and *Gemini-3-Pro* from Google. We also compare against a large-scale open-source model *Qwen3-VL-235B-A22B* from Alibaba, as well as two open-source geolocation methods GLOBE (Li et al., 2025a) and GeoVista (Wang et al., 2025b). Our method is built upon *Qwen3-VL-30B-A3B-Instruct*.

Datasets. To evaluate our method for worldwide geolocation capability, in addition to the proposed MAPBench, we also include two recently released benchmarks IMAGEO-Bench (Li et al., 2025b) and GeoBench (Wang et al., 2025b). In particular, we use an IMAGEO-2 subset as it exhibits greater difficulty in our experiments. For

Method	GeoBench ($Acc@Dis$, %)						IMAGEO-2-test ($Acc@Dis$, %)					
	Fine 500m	Local 2km	District 10km	City 25km	Region 200km	Country 750km	Fine 500m	Local 2km	District 10km	City 25km	Region 200km	Country 750km
<i>Closed Source Model</i>												
GPT-o3	33.08	50.75	61.99	64.45	67.67	73.45	9.66	18.76	27.41	30.85	47.06	67.04
GPT-5	33.30	46.90	59.64	63.17	67.13	75.05	11.14	19.91	28.12	32.62	50.62	72.78
Gemini-3-Pro (w/. Google Search/Map)	37.79	47.22	51.61	53.64	56.32	59.10	16.33	27.33	33.22	37.00	48.78	62.67
<i>Open Source Model</i>												
Qwen3-VL-235B-A22B	19.38	46.68	66.60	71.52	78.05	91.54	1.78	5.66	11.88	15.76	34.07	62.38
GLOBE-7B	11.21	43.69	69.15	71.72	78.50	88.78	0.33	1.33	4.77	7.77	31.74	65.37
GeoVista-7B (w/. Google Search)	6.85	26.55	45.50	51.17	54.81	58.35	0.22	1.11	3.77	5.66	12.54	20.08
Qwen3-VL-30B-A3B	12.21	40.47	66.60	71.52	76.02	90.90	1.11	3.22	8.77	12.99	34.52	65.82
+ Thinking with Map	49.82	59.05	66.64	68.28	71.72	81.36	17.75	19.33	21.55	23.72	31.93	47.36
+ Reinforcement Learning	52.57	64.01	72.83	74.53	77.92	86.62	18.64	20.50	23.77	27.19	42.59	72.41
+ Parallel×2 & Verifier	55.61	67.06	75.23	76.17	79.44	87.38	19.64	21.86	25.53	29.08	45.06	74.14
+ Parallel×4 & Verifier	57.94	69.16	76.17	77.57	80.84	89.02	20.53	22.64	26.19	30.19	46.06	75.69

Table 4: Comparison of *Thinking with Map* with open- and closed-source models on GeoBench and IMAGEO. Results are reported as accuracy at multiple granularities ($Acc@Dis$). The **bold** indicates the best.

RL Method	MAPBench-test-all ($Acc@Dis$, %)					
	Fine 500m	Local 2km	District 10km	City 25km	Region 200km	Country 750km
Qwen3-VL-30B-A3B	1.12	6.67	24.29	28.01	39.82	60.07
+ image_zoom_tool	1.48	6.81	23.27	26.53	35.36	53.60
+ web_search_tool	1.77	9.55	26.05	29.34	36.73	49.73
+ map_tool	16.16	18.80	25.07	28.11	33.80	44.61

Table 5: The ablation study on tool types.

RL training, we use the MAPBench training set and 2,000 examples from IMAGEO-2, achieving globally covered samples. More details are in Appendix A.

Evaluation. To analyze the model localization accuracy at different granularities, we report $acc@dis$ at six levels (500m@Fine, 2km@Local, 10km@District, 25km@City, 200km@Region and 750km@Country), with distance thresholds matching the reward settings. Specifically, a prediction is considered correct if its distance to the ground truth is below the corresponding threshold.

Settings. For closed-source models, we query them directly via APIs. Some of them have built-in tool-use capabilities, such as image manipulation tools of *GPT-o3* and Google Search / Google Maps grounded mode of *Gemini-3-Pro*. For the two open-source geolocalization methods, we follow the original papers to set the corresponding inference hyperparameters, and equip *GeoVista-7B* with `image_zoom_tool` and `web_search_tool` via a unified tool interface. If not specified, we use *Qwen3-VL-235B-A22B* as the verifier for the results of parallel sampling. More details are in Appendix B.1.

