

GAEA: A Geolocation Aware Conversational Assistant

Ron Campos^{*}, Ashmal Vayani^{*}, Parth Parag Kulkarni^{*}, Rohit Gupta,
Aizan Zafar, Aritra Dutta, Mubarak Shah

Center for Research in Computer Vision, University of Central Florida

Abstract

Image geolocalization, in which an AI model traditionally predicts the precise GPS coordinates of an image, is a challenging task with many downstream applications. However, the user cannot utilize the model to further their knowledge beyond the GPS coordinates; the model lacks an understanding of the location and the conversational ability to communicate with the user. In recent days, with the tremendous progress of large multimodal models (LMMs)—proprietary and open-source—researchers have attempted to geolocalize images via LMMs. However, the issues remain unaddressed; beyond general tasks, for more specialized downstream tasks, such as geolocalization, LMMs struggle. In this work, we propose solving this problem by introducing a conversational model, GAEA, that provides information regarding the location of an image as the user requires. No large-scale dataset enabling the training of such a model exists. Thus, we propose GAEA-1.4M, a comprehensive dataset comprising over 800k images and approximately 1.4M question-answer pairs, constructed by leveraging OpenStreetMap (OSM) attributes and geographical context clues. For quantitative evaluation, we propose a diverse benchmark, GAEA-Bench, comprising 3.5k image-text pairs to evaluate conversational capabilities equipped with diverse question types. We consider 11 state-of-the-art open-source and proprietary LMMs and demonstrate that GAEA significantly outperforms the best open-source model, LLaVA-OneVision, by 18.2% and the best proprietary model, GPT-4o, by 7.2%. Our dataset, model and codes are available <https://ucf-crcv.github.io/GAEA>.

1. Introduction

Image geolocalization [25, 34, 52, 58, 63] is a notoriously challenging task, in which, conventional AI models predict the precise GPS coordinate of an image taken anywhere on Earth. Moreover, seasonal changes, geographical and climatic diversity, changes in solar zenith angle, and

^{*}equally contributing first authors







	Location: Bangkok, Thailand (-13.7465, 100.4924) Input Question: Can you name a coffee shop near this image with takeaway services? Ground Truth: Cafe Amazon is a coffee shop with takeaway services, near the location of this image.
	The image you've shared is of a statue at Wat Pho... While Wat Pho is a significant landmark, it's not a coffee shop...
	Unfortunately, I'm unable to provide a specific coffee shop near the image with takeaway services as I don't have access to real-time information.
	For coffee lovers, many local cafes offer takeaway services, such as Gallery Drip Coffee and Cafe Baan .
	While I can identify the image as a statue at Wat Pho... I do not have access to information about coffee shops...
	In Bangkok, Thailand , near the location of this image, you can find Cafe Amazon , which offers both coffee and takeaway options.

Figure 1. We compare the performance of various LMMs on the geographically-grounded visual-question-answering task, included in our new GAEA-Bench benchmark. Most LMMs can describe the Wat Pho statue, but only GAEA, our Geolocation Aware Assistant, retrieves the correct nearby cafe, *Cafe Amazon*.

lack of diverse image distributions make the geolocalization task more challenging. Although challenging, geolocalization has direct applications in multiple domains, including tourism [16], navigation, urban planning [48], and security [22, 52], among others.

Recently, the CLIP-inspired image-to-GPS retrieval approach, GeoCLIP [52], has shown significant performance in global-scale image geolocalization. To further mitigate the performance gap, and to increase the generalization capacity of the models, interestingly, a new wave of works infuse *human-level cognition and inference capacity* in their model training [25, 34, 58]. E.g., PIGEON is trained on data from the popular geolocalization game GeoGuesser [2]; a recent vision-language model, GeoReasoner [34] uses user- and administrator-maintained approximately 3K textual clues from GeoGuesser and Tuxun gaming platforms.

These focused geolocalization models lack a geograph-

ical understanding of the predicted locations beyond their GPS coordinates. They cannot provide additional information that might be invaluable for applications such as tourism, navigation, urban planning, etc. Even if the models possess that understanding, they do not possess the conversational ability to convey that information and fail to meet the user’s needs. In contrast, despite having the conversational capability, visually and textually prompted large language models (LLMs) [21, 50, 59] and their multimodal variants, popularly referred to as large multimodal models (LMMs) [10, 11, 14, 37, 49], fail to capture fine-grained nuances from an image in specialized downstream tasks such as geolocalization, making their predictions vastly imprecise and worse than random guesses in many cases; see Figure 1.

Motivated by these aspects, in this paper, we propose GAEA, an open-source conversational model with global-scale geolocalization capability. To the best of our knowledge, this is the first work in the ground-view geolocalization domain that introduces an open-source conversational chatbot, where the user can obtain image geolocalization, relevant description of the image, and engage in a meaningful conversation about the surrounding landmarks, natural attractions, restaurants or coffee shops, medical or emergency facilities, and recreational areas.

However, training an open-source LMM with conversational capacity is not straightforward. These models are data-hungry, and their training is computationally intensive. Unfortunately, no dataset can facilitate the training of such a model. To this end, we meticulously curate a GAEA-1.4M—a high-quality conversational VQA pair equipped with diversity in scene understanding and image captions for training and instruction tuning the LMMs on the street-level geolocalization task. GAEA-1.4M is a comprehensive dataset consisting over 800k images from *MP-16* [31], *GLD-v2* [57], and *CityGuessr68k* [30] covering locations around the Earth. We augment these images with rich metadata from the OpenStreetMap (OSM) [40] at a 1 km radius, a first effort of its kind. OSM attributes contain details about the surrounding area, nearby landmarks, accessible services, and historical buildup of the region. Additionally, the geolocalizable explanatory captions set contains 385K image-QA pairs and is equipped with their corresponding knowledge and reasoning captions. These knowledge and reasoning captions are constructed using a set of geographical context clues from *GeoGuessr* [2] that enable the model to gain a holistic understanding of the location. Taken together, GAEA-1.4M is the largest and most comprehensive collection of geolocalizable and conversational QA pairs. Using this, we train our conversational chatbot, GAEA.