5.2 Main Results

As shown in Tables 3 and 4, our proposed *Thinking with Map* method achieves the best performance comparing with all open- and closed-source models on most metrics across four test sets. In particular, for fine localization $Acc@500m$, our method outperforms the best closed-source model *Gemini-3-Pro* on MAPBench-test-hard by a large margin, from 4.02% to 14.86%. The substantial gains on GeoBench and IMAGEO-2-test also show improving $Acc@500m$ from 37.79% to 57.94% and 16.33% to 20.53%, respectively. Due to the base model used in existing open-source geolocalization methods are relatively small (7B), their performance also cannot match closed-source models. On the other hand, our task directly predicts latitude and longitude, which differs from the models original training targets and can hurt performance.

In our experiments, we find that the capability of base model can determine coarse-grained localization performance (e.g., $Acc@25km$ and $Acc@200km$), while the search and map tools can greatly enhance fine-grained localization performance (e.g., $Acc@500m$). For example, on MAPBench-test-hard, all base models achieve nearly 0% accuracy for fine-localization, while only *Gemini-3-Pro* with Google Search/Map grounded mode and our method reach 4.02% and 14.86% respectively. However, directly integrating map tools can also lead to negative effects. Noisy information from the map tools (e.g., wrong search results) may introduce substantial bias in coarse localization, which is reflected by the performance drop in “+ Thinking with Map” row. This performance drop is addressed after reinforcement

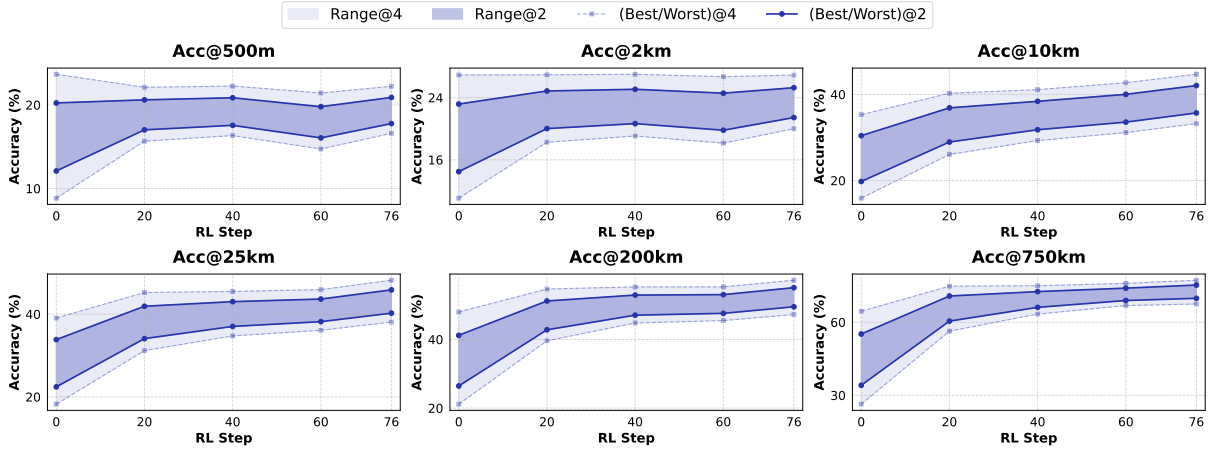


Figure 5: The evolution of pass@K accuracy across RL training steps on MAPBench.

Verifier Model	MAPBench-test-easy ($Acc@Dis$, %)						MAPBench-test-hard ($Acc@Dis$, %)					
	Fine 500m	Local 2km	District 10km	City 25km	Region 200km	Country 750km	Fine 500m	Local 2km	District 10km	City 25km	Region 200km	Country 750km
Verifier@2												
Qwen3-VL-30B-A3B	43.48	53.18	77.93	80.27	83.95	89.97	13.64	16.34	28.34	32.95	42.99	67.42
Qwen3-VL-235B-A22B	43.65	54.35	79.93	82.27	85.12	90.64	13.70	16.45	28.98	33.79	44.32	68.86
GPT-5	43.81	54.01	79.93	82.27	85.79	91.47	13.86	16.61	28.45	33.21	43.89	68.06
Verifier@4												
Qwen3-VL-30B-A3B	44.15	53.85	79.26	81.10	85.12	90.13	14.65	17.03	28.98	33.74	44.32	68.48
Qwen3-VL-235B-A22B	44.98	55.02	80.27	82.27	85.79	91.30	14.86	17.40	29.88	34.37	45.21	68.85
GPT-5	45.82	54.85	80.94	83.11	86.96	92.31	14.86	17.19	29.88	34.58	44.79	68.96

Table 6: The ablation study on verifier models. Verifier@N means verifier with N parallel samples.

learning training. Notably, our *Thinking with Map* method already outperforms the other approaches even before incorporating parallel TTS.

When incorporating parallel TTS, our method achieves further performance gains, and the improvement is positively correlated with the number of parallel samples. This gain trend is consistent with that of the base model with parallel TTS in Figure 4.

5.3 Quantitative Analysis

Different Tools. Here we explore how different types of tools affect the geolocalization task. We use *Qwen3-VL-30B-A3B-Instruct* as the base model and integrate three types of tools separately. The results in Table 5 align with our earlier discussion in § 5.2. All three tool types improve fine-grained localization ($< 2km$), but hurt coarse-grained localization ($> 200km$). Among them, *image_zoom_tool* and *web_search_tool* bring very marginal improvements, whereas *map_tool* yields a clear gain from 1.12% to 16.16% on Acc@500m.

Evolution of Pass@K across RL. Many recent studies (Yue et al., 2025) explore the impact of RL-based post-training on LVLMs. Here we evaluate the effect of RL on the geolocalization task by examining the evolution of pass@K accuracy throughout RL training, as shown in Figure 5. As RL training progresses, the prediction accuracy at all granularities shows lower variance, as Range@2/4 becoming smaller. This trend is consistent with the view that RL helps optimize performance from pass@K toward pass@1. Notably, accuracy at larger distance thresholds (i.e., $Dis > 10km$) shows a clear upward trend under best@N. This suggests that RL also helps the model achieve stronger pass@K from pass@N ($K < N$). However, Best@500m shows little to no improvement, and can even limit exploration.

Different Verifier Models. To further validate the role of the verifier and investigate what makes a better verifier in parallel TTS, we experiment with different verifier models in Table 6. The results show that when the parallel size $N = 2$, the choice of model has only a minor impact, and a 30B model

is already sufficient to serve as a strong verifier. As the parallel size increases, the verifying task becomes harder, and the impact of model capacity becomes correspondingly more important.

6 Conclusion

In this work, we propose a map-augmented agent for image geolocalization, to enable model *Thinking with Map*. We model this process as an agent-in-the-map loop of proposing hypotheses, map retrieval, cross-validation, and decision convergence. Based on this, we propose a two-stage optimization approach that combines agentic RL and parallel test-time scaling to gain pass@N capability within a single query. Experimental results show that our method outperforms all open- and closed-source models on most metrics.

Limitation

In this work, we equip the agent with map tools, enabling the LVLM agent to do geolocalization by iteratively interacting within a structured map environment. Although the model can perform evidence-grounded reasoning with map tools, we find that its map-use ability still falls far short of human performance. For example, we do not observe the model inferring orientation from relative spatial relationships, which is a common strategy humans use when estimating locations. For agentic RL, our training data remain very limited, which constrains the model to learn in open environments. One promising avenue for future work is to investigate what emergent capabilities arise when scaling up this RL paradigm. Finally, we consider parallel TTS a pragmatic interim solution that compensates for the current limitations of a single agent. How to build a single agent with stronger long-horizon problem-solving capabilities remains an open problem.

Acknowledgment

We acknowledge the helpful discussion with Kaibin Tian, the author of SeekWorld (Tian et al., 2025), and our intern colleagues, Shidong Yang and Zengbin Wang for their assistance.

References

Relja Arandjelovic, Petr Gronat, Akihiko Torii, Tomas Pajdla, and Josef Sivic. 2016. Netvlad: Cnn architecture for weakly supervised place recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5297–5307.

Guillaume Astruc, Nicolas Dufour, Ioannis Siglidis, Constantin Aronsohn, Nacim Bouia, Stephanie Fu, Romain Loiseau, Van Nguyen Nguyen, Charles Raude, Elliot Vincent, and 1 others. 2024. Openstreetview-5m: The many roads to global visual geolocation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21967–21977.

Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, Wenbin Ge, Zhi-fang Guo, Qidong Huang, Jie Huang, Fei Huang, Binyuan Hui, Shutong Jiang, Zhaohai Li, Mingsheng Li, and 45 others. 2025. *Qwen3-VL Technical Report*. *arXiv preprint*. ArXiv:2511.21631 [cs].

Gabriele Berton and Carlo Masone. 2025. Megaloc: One retrieval to place them all. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 2861–2867.

Gabriele Berton, Carlo Masone, and Barbara Caputo. 2022. Rethinking visual geo-localization for large-scale applications. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4878–4888.

Xiangxiang Chu, Hailang Huang, Xiao Zhang, Fei Wei, and Yong Wang. 2025. *GPG: A Simple and Strong Reinforcement Learning Baseline for Model Reasoning*. *arXiv preprint*. ArXiv:2504.02546 [cs].

Brandon Clark, Alec Kerrigan, Parth Parag Kulkarni, Vicente Vivanco Cepeda, and Mubarak Shah. 2023. *Where We Are and What We’re Looking At: Query Based Worldwide Image Geo-localization Using Hierarchies and Scenes*. *arXiv preprint*. ArXiv:2303.04249 [cs].