To quantitatively evaluate the conversational capability of LMMs and address the scarcity of benchmark datasets in a geolocalization setting, we propose GAEA-Bench, a

diverse set of 3.5k conversational question samples.

GAEA-bench comprises multiple-choice (MCQs) and true/false (T/Fs) for checking a model’s understanding and choosing capability, short questions (SVQAs) for testing a model’s knowledge, and long questions (LVQAs) for evaluating a model’s descriptive and in-depth explanation ability about the location in question.

We summarize the main contributions as follows:

- We propose GAEA-1.4M (§Section 3), a new dataset for training conversational image geolocalization models.
- For evaluating conversational capabilities in a geolocalization setting (§Section 5), we propose GAEA-Bench, a novel benchmark of 3.5k samples with various question types, including short, long, MCQs, and T/F.
- We propose GAEA, a conversational chatbot (§4) that extends beyond global-scale geolocalization, providing information about the location described by an image.
- We quantitatively compare the performance of our model to 8 state-of-the-art open-source and 3 proprietary LMMs, including GPT-4o [9] and Gemini-2.0-Flash [49].

2. Related Work

Large Multimodal Models (LMMs) have been at the forefront of computer vision research; geo-localizable LMMs are in their nascent stages. Multimodal learning unifies different modalities by a common representation. By *contrastively* fitting text and images into the same feature space, CLIP [41] has revolutionized multimodal learning. LLMs like GPT2 [42] made strides in representing text and next token prediction. Visual question answering (VQA) was of interest before, but after LLaVA [36] and BLIP2 [33] combined the conversational aspects of LLMs and the representational capabilities of CLIP-like models, many problems of VQA are addressed. After that, numerous modern works emerged, such as GeoChat [29], Qwen2.5-VL [12], LLaMA-3.2 Vision [10], and LLaVA-OneVision [32] as well as proprietary models like GPT4 [9] and Gemini [49]. Although most of these models are excellent for general VQA, they perform poorly on specialized tasks in fields like medicine, statistics, and geolocalization. This inspires the need for specialized LMMs that can address specific tasks.

Geo-localization is a crucial field in vision research with essential applications in forensics, social media, and exploration; see [16, 48, 52]. On a global scale, Weyand et al. [56] first introduced a classification-based approach on the Im2GPS [26] dataset. Vo et al. [53] introduced classification in multiple hierarchies, while CPLaNet [47] introduced a combinatorial partitioning technique for combining coarse hierarchies to predict finer ones. Over the years many other works like ISNs [27], TransGeo [63], TransLocator [54], and GeoDecoder [15] have made significant advancements in this classification-based worldwide geolocalization by in-

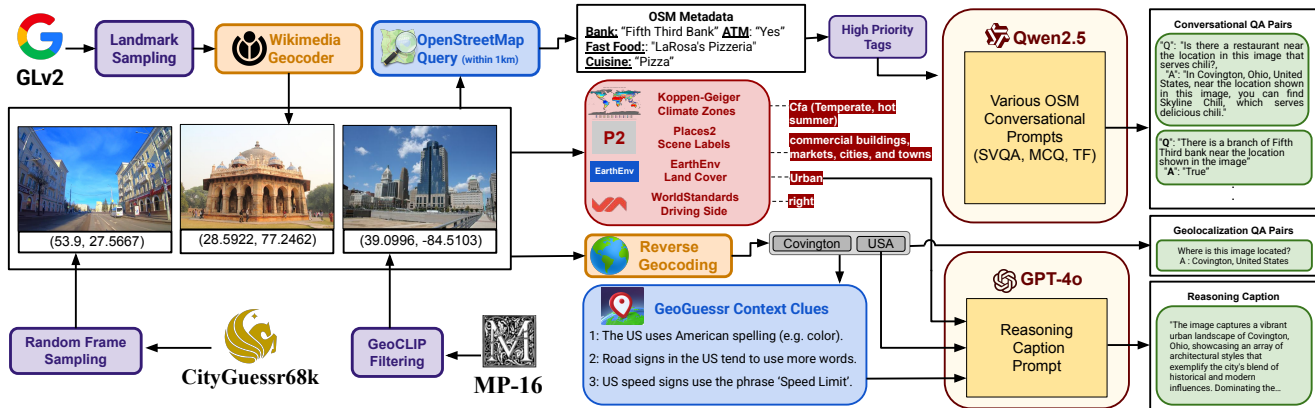


Figure 2. **Data Collection and Annotation Pipeline.** GAEA-1.4M includes geographically diverse visual samples from various data sources, such as MP-16 [31], GLD-v2 [57], and CityGuesser68k [30] (left). We also incorporate additional metadata and auxiliary context for each image from OpenStreetMap (OSM), ranging from climate zones to geographical clues about the country (middle). Using open-source LLMs and GPT-4o, we generate four diverse question-answer pairs across geolocation, reasoning, and conversational subsets (right).

roducing scene-based specialized encoders and hierarchical evaluation, auxiliary scene recognition, and twin encoder approach, and a query-based encoder-decoder architecture, respectively. PIGEON [25], the most recent work, leverages the image representation capabilities of the CLIP vision encoder and a unique clustering approach to improve its geo-localization performance. Image-to-image retrieval models tend to be more accurate than their classification-based counterparts, but are infeasible on a global scale due to their requirement for large reference image galleries. GeoCLIP [52] was the first work to incorporate the contrastive multimodal learning between images and raw GPS information that revolutionized this domain by introducing a more accurate retrieval-based model for a global scale.

These specialized models work well for worldwide image geo-localization but lack the conversational aspect that can aid an individual in gaining a holistic understanding of a location portrayed in an image. GeoReasoner [34] attempts to incorporate an inherent geospatial understanding into a LMM by looking at specific information displayed in the image. It also introduces the idea of *locatability*, which can determine the extent of that information present in the image, which may improve the reasoning capability of the model. The model, however, lacks the *conversational aspect*, and the locatability-based filtering of data might hurt its generalization capability. We address these issues in GAEA by primarily focusing on its conversational ability.

The generalizability of GAEA comes from its training data. All specialized geo-localization models that function on a global scale train their model on MP-16 [31], which is a large-scale worldwide dataset. However, it lacks the verbal context required in LMM training. Hence, we introduce a new conversational dataset GAEA-1.4M; see details in §Section 3. Additionally, we introduce the first conversa-

tional benchmark in §Section 4 to evaluate Geolocation LMMs and an evaluation pipeline to judge the efficacy of such models.

3. Dataset Curation and Annotation

3.1. GAEA-1.4M

The GAEA-1.4M dataset provides comprehensive global coverage, featuring both rich conversational and diverse geolocalization sets. It includes various QA formats, such as MCQs, True/False, and open-ended VQA (long and short), from more than 234 countries/territories, grouped under conversational and geolocalization groups. Spanning 40k cities across 7 continents, GAEA-1.4M is structured into two key groups: conversational and geolocalization. With over 1.4 million QA pairs, it captures the geographical diversity of both underrepresented and widely recognized regions worldwide. Figure 2 shows our complete curation pipeline, which we will discuss piecewise.

Acquiring Diverse Geo-localizable Images. We sample geographically diverse visual data from MediaEval 2016 (MP-16) [31], Google Landmarks v2 (GLDv2) [57], and CityGuesser [30] to curate GAEA-1.4M.

MP-16 contains over 4.6 million geotagged Flickr images, including indoor and outdoor scenes. For our street-view geolocalization subset, we filter out indoor images, retaining 3 million outdoor images. However, some of these images are non-geolocalizable, such as close-up shots of doors, grass, or wires, which are excluded from the final dataset. To filter out non-geolocalizable images, we process all 3 million samples using GeoCLIP [52], which is trained on the full MP-16 dataset and effectively identifies non-geolocalizable outlier images. GeoCLIP assigns a confidence score based on its ability to predict GPS coordi-





Data Source: MP-16	Data Source: GLDv2	Data Source: MP-16	Data Source: GLDv2
<p>Location: Cuenca, Spain (40.074511, -2.134486)</p> <p>OSM:</p> <ol style="list-style-type: none"> 1. <i>Pharmacy:</i> Farmacia Garcia Honduvilla 2. <i>Bank:</i> Santander Bank 3. <i>Fast Food:</i> La que ha liado el pollito 4. <i>Cuisine:</i> Chicken 	<p>Location: Zürich, Switzerland (47.3964, 8.54778)</p> <p>OSM:</p> <ol style="list-style-type: none"> 1. <i>River:</i> Spitalerbach 2. <i>University:</i> Universität Zürich Irchel 3. <i>Bus Route:</i> VBZ route: Bus 69: Zürich, Milchbuck --> ETH Hönggerberg 	<p>Location: New Orleans, United States (29.946846, -90.062592)</p> <p>OSM:</p> <ol style="list-style-type: none"> 1. <i>Restaurant:</i> Felipe's Taqueria 2. <i>Cuisine:</i> Mexican 3. <i>Cafe:</i> Cafe Fleur De Lis 	<p>Location: Rio de Janeiro, Brazil (-22.93836, -43.25942)</p> <p>OSM:</p> <ol style="list-style-type: none"> 1. <i>River:</i> Rio Andaraí 2. <i>Bus Route:</i> STPC Borel 
<p>Question: Is there a fast-food option near the location in this image that serves chicken? If so, can you name it?</p> <p>Answer: In Cuenca, Spain, near the location depicted in this image, you can find La Que Ha Liado El Pollito, which specializes in chicken dishes.</p>	<p>Question: Is there a specific bus route near the location in this image that could help tourists get around?</p> <p>Answer: In Zürich, Kreis 6, Switzerland, the Bus 69 runs from Zürich, Milchbuck to ETH Hönggerberg, providing convenient transportation options for visitors.</p>	<p>Question: Is there a cafe near the location of this image in New Orleans? Choose one option: A) Café Du Monde B) Starbucks Reserve C) Café Fleur De Lis D) Blue Bottle Coffee</p> <p>Answer: C. Café Fleur De Lis</p>	<p>Question: Answer the following question with either True or False: This image is located in São Paulo, Brazil.</p> <p>Answer: False</p>

Figure 3. **Qualitative samples** from our GAEA-1.4M training set showcasing various question-types, including multiple-choice, true/false, short and long VQAs generated using a proprietary model, GPT-4o [39]. We carefully select geographical tags from OSM metadata to generate conversational question-answer pairs.

nates, with higher scores indicating geo-localizability. We set a confidence threshold of 0.75 and computed the distance between the ground truth MP-16 GPS coordinates and the GeoCLIP’s predicted location. We retain the images if this distance is less than 500 km; see additional ablations on different thresholds and distance metrics in the Appendix.

To achieve a balanced geographical distribution in GAEA-1.4M, we use the 10th hierarchy of S2-Cells [5] to partition our filtered MP-16 dataset into 16,753 spatial grid cells. S2-Cells enable hierarchical spatial indexing, ensuring diverse global coverage while preventing overrepresenting densely imaged regions. We randomly sample up to 200 images from each cell, resulting in a final set of over 750k distinct samples.

GLDv2 [57] is a fine-grained landmark recognition dataset featuring natural and human-made landmarks across diverse time zones, climates, and lighting conditions. Given the significance of landmark geolocation for real-world applications, we randomly sample 50K distinct landmarks from GLDv2. These highly recognizable landmarks offer rich geographic and cultural context. Each image is linked to Wikipedia metadata, from which we extract GPS coordinates using the `Wikimedia` [6] API. We then apply the `reverse_geocoder` Python library to determine each landmark’s corresponding city and country.

CityGuessr68k [30] focuses on global video-based geolocation emphasizing urban regions and hierarchical prediction across 166 major cities. To incorporate this diversity, we randomly sample one frame from each of the 54k training videos and include them in our dataset. These three sources provide over 852k geographically diverse geolocalizable images, forming GAEA-1.4M dataset.

3.1.1. Meta-data curation for dataset annotation

After acquiring all visual samples for our GAEA-Conversational Assistant, we churn the metadata for each image for a comprehensive QA-pair generation.

Churning OSM metadata. OpenStreetMap (OSM) [40] is

a collaborative open-source mapping platform that provides extensive geographical data. In our work, OSM plays a central role by enriching geolocalization and conversational capabilities. We retrieve metadata from a 1 km radius around the GPS coordinates of 850k images, leveraging OSM’s detailed, publicly annotated tags. These tags cover many real-world elements, including amenities, transportation, hotels, and restaurants, making them invaluable for our ground-view geolocalization and QA generation.