Yanqi Dai, Yuxiang Ji, Xiao Zhang, Yong Wang, Xiangxiang Chu, and Zhiwu Lu. 2026. Harder is better: Boosting mathematical reasoning via difficulty-aware grpo and multi-aspect question reformulation. *arXiv preprint arXiv:2601.20614*.

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, and 181 others. 2025. *DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning*. *arXiv preprint*. ArXiv:2501.12948 [cs].

Guanting Dong, Hangyu Mao, Kai Ma, Licheng Bao, Yifei Chen, Zhongyuan Wang, Zhongxia Chen, Jia-zhen Du, Huiyang Wang, Fuzheng Zhang, Guorui Zhou, Yutao Zhu, Ji-Rong Wen, and Zhicheng Dou. 2025. *Agentic Reinforced Policy Optimization*. *arXiv preprint*. ArXiv:2507.19849 [cs].

Lang Feng, Zhenghai Xue, Tingcong Liu, and Bo An. 2025. *Group-in-Group Policy Optimization for LLM Agent Training*. *arXiv preprint*. ArXiv:2505.10978 [cs].

- Google DeepMind. 2025a. [Advanced version of gemini with deep think officially achieves gold-medal standard at the international mathematical olympiad](#). Blog. Accessed: 2025-12-25.
- Google DeepMind. 2025b. [Gemini 3 pro model card](#). Model card. Accessed: 2025-12-25.
- Izzeddin Gur, Hiroki Furuta, Austin Huang, Mustafa Safdari, Yutaka Matsuo, Douglas Eck, and Aleksandra Faust. 2023. A real-world webagent with planning, long context understanding, and program synthesis. *arXiv preprint arXiv:2307.12856*.
- Lukas Haas, Michal Skreta, Silas Alberti, and Chelsea Finn. 2024. [PIGEON: Predicting Image Geolocations](#). *arXiv preprint*. ArXiv:2307.05845 [cs].
- James Hays and Alexei A Efros. 2008. Im2gps: estimating geographic information from a single image. In *2008 IEEE conference on computer vision and pattern recognition*, pages 1–8. IEEE.
- Jinwei Hu, Xinmiao Huang, Youcheng Sun, Yi Dong, and Xiaowei Huang. 2026. Lying with truths: Open-channel multi-agent collusion for belief manipulation via generative montage. *arXiv preprint arXiv:2601.01685*.
- Jingyuan Huang, Jen-tse Huang, Ziyi Liu, Xiaoyuan Liu, Wenxuan Wang, and Jieyu Zhao. 2025. [AI Sees Your Location, But With A Bias Toward The Wealthy World](#). *arXiv preprint*. ArXiv:2502.11163 [cs].
- Yuxiang Ji, Boyong He, Zhuoyue Tan, and Liaoni Wu. 2025a. Game4loc: A uav geo-localization benchmark from game data. In *AAAI*.
- Yuxiang Ji, Boyong He, Zhuoyue Tan, and Liaoni Wu. 2025b. Mmgeo: Multimodal compositional geo-localization for uavs. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 25165–25175.
- Yuxiang Ji, Ziyu Ma, Yong Wang, Guanhua Chen, Xiangxiang Chu, and Liaoni Wu. 2025c. [Tree Search for LLM Agent Reinforcement Learning](#). *arXiv preprint*. ArXiv:2509.21240 [cs].
- Pengyue Jia, Yiding Liu, Xiaopeng Li, Yuhao Wang, Yantong Du, Xiao Han, Xuetao Wei, Shuaiqiang Wang, Dawei Yin, and Xiangyu Zhao. 2024. [G3: An Effective and Adaptive Framework for Worldwide Geolocalization Using Large Multi-Modality Models](#). *arXiv preprint*. ArXiv:2405.14702 [cs].