OSM data is multilingual, which is a key challenge. To ensure accessibility, we use GPT-4o [9] to translate these annotations into English. Additionally, many retrieved tags consisted of plain numbers or non-meaningful entries, which we systematically filtered out to retain only informative and contextually relevant metadata. To our knowledge, this is the first work to utilize OSM’s rich metadata to develop a conversational chatbot for ground-view geolocalization. Figure 3 shows high-priority meta-tags derived from OSM churned data.

Curating Country-Specific Geographical Clues. We web-crawled diverse clues from `Plonkit`[4], an open-source community resource for the `GeoGuessr` [2] game, which has over 65 million players. Similar datasets have been used in recent works [25, 34]. We obtained 129 country clues but found gaps for some countries, such as New Zealand and France. To address this, we curated clues for 58 additional countries using GPT-4o, aligning them with `Plonkit`’s style, resulting in altogether 187 countries. These clues are incorporated into our dataset for generating reasoning-based QAs. For examples of the type of clues utilized, see Figure 9 in the Supplementary material.

Additional Metadata. For auxiliary context, we group our country-specific, geographically diverse dataset in 31 Köppen-Geiger climate zones [13]. We obtain the traffic direction data through `WorldStandards` [7] and `Land Cover Use` statistics from `EarthEnv` [1]. Additionally, we compute scene labels for each image using the `Places2` [62] database.

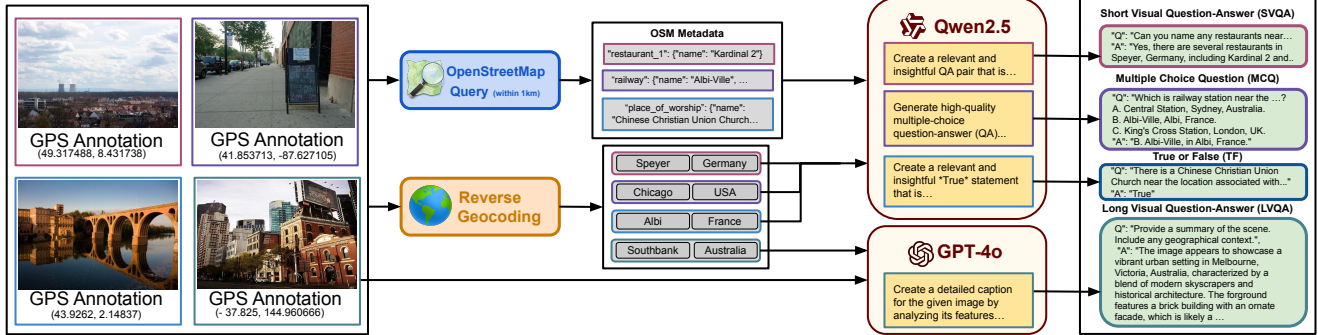


Figure 4. **Overview of GAEA-Bench.** GAEA-Bench is designed to evaluate the conversational abilities of various LMMs across different question types, including MCQs, T/F, and both short and long VQAs. We have carefully selected a subset of 3.5k samples from MP-16 [31] and generated corresponding OSM metadata to generate QA pairs using GPT-4o [39]. GAEA-Bench aims to fill the gap in conversational benchmarks by incorporating geolocalization capabilities.

3.1.2. Question-Answer (QA) Pairs Generation

GAEA-1.4M is carefully curated to enhance ground-view geolocalization through diverse, context-rich QA pairs; see Figure 2. Comprising over 800k distinct images and around 1.4 million QA pairs, it stands as the largest and most comprehensive dataset for this task; see Figure 3. Unlike existing works, such as [20, 34], which are limited to JSON structures and fewer question types, our work emphasizes the conversational capabilities of the model, providing a broader range of QA formats. The dataset is divided into three subsets—Conversational, Reasoning, and Geo-Localization, each designed to capture different aspects of geographic understanding. These subsets feature various question formats, including multiple-choice [38], true/false, and open-ended questions (SVQA and LVQA) [51]. Below, we detail the curation process for each subset.

Conversational QA Generation. We generate conversational QA pairs using OSM metadata from the sampled MP-16 and GLDv2 subsets. We prompt Qwen-2.5-14B [59] with enriched OSM attributes to create diverse question formats, including short-form, multiple-choice, and true/false questions. These OSM tags cover various categories such as amenities, food places, financial institutions, government offices, accommodation, transportation, healthcare, religious sites, education, and waterways. This subset comprises over 380k questions. Figure 3 showcases qualitative samples of QA pairs curated using OSM churned data.

Geolocalization Questions. To enhance the geolocalization capabilities of our GAEA model, we introduce large-scale meta-geographic information through geolocation-specific QA pairs. This subset consists of 820k image-question pairs designed to help the model predict the correct location of an image. We curate 50k geolocation questions from GLDv2, each corresponding to a distinct landmark, leveraging their global recognition to improve location-based reasoning. Additionally, we incorporate 54K geolo-

cation QA pairs from CityGuessr, which focuses on urban environments, and 720k from MP-16, ensuring broad geographic coverage. This results in a diverse and well-distributed geolocation QA dataset spanning 234 countries and territories, 40k cities, and 7 continents.

Reasoning Questions. We generate detailed image-caption QA pairs (Long-VQA) to enhance fine-grained reasoning in our GAEA model. We prompt GPT-4o [9] with each image, its scene labels, and country-specific geographical attributes, including GeoGuessr clues, traffic-side driving information, Köppen-Geiger climate zone, and land cover data. While scene labels are unique to each image, the other attributes provide country-level context. GPT-4o integrates this information to generate contextually rich and highly correlated captions with the provided geographic labels. These reasoning-based captions strengthen the model’s geolocalization and conversational capabilities and induce a rich semantic understanding in our model by infusing *human-level cognition and inference capability*, enabling the model to emphasize why particular image features might be associated with specific geographical contexts, reducing disinformation [44, 45]. In total, we curate 237k knowledge-driven LVQA pairs.