- Pengyue Jia, Yingyi Zhang, Xiangyu Zhao, and Sharon Li. 2025. [GeoArena: An Open Platform for Benchmarking Large Vision-language Models on WorldWide Image Geolocalization](#). *arXiv preprint*. ArXiv:2509.04334 [cs].
- Xin Lai, Junyi Li, Wei Li, Tao Liu, Tianjian Li, and Hengshuang Zhao. 2025. [Mini-o3: Scaling Up Reasoning Patterns and Interaction Turns for Visual Search](#). *arXiv preprint*. ArXiv:2509.07969 [cs].
- Ling Li, Yu Ye, Yao Zhou, Bingchuan Jiang, and Wei Zeng. 2024. Georeasoner: Geo-localization with reasoning in street views using a large vision-language model. *arXiv preprint arXiv:2406.18572*.
- Ling Li, Yao Zhou, Yuxuan Liang, Fugee Tsung, and Jiaheng Wei. 2025a. [Recognition through Reasoning: Reinforcing Image Geo-localization with Large Vision-Language Models](#). *arXiv preprint*. ArXiv:2506.14674 [cs].
- Lingyao Li, Runlong Yu, Qikai Hu, Bowei Li, Min Deng, Yang Zhou, and Xiaowei Jia. 2025b. [From Pixels to Places: A Systematic Benchmark for Evaluating Image Geolocalization Ability in Large Language Models](#). *arXiv preprint*. ArXiv:2508.01608 [cs].
- Renda Li, Hailang Huang, Fei Wei, Feng Xiong, Yong Wang, and Xiangxiang Chu. 2025c. Adacurl: Adaptive curriculum reinforcement learning with invalid sample mitigation and historical revisiting. *arXiv preprint arXiv:2511.09478*.
- Xiaoxi Li, Wenxiang Jiao, Jiarui Jin, Guanting Dong, Jiajie Jin, Yinuo Wang, Hao Wang, Yutao Zhu, Ji-Rong Wen, Yuan Lu, and Zhicheng Dou. 2025d. [Deep-Agent: A General Reasoning Agent with Scalable Toolsets](#). *arXiv preprint*. ArXiv:2510.21618 [cs].
- Tengxiao Liu, Zifeng Wang, Jin Miao, I.-Hung Hsu, Jun Yan, Jiefeng Chen, Rujun Han, Fangyuan Xu, Yanfei Chen, Ke Jiang, Samira Daruki, Yi Liang, William Yang Wang, Tomas Pfister, and Chen-Yu Lee. 2025. [Budget-Aware Tool-Use Enables Effective Agent Scaling](#). *arXiv preprint*. ArXiv:2511.17006 [cs].
- Eric Müller-Budack, Kader Pustu-Iren, and Ralph Ewerth. 2018. [Geolocation Estimation of Photos Using a Hierarchical Model and Scene Classification](#). In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Computer Vision – ECCV 2018*, volume 11216, pages 575–592. Springer International Publishing, Cham. Series Title: Lecture Notes in Computer Science.
- Hyeonwoo Noh, Andre Araujo, Jack Sim, Tobias Weyand, and Bohyung Han. 2017. Large-scale image retrieval with attentive deep local features. In *Proceedings of the IEEE international conference on computer vision*, pages 3456–3465.
- OpenAI. 2025. [Openai o3-mini system card](#). System card. Accessed: 2025-12-25.
- Zhaofang Qian, Hardy Chen, Zeyu Wang, Li Zhang, Zijun Wang, Xiaoke Huang, Hui Liu, Xianfeng Tang, Zeyu Zheng, Haoqin Tu, Cihang Xie, and Yuyin Zhou. 2025. [Where on Earth? A Vision-Language Benchmark for Probing Model Geolocation Skills Across Scales](#). *arXiv preprint*. ArXiv:2510.10880 [cs].
- Zile Qiao, Guoxin Chen, Xuanzhong Chen, Donglei Yu, Wenbiao Yin, Xinyu Wang, Zhen Zhang, Baixuan

- Li, Huifeng Yin, Kuan Li, Rui Min, Minpeng Liao, Yong Jiang, Pengjun Xie, Fei Huang, and Jingren Zhou. 2025. [WebResearcher: Unleashing unbounded reasoning capability in Long-Horizon Agents](#). *arXiv preprint*. ArXiv:2509.13309 [cs].
- Paul-Edouard Sarlin, Cesar Cadena, Roland Siegwart, and Marcin Dymczyk. 2019. From coarse to fine: Robust hierarchical localization at large scale. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12716–12725.
- Seed. 2025. [Seed-1.8 model card](#). Model card. Accessed: 2025-12-25.
- Paul Hongsuck Seo, Tobias Weyand, Jack Sim, and Bohyung Han. 2018. [CPLaNet: Enhancing Image Geolocalization by Combinatorial Partitioning of Maps](#). *arXiv preprint*. ArXiv:1808.02130 [cs].
- 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](#). *arXiv preprint*. ArXiv:2402.03300 [cs].
- Zhaochen Su, Peng Xia, Hangyu Guo, Zhenhua Liu, Yan Ma, Xiaoye Qu, Jiaqi Liu, Yanshu Li, Kaide Zeng, Zhengyuan Yang, and 1 others. 2025. Thinking with images for multimodal reasoning: Foundations, methods, and future frontiers. *arXiv preprint* [arXiv:2506.23918](#).
- Hajime Taira, Masatoshi Okutomi, Torsten Sattler, Mircea Cimpoi, Marc Pollefeys, Josef Sivic, Tomas Pajdla, and Akihiko Torii. 2018. Inloc: Indoor visual localization with dense matching and view synthesis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7199–7209.
- Yunhao Tang, Kunhao Zheng, Gabriel Synnaeve, and Rémi Munos. 2025. Optimizing language models for inference time objectives using reinforcement learning. *arXiv preprint* [arXiv:2503.19595](#).
- Kimi Team, Yifan Bai, Yiping Bao, Guanduo Chen, Jiahao Chen, Ningxin Chen, Ruijue Chen, Yanru Chen, Yuankun Chen, Yutian Chen, Zhuofu Chen, Jialei Cui, Hao Ding, Mengnan Dong, Angang Du, Chenzhuang Du, Dikang Du, Yulun Du, Yu Fan, and 150 others. 2025. [Kimi K2: Open Agentic Intelligence](#). *arXiv preprint*. ArXiv:2507.20534 [cs].
- Bart Thomee, David A Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. 2016. Yfcc100m: The new data in multimedia research. *Communications of the ACM*, 59(2):64–73.
- Kaibin Tian, Zijie Xin, and Jiazhen Liu. 2025. SeekWorld: Geolocation is a natural RL task for o3-like visual clue-tracking. <https://github.com/TheEighthDay/SeekWorld>. GitHub repository.
- Nam Vo, Nathan Jacobs, and James Hays. 2017. Revisiting im2gps in the deep learning era. In *Proceedings of the IEEE international conference on computer vision*, pages 2621–2630.
- Christian Walder and Deep Karkhanis. 2025. Pass@k policy optimization: Solving harder reinforcement learning problems. *arXiv preprint* [arXiv:2505.15201](#).
- Yifei Wang, Feng Xiong, Yong Wang, Linjing Li, Xiangxiang Chu, and Daniel Dajun Zeng. 2025a. Position bias mitigates position bias: Mitigate position bias through inter-position knowledge distillation. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 1495–1512.
- Yikun Wang, Zuyan Liu, Ziyi Wang, Pengfei Liu, Han Hu, and Yongming Rao. 2025b. [GeoVista: Web-Augmented Agentic Visual Reasoning for Geolocalization](#). *arXiv preprint*. ArXiv:2511.15705 [cs].
- Zhiqiang Wang, Dejia Xu, Rana Muhammad Shahroz Khan, Yanbin Lin, Zhiwen Fan, and Xingquan Zhu. 2024a. [LLMGeo: Benchmarking Large Language Models on Image Geolocation In-the-wild](#). *arXiv preprint*. ArXiv:2405.20363 [cs].
- Zhiqiang Wang, Dejia Xu, Rana Muhammad Shahroz Khan, Yanbin Lin, Zhiwen Fan, and Xingquan Zhu. 2024b. Llmgeo: Benchmarking large language models on image geolocation in-the-wild. *arXiv preprint* [arXiv:2405.20363](#).
- Zihan Wang, Kangrui Wang, Qineng Wang, Pingyue Zhang, Linjie Li, Zhengyuan Yang, Xing Jin, Kefan Yu, Minh Nhat Nguyen, Licheng Liu, Eli Gottlieb, Yiping Lu, Kyunghyun Cho, Jiajun Wu, Li Fei-Fei, Lijuan Wang, Yejin Choi, and Manling Li. 2025c. [RAGEN: Understanding Self-Evolution in LLM Agents via Multi-Turn Reinforcement Learning](#). *arXiv preprint*. ArXiv:2504.20073 [cs].
- Hao Wen, Yifan Su, Feifei Zhang, Yunxin Liu, Yunhao Liu, Ya-Qin Zhang, and Yuanchun Li. 2025. [ParaThinker: Native Parallel Thinking as a New Paradigm to Scale LLM Test-time Compute](#). *arXiv preprint*. ArXiv:2509.04475 [cs].
- Tobias Weyand, Andre Araujo, Bingyi Cao, and Jack Sim. 2020. Google landmarks dataset v2-a large-scale benchmark for instance-level recognition and retrieval. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2575–2584.
- Tobias Weyand, Ilya Kostrikov, and James Philbin. 2016. [PlaNet - Photo Geolocation with Convolutional Neural Networks](#). volume 9912, pages 37–55. ArXiv:1602.05314 [cs].
- Feng Xiong, Hongling Xu, Yifei Wang, Runxi Cheng, Yong Wang, and Xiangxiang Chu. 2025. Hs-star: Hierarchical sampling for self-taught reasoners via difficulty estimation and budget reallocation. *arXiv preprint* [arXiv:2505.19866](#).