3.2. GAEA-Bench

Existing benchmarks for evaluating geolocalization tasks mainly focus on retrieval and classification-based methods, such as IM2GPS [26], IM2GPS3k [53], and GSW15k [18], which assess the distance between ground-truth and predicted GPS coordinates. However, there is a lack of conversational benchmarking datasets to evaluate the geolocalization and conversational capabilities of LMMs. We introduce GAEA-Bench, a geographically diverse and conversationally rich multimodal benchmark to address these shortcomings. GAEA-Bench is designed to assess LMMs across various question types, including MCQs, true/false, and long and short VQAs while integrating geolocalization

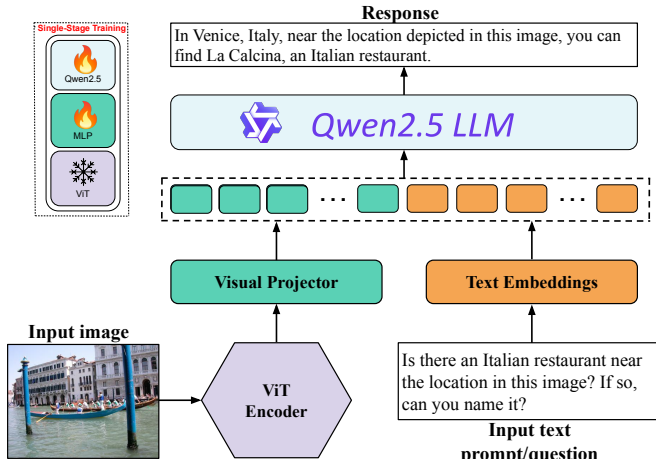


Figure 5. **Overview of the GAEA model architecture and workflow.** An input image is first processed by a Vision Transformer (ViT) encoder, whose output is projected through a visual projector to obtain visual embeddings. Simultaneously, the input text prompt is converted into text embeddings. The combined visual and textual embeddings are then fed into the Qwen2.5 LLM space, which generates a response based on the multimodal input. We follow the single-stage training approach, unfreezing MLP, and performing LoRA fine-tuning in the same stage.

tasks. It includes 3.5k image-text QA pairs that provide a rich geographical context for each image.

GAEA-Bench Curation. We curate a non-overlapping subset of highly geolocalizable MP-16 images, manually filtering out the non-geolocalizable ones. Using OpenStreetMaps (OSM), we generate metadata within a 1km radius and curate 975 short-form (SVQA), 978 multiple-choice (MCQ), and 978 true/false (T/F) questions. For long-form questions (LVQAs), we follow a similar process for generating reasoning questions in GAEA-Bench, resulting in an additional 383 questions. In total, we curate 3,314 diverse image-text QA pairs. To ensure that the GAEA-Bench remains independent of the training set, we select geographically distinct locations for its 3.5k samples. We show our GAEA-Bench annotation and curation process in Figure 4. The OSM metadata are fetched for each image and are passed to Qwen2.5-14B for generating several QA pairs, including SVQA, MCQ, and T/F.

4. Model Architecture

GAEA follows the architecture of the open-source model, Qwen2.5-VL [12], which seamlessly integrates (1) a vision encoder, (2) a vision-to-language projector, and (3) a language model. The re-engineered vision-transformer (ViT) architecture incorporates 2D-RoPE and window attention. The projector is a two-layer multi-layer perception (MLP) to align raw patch features from the ViT and provides the final representation $\mathbf{E}^{\text{Joint}}$ by concatenating the image embeddings, \mathbf{E}^{Img} with the text embeddings, \mathbf{E}^{Text} such that

$$\mathbf{E}^{\text{Joint}} = [\mathbf{E}^{\text{Img}}, \mathbf{E}^{\text{Text}}]; \text{ see Figure 5.}$$

Training Details. We perform single-stage fine-tuning of Qwen2.5VL on our GAEA Conversational Assistant dataset. The model is trained across all three subsets—*geolocalization*, *reasoning*, and *conversational*—covering both open-ended QA formats (short and long answers) and decision-based questions (multiple-choice and true/false). This fine-tuning process enables the model to integrate rich geographical cues, contextual metadata, and image-specific attributes, enhancing its spatial reasoning, location inference, and multimodal conversational capabilities. We employ LoRA fine-tuning [28] with a rank of $r = 16$ and $\alpha = 32$ along with the unfrozen vision-to-language MLP projector. To handle varying image resolutions, we apply dynamic resolution processing: Images below 448×448 are upsampled, while those exceeding 1000×1000 are downsampled, similar to [12]. The model is trained for one epoch over 12,600 steps.

5. Benchmarking and Evaluations

GAEA-1.4M training set comprises four distinct question types: Multiple Choice Questions (MCQs), True/False (T/F), and Short and Long Visual Question Answering (VQA). GAEA is meticulously trained to ensure conversational fluency while possessing the capability to geolocalize visual samples. Current evaluation frameworks primarily focus on standard geo-localization datasets, measuring accuracy using distance-based metrics at various scales, including Street (1 km), City (25 km), Region (200 km), Country (750 km), and Continent (2,500 km). However, these methods fail to assess the conversational capabilities of LMMs. To address this gap, we define our evaluation process in three key dimensions: (a) Conversational accuracy, (b) Quantitative geo-localization accuracy, and (c) Classification accuracy.