An Yan, Zhankui He, Jiacheng Li, Tianyang Zhang, and Julian McAuley. 2023. Personalized showcases: Generating multi-modal explanations for recommendations. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2251–2255.

Hongji Yang, Xiufan Lu, and Yingying Zhu. 2021. Cross-view Geo-localization with Layer-to-Layer Transformer. In *Advances in Neural Information Processing Systems*, volume 34, pages 29009–29020. Curran Associates, Inc.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023. ReAct: Synergizing Reasoning and Acting in Language Models. *arXiv preprint*. ArXiv:2210.03629 [cs].

Zhenlong Yuan, Xiangyan Qu, Chengxuan Qian, Rui Chen, Jing Tang, Lei Sun, Xiangxiang Chu, Dapeng Zhang, Yiwei Wang, Yujun Cai, and 1 others. 2025. Video-star: Reinforcing open-vocabulary action recognition with tools. *arXiv preprint arXiv:2510.08480*.

Yang Yue, Zhiqi Chen, Rui Lu, Andrew Zhao, Zhaokai Wang, Shiji Song, and Gao Huang. 2025. Does reinforcement learning really incentivize reasoning capacity in llms beyond the base model? *arXiv preprint arXiv:2504.13837*.

Tong Zheng, Hongming Zhang, Wenhao Yu, Xiaoyang Wang, Runpeng Dai, Rui Liu, Huiwen Bao, Chengsong Huang, Heng Huang, and Dong Yu. 2025a. Parallel-R1: Towards Parallel Thinking via Reinforcement Learning. *arXiv preprint*. ArXiv:2509.07980 [cs].

Yan-Tao Zheng, Ming Zhao, Yang Song, Hartwig Adam, Ulrich Buddemeier, Alessandro Bissacco, Fernando Brucher, Tat-Seng Chua, and Hartmut Neven. 2009. Tour the world: Building a web-scale landmark recognition engine. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1085–1092.

Ziwei Zheng, Michael Yang, Jack Hong, Chenxiao Zhao, Guohai Xu, Le Yang, Chao Shen, and Xing Yu. 2025b. Deepeyes: Incentivizing "thinking with images" via reinforcement learning. *arXiv preprint arXiv:2505.14362*.

King Zhu, Hanhao Li, Siwei Wu, Tianshun Xing, Dehua Ma, Xiangru Tang, Minghao Liu, Jian Yang, Jiaheng Liu, Yuchen Eleanor Jiang, Changwang Zhang, Chenghua Lin, Jun Wang, Ge Zhang, and Wangchunshu Zhou. 2025. Scaling Test-time Compute for LLM Agents. *arXiv preprint*. ArXiv:2506.12928 [cs].

Wei Zhu, Zhiwen Tang, and Kun Yue. 2026. Symphony: Synergistic multi-agent planning with heterogeneous language model assembly. *arXiv preprint arXiv:2601.22623*.




Figure 6: Reward dynamics across RL training.

Appendix

A Datasets

Here we provide more details of our proposed MAPBench. We uniformly and randomly sample 5,000 valid POIs across 20 cities in China, and for each POI we randomly select either a street-view or storefront photo, forming a final set of 5,000 images. This simple construction process ensures that the samples are both up-to-date and broadly coverage.

Considering the worldwide coverage and timeliness of the image sources, in addition to our proposed MAPBench, we also use two recently released datasets for global images:

- **IMAGEO-2** is a subset of IMAGEO-Bench (Li et al., 2025b), and constructed from crowd-sourced images from Google Map POIs. The original data are released by Yan et al. (2023), then compiled and filtered to final 2,929 images. We use 2,027 randomly sampled instances for training (as IMAGEO-2-train) and the remaining 902 instances for testing (as IMAGEO-2-test).
- **GeoBench** (Wang et al., 2025b) is a recently released datasets composed of three types images, including 512 normal photos, 512 panoramas and 108 satellite images. The normal photos are sourced from Internet, the panoramas are collected via the Mapillary API, and the satellite images come from Sentinel-2 Level-2A imagery accessed through Microsoft Planetary Computer. We use all the data for testing.

B Experiment Details

B.1 Implementation Details

Our agentic RL training is implemented on VeRL codebase. The specific hyperparameter settings

Config	Setting
<i>RL Training</i>	
optimizer	AdamW
learning rate	1e-6
KL coefficient	0.001
training epoch	2
training batch size	64
PPO mini batch size	16
max response length	4096
max tool response length	1024
max turns	8
group size	16
<i>Parallel Testing</i>	
top K	60
top P	0.95
temperature	1.0

Table 7: Hyperparameters for *Thinking with Map* RL training and parallel testing.

for RL training and parallel testing are shown in Table 7. The RL training and other experiments are conducted on 32 NVIDIA H20 GPUs.

Here we provide the prompt template for *Thinking with Map* and other base models as follows. They all pose a straightforward geolocation task and require the final answer to be returned in a fixed JSON format. The only difference is that the former additionally provides guidance on tool use. The verifier prompt consists of the original geolocation query together with multiple parallel *Thinking with Map* trajectories. The prediction format matches the requirements for the single-agent and base-model setting, that the answer must be in the same fixed JSON format.