5.1. Evaluation and Metrics

Conversational Evaluation. Most geolocation-specific models operate as “black box” systems, providing GPS coordinates without offering any reasoning or justification behind their outputs. In contrast, GAEA is the first model of its kind, explicitly trained on 1.4 million instructions, which include a significant number of knowledge-reasoning question-answer pairs. This enables GAEA to integrate world knowledge, such as geographical clues, conversational meta-tags, and advanced reasoning capabilities, making its geolocation predictions more transparent and insightful. To address the challenges of complex conversational evaluation, we benchmark 12 state-of-the-art open-source and closed-source LMMs on GAEA-Bench, meticulously curated to evaluate LMMs on diverse question types, including multiple-choice, true/false, and open-ended questions (short and long VQAs). See §Section 9.1 for the base-

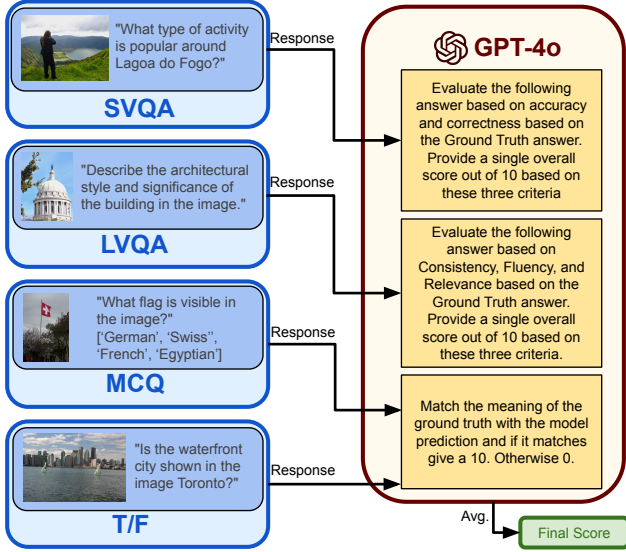


Figure 6. **Evaluation pipeline** for conversational benchmarking on GAEA-Bench, highlighting various question types we introduce in our GAEA-Bench. Each question type is evaluated with various defined criteria using GPT-4o as a judge. For instance, SVQA is evaluated against *accuracy* and *correctness*, and LVQA is evaluated on *Consistency*, *Fluency*, and *Relevancy* criteria.

lines used in this work.

We employ different prompts for each type of question. We use GPT-4o as a judge and prompt it to score responses to various types of questions with different criteria. We use *accuracy* for MCQs and T/F, *correctness* for SVQA, and *consistency*, *relevance*, and *geographical correctness* for long VQAs (LVQAs); see the evaluation pipeline in Figure 6. Here, *correctness* refers to how closely the model’s output matches the location and the correct answer in the ground-truth response [51]. For LVQA, the *consistency* metric evaluates the fluency and readability of the model’s prediction [46, 50, 51], while *geographical correctness* assesses whether the model’s prediction accurately identifies the correct city and country, directly matching the ground-truth answer. This is further discussed in §Section 9.2, and Figure 12.

Quantitative Geo-localization Evaluation. We compared the performance of GAEA against six state-of-the-art (SoTA) geo-localization models, namely PlaNet [56], CPlaNet [47], ISNs [27], TransLocator [54], GeoDecoder [18], and PIGEON [25] on three standard geo-localization benchmarks including IM2GPS [26], IM2GPS3k [53], GWS15k [18]. We prompt various LMMs to output the corresponding city and country to which the image belongs. We retrieve GPS coordinates using GeoPy [3] and compute distance with ground truth. We compare the output with distance thresholds of 1 km, 25 km, 200 km, 750 km, and 2,500 km; see Table 2.

Classification Accuracy. Figure 7 illustrates the *classifi-*

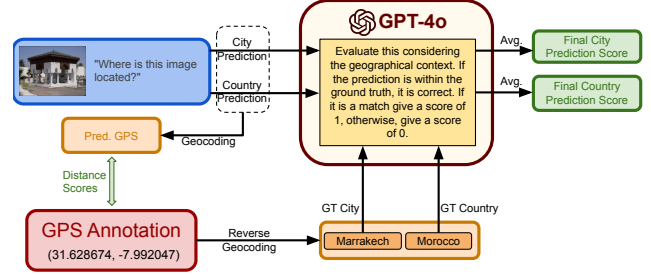


Figure 7. **Classification and distance threshold accuracy computation pipeline** simultaneously evaluates geolocation performance at *city* and *country* level by comparing model predictions with ground truth annotations derived from reverse-geocoding GPS coordinates and accuracy at different distance thresholds by geocoding predictions of the model.

cation accuracy at the city and country levels and *distance* threshold accuracy computation pipeline. For this evaluation, we benchmark on three datasets: GeoDE [43], DollarStreet [23], and CityGuessr68k [30]. From GeoDE, we sampled 22k images based on 16 meta-tags having geolocalizable features. From DollarStreet, we manually sampled 1.3k images, removing indoor and non-geolocalizable samples. Since its metadata contains only country-level information, we evaluate this dataset solely for country classification. Additionally, we use the validation set of 14k images from CityGuessr and all 22k GeoDE samples for city and country classification tasks.

5.2. Results and Discussion

GAEA-Bench Evaluation. Table 1 presents the performance of 12 recent LMMs on GAEA-Bench. The results offer several insights: (i) Our proposed model, GAEA, achieves the highest average performance across decision-making questions (T/F and MCQs) and Short VQAs. Among proprietary models, GPT-4o [9] overall performs the best, with an accuracy of 60.3%, excelling particularly in Long VQAs—outperforming GAEA by 1.9% in this category. However, both open-source and proprietary models struggle with short-form questions. E.g., GPT-4o’s accuracy drops from 66.3% on long questions to 49.5% on short questions. (ii) GAEA outperforms all LMMs with an average accuracy of 67.5%, surpassing GPT-4o by 7.2% and outperforming the second-best open-source model, LLaVA-OneVision [32], by 18.2%. (iii) Several open-source models, including LLaMA-3.2-11B [21], GLM-4V-9B [24], and Phi-3.5-Vision [8], achieve comparable overall performance. (iv) LMMs perform better on decision-making questions (MCQs and T/F) than open-ended questions; see Figure 17. E.g., LLaVA-OneVision experiences a 26.5% drop in accuracy on SVQA compared to MCQ questions. The low performance on free-form questions underscores the challenge of using short questions in effectively assessing conversational capabilities in the GAEA-Bench. Fig-

Model Name	Performance on Different QA Formats				
	LVQA	SVQA	MCQ	TF	Average
GeoChat-7B [29]	23.9 _{-40.5}	16.4 _{-36.2}	54.5 _{-18.2}	32.1 _{-46.4}	33.2 _{-34.3}
LLaVA-Next-Mistral-7B [35]	48.0 _{-16.4}	23.2 _{-29.4}	29.2 _{-43.5}	56.7 _{-21.8}	37.7 _{-29.8}
Phi-3.5-Vision-Instruct [8]	48.7 _{-15.7}	14.3 _{-38.3}	54.7 _{-18.0}	57.4 _{-21.1}	42.9 _{-24.6}
LLaMA-3.2-Vision-11B [10]	48.6 _{-15.8}	29.9 _{-22.7}	53.2 _{-19.5}	47.1 _{-31.4}	44.0 _{-23.5}
GLM-4V-9B [24]	41.6 _{-22.8}	29.7 _{-22.9}	56.7 _{-16.0}	50.6 _{-27.9}	45.2 _{-22.3}
Qwen2.5-VL [12]	57.0 _{-7.4}	29.9 _{-22.7}	48.1 _{-24.6}	59.4 _{-19.1}	47.1 _{-20.4}
InternVL2-8B [60]	54.5 _{-9.9}	31.0 _{-21.6}	55.7 _{-17.0}	56.8 _{-21.7}	48.6 _{-18.9}
LLaVA-OV-7B [32]	54.1 _{-10.3}	31.5 _{-21.1}	58.0 _{-14.7}	56.4 _{-22.1}	49.3 _{-18.2}
Gemini-2.0-Flash [49]	58.3 _{-6.1}	34.7 _{-17.9}	57.2 _{-15.5}	56.7 _{-21.8}	50.5 _{-17.0}
GPT-4o-mini [9] *	61.8 _{-2.6}	34.1 _{-18.5}	54.9 _{-17.8}	33.9 _{-44.6}	43.4 _{-24.1}
GPT-4o [9] *	66.3 _{+1.9}	49.5 _{-3.1}	59.4 _{-13.3}	69.6 _{-8.9}	60.3 _{-7.2}
GAEA (Ours)	64.4	52.6	72.7	78.5	67.5

Table 1. We benchmark 11 open-source and proprietary LMMs on GAEA-Bench. Notably, GAEA outperforms all open-source models and fares higher than the proprietary models on decision-making questions (*MCQs* and *TFs*). We provide the relative performance change for each model compared to GAEA. * - We use GPT-4o as a judge for evaluation, and it has been documented that LLMs as judges prefer their long-form output [55, 61], hence the scores for these models are likely overestimated.




Figure 8. Classification accuracy for both city and country labels, where GAEA surpasses several recent LMMs in performance.

ures 14 to 16 shows comparisons with several LMMs.

Standard Geo-localization Evaluation. Table 2 compares GAEA’s performance with various specialized encoder-only methods across three standard geolocation benchmarks. While GAEA is trained on a large-scale conversational dataset with geolocation capabilities, it achieves competitive results against specialized models. We evaluate against GaGA [20], which is trained on a dataset five times larger than ours, on IM2GPS and IM2GPS3k. However, we exclude comparisons on GWS15k due to differences in dataset curation. We contacted the authors of [18] for the original GWS15k benchmark for fair evaluation.

On the IM2GPS3k benchmark, GAEA consistently outperforms all existing LMMs, including the domain-specialized PIGEON model [25], across all four distance thresholds. Notably, GAEA surpasses PIGEON by 2.2% at the 200 km threshold and by approximately 1.4% at the continent level. Compared to GeoCLIP [52], GAEA achieves gains of 5.3% at the region level and 2.5% at the city level. Additionally, GAEA significantly outperforms the open-source models, GaGA [20], with improvements of 6.1% at the country level and 8.0% at the regional level and GeoReasoner [34] with 10.5% at city, 15.6% at region, and 13.9% at a 2500 km threshold. On IM2GPS, GAEA

Benchmark	Model	City	Region	Country	Continent
		25 km	200 km	750 km	2500 km
IM2GPS [26]	PlaNet [56]	24.5	37.6	53.6	71.3
	CPlaNet [47]	37.1	46.4	62.0	78.5
	ISNs [27]	43.0	51.9	66.7	80.2
	TransLocator [54]	48.1	64.6	75.6	86.7
	GeoCLIP [52]	41.8	60.8	77.2	89.9
	GeoDecoder [18]	50.2	69.0	80.0	89.1
	PIGEON [25]	40.9	63.3	82.3	91.1
	GeoReasoner [34]	24.9	48.1	65.8	82.3
	GaGA [20]	38.8	54.8	75.1	87.7
	GAEA (Ours)	43.0	57.4	77.2	89.5
IM2GPS3k [53]	PlaNet [56]	24.8	34.3	48.4	64.6
	CPlaNet [47]	26.5	34.6	48.6	64.6
	ISNs [27]	28.0	36.6	49.7	66.0
	TransLocator [54]	31.1	46.7	58.9	80.1
	GeoDecoder [18]	33.5	45.9	61.0	76.1
	GeoCLIP [52]	34.5	50.7	69.7	83.8
	PIGEON [25]	36.7	53.8	72.4	85.3
	GeoReasoner [34]	26.5	40.4	57.7	72.8
	GaGA [20]	33.0	48.0	67.1	82.1
	GAEA (Ours)	36.9	56.0	73.2	86.7
GWS15k [18]	ISNs [27]	0.6	4.2	15.5	38.5
	TransLocator [54]	1.1	8.0	25.5	48.3
	GeoDecoder [18]	1.5	8.7	26.9	50.5
	GeoCLIP [52]	3.1	16.9	45.7	74.1
	GAEA (Ours)	3.7	16.7	43.3	73.5

Table 2. We benchmark the performance of various specialized models on standard geolocation datasets. GAEA demonstrates competitive results, outperforming GaGA on multiple distance thresholds in both IM2GPS and IM2GPS3k.

outperforms both the specialized LMMs, GeoReasoner and GaGA on all four distance metrics. It also surpasses PIGEON and GeoCLIP models at the city level by 2.1% and 1.2%, respectively, while maintaining competitive performance across other thresholds. We also evaluate GAEA on GWS15k, one of the most challenging datasets, which includes non-geolocalizable landmarks. GAEA outperforms GeoCLIP [52] and GeoDecoder [18] on city-level distance and achieves comparable performance at the region and country levels. Figure 8 presents GAEA’s **Classification Accuracy** on three new datasets: CityGuessr68k-val [30], GeoDE [43], and DollarStreet [23]. GAEA outperforms recent LMMs, including LLaVA-OneVision [32], InternVL [17], and GLM-4V-9B [24], on both city- and country-level classification. These results highlight GAEA’s extensive geographical coverage and strong geolocation capabilities.