B.2 Training Dynamics of RL

To better understand the benefits of agentic RL, we show the reward dynamics over RL training steps in Figure 6. From the reward curve, we find that the training reward increases from 0.25 in the early stage to 0.45 by the end, showing an overall upward trend. This further demonstrates the positive effect of RL on localization accuracy. In the second epoch (i.e., the latter half of training), the reward gradually oscillates and approaches to stable, which suggests that more data may be needed.

B.3 Ablation Study on RL Algorithm

We also try other RL algorithms for *Thinking with Map* agentic training, in particular Pass@K-GRPO (Tang et al., 2025) and PKPO (Walder and Karkhanis, 2025). Results in Table 9 show that although these methods explicitly optimize for pass@K, they perform substantially worse than vanilla GRPO on our task. Therefore, we still use GRPO-trained model for parallel TTS.

B.4 More Ablation Studies on Verifier Models

Here we provide more ablation studies of verifier models on GeoBench and IMAGEO-2-test. As shown in Table 8, unlike the results on MAPBench, using a verifier based on a different base model (e.g., GPT-5) can even outperform the corresponding Best@N (Oracle). This suggests that the verifier is not merely selecting among existing candidates. In few cases, it also identifies more plausible answers along the *Thinking with Map* trajectory.

Prompt Template for Thinking with Map

<image>You are given an image, and your task is to use your exceptional skills to determine the precise coordinates of the location depicted.

Carefully examine the image, taking note of any distinctive features, POIs, landmarks, vegetation, or other elements that could serve as clues.

When extra information is needed to search for a location or confirm precise coordinates, you can use the given tools to get the information from search engine and maps.

Once you have gathered sufficient evidence, provide your best inference for the coordinates in the following JSON format: {"lat": latitude, "lon": longitude, "city": city, "country": country}.

Use signed values for latitude and longitude to indicate N/S and E/W. If you cannot narrow it down, then provide your best guess.

Verifier Model	GeoBench (<i>Acc@Dis</i> , %)						IMAGEO-2-test (<i>Acc@Dis</i> , %)					
	Fine 500m	Local 2km	District 10km	City 25km	Region 200km	Country 750km	Fine 500m	Local 2km	District 10km	City 25km	Region 200km	Country 750km
Verifier@2												
Qwen3-VL-30B-A3B	56.78	66.82	75.47	76.40	79.44	87.38	19.76	21.75	25.19	28.41	45.62	74.25
Qwen3-VL-235B-A22B	55.61	67.06	75.23	76.17	79.44	87.38	19.64	21.86	25.53	29.08	45.06	74.14
GPT-5	60.51	72.90	80.37	81.31	84.11	90.65	21.64	24.20	28.52	31.63	49.06	75.69
Best@2	57.48	69.86	77.34	78.27	80.84	88.79	19.76	22.09	26.42	30.52	48.72	78.36
Verifier@4												
Qwen3-VL-30B-A3B	57.71	69.86	76.64	77.80	81.07	89.02	20.31	22.09	25.97	29.41	45.84	74.14
Qwen3-VL-235B-A22B	57.94	69.16	76.17	77.57	80.84	89.02	20.53	22.64	26.19	30.19	46.06	75.69
GPT-5	63.32	75.00	82.01	83.64	86.45	92.76	22.09	24.64	29.19	33.07	49.39	77.36
Best@4	61.92	73.13	78.50	79.44	82.48	89.95	22.31	24.42	28.52	33.74	53.05	82.24

Table 8: The ablation study of verifier models on GeoBench and IMAGEO. Verifier@N means verifier with N parallel samples.

RL Algorithm	AMAP-test-all (<i>Acc@Dis</i> , %)					
	Fine 500m	Local 2km	District 10km	City 25km	Region 200km	Country 750km
GRPO	19.33	23.36	38.89	43.07	52.29	72.57
Pass@K-GRPO	16.28	19.35	27.52	31.91	40.46	60.48
PKPO	16.97	19.35	26.43	30.15	36.68	50.41

Table 9: The ablation study of RL algorithm.

Prompt Template for Base Model

<image>You are given an image, and your task is to use your exceptional skills to determine the precise coordinates of the location depicted.

Carefully examine the image, taking note of any distinctive features, POIs, landmarks, vegetation, or other elements that could serve as clues.

After showing your thinking, provide your final answer in the JSON format:

```
{"lat": latitude, "lon": longitude, "city": city, "country": country}
```

Use signed values for latitude and longitude to indicate N/S and E/W. If you cannot narrow it down, then provide your best guess.

If candidates disagree, pick the most evidence-consistent and geographically plausible one.

After thinking, provide your final answer in the JSON format:

```
{"lat": latitude, "lon": longitude, "city": city, "country": country}
```

Use signed values for latitude and longitude to indicate N/S and E/W.

Prompt Template for Verifier

You are a strict geo-localization solver. You will be given an image, the original task, and multiple candidate answers from other agents. Synthesize the best final location.