6. Conclusion

We introduced GAEA, the first interactive conversational model with specialized geolocation capabilities, explicitly trained on a large-scale conversational dataset, GAEA-1.4M. We meticulously designed the dataset to enhance GAEA’s reasoning, conversational abilities, and geolocation accuracy. We curated geolocalizable images from MP-16, GLDv2, and CityGuessr68k, enriching them with auxiliary context and metadata, such as geographic clues and climate zones. In addition to a high-quality instruction set, we present GAEA-Bench, a comprehensive benchmark that evaluates LMMs across multiple question types, including MCQs, True/False, short- and long-VQAs. Our results

show that GAEA outperforms recent LMMs on GAEA-Bench, demonstrating strong geolocation and conversational capabilities by leveraging OpenStreetMap (OSM) data. These findings establish GAEA as a strong baseline for future research in geolocalization.

7. Acknowledgement

This work was supported by “MFC Lockheed Martin, Orlando”. We would also like to thank David Shatwell, Manu S Pillai, Praveen Tirupattur, Brian Dina, and Suranadi Dodampaganmage for their insightful discussions and contributions.

References

- [1] . EarthEnv. <https://www.worldstandards.eu/cars/list-of-left-driving-countries/>, . 4, 12
- [2] . GeoGuessr. <https://www.geoguessr.com/>, . 1, 2, 4, 12
- [3] . GeoPy. <https://geopy.readthedocs.io/en/stable/>, . 7
- [4] . Plonkit. <https://www.plonkit.net/>, . 4, 12
- [5] . S2-Cells. <https://code.google.com/archive/p/s2-geometry-library/>, . 4
- [6] . WikiMedia. <https://commons.wikimedia.org/w/api.php>, . 4
- [7] . WorldStandards. <https://www.worldstandards.eu/cars/list-of-left-driving-countries/>, . 4, 12
- [8] Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, et al. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219*, 2024. 7, 8, 14
- [9] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. 2, 4, 5, 7, 8, 14
- [10] Meta AI. Llama 3.2: Vision and edge models. Meta AI Blog, 2024. 2, 8, 14
- [11] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. In *Advances in neural information processing systems*, pages 23716–23736, 2022. 2
- [12] 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. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025. 2, 6, 8, 14
- [13] Hylke E Beck, Niklaus E Zimmermann, Tim R McVicar, Noemi Vergopolan, Alexis Berg, and Eric F Wood. Present and future köppen-geiger climate classification maps at 1-km resolution. *Scientific data*, 5(1):1–12, 2018. 4, 12
- [14] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. In *Advances in neural information processing systems*, pages 1877–1901, 2020. 2
- [15] Denis Carriere. Geocoder: Simple, consistent. <https://geocoder.readthedocs.io/>. Accessed: [Insert Date]. 2
- [16] Athanasios Chalvatzaras, Ioannis Pratikakis, and Angelos A Amanatiadis. A survey on map-based localization techniques for autonomous vehicles. *IEEE Transactions on intelligent vehicles*, 8(2):1574–1596, 2022. 1, 2
- [17] Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 24185–24198, 2024. 8, 14
- [18] Brandon Clark, Alec Kerrigan, Parth Parag Kulkarni, Vicente Vivanco Cepeda, and Mubarak Shah. Where we are and what we’re looking at: Query based worldwide image geo-localization using hierarchies and scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23182–23190, 2023. 5, 7, 8, 14
- [19] Tri Dao, Dan Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. Flashattention: Fast and memory-efficient exact attention with io-awareness. *Advances in neural information processing systems*, 35:16344–16359, 2022. 15
- [20] Zhiyang Dou, Zipeng Wang, Xumeng Han, Chenhui Qiang, Kuiran Wang, Guorong Li, Zhibei Huang, and Zhenjun Han. Gaga: Towards interactive global geolocation assistant. *arXiv preprint arXiv:2412.08907*, 2024. 5, 8, 14
- [21] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The Llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024. 2, 7
- [22] Aritra Dutta, Srijan Das, Jacob Nielsen, Rajat Subhra Chakraborty, and Mubarak Shah. Multiview aerial visual recognition (mavrec): Can multi-view improve aerial visual perception? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22678–22690, 2024. 1
- [23] William Gaviria Rojas, Sudnya Diamos, Keertan Kini, David Kanter, Vijay Janapa Reddi, and Cody Coleman. The dollar street dataset: Images representing the geographic and socioeconomic diversity of the world. *Advances in Neural Information Processing Systems*, 35:12979–12990, 2022. 7, 8, 14
- [24] Team GLM, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Diego Rojas, Guanyu Feng, Hanlin Zhao, Hanyu Lai, Hao Yu, Hongning Wang, Jiadai Sun, Jiajie Zhang, Jiale Cheng, Jiayi Gui, Jie Tang, Jing Zhang, Juanzi Li, Lei Zhao, Lindong Wu, Lucen Zhong, Mingdao Liu, Minlie Huang, Peng Zhang, Qinkai Zheng, Rui Lu, Shuaiqi Duan, Shudan Zhang, Shulin Cao, Shuxun Yang, Weng Lam Tam, Wenyi Zhao, Xiao Liu, Xiao Xia, Xiaohan Zhang, Xiaotao Gu, Xin Lv, Xinghan Liu, Xinyi Liu, Xinyue Yang,

- Xixuan Song, Xunkai Zhang, Yifan An, Yifan Xu, Yilin Niu, Yuantao Yang, Yueyan Li, Yushi Bai, Yuxiao Dong, Zehan Qi, Zhaoyu Wang, Zhen Yang, Zhengxiao Du, Zhenyu Hou, and Zihan Wang. Chatglm: A family of large language models from glm-130b to glm-4 all tools, 2024. 7, 8, 14
- [25] Lukas Haas, Michal Skreta, Silas Alberti, and Chelsea Finn. Pigeon: Predicting image geolocations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12893–12902, 2024. 1, 3, 4, 7, 8, 14
- [26] James Hays and Alexei A Efros. Im2gps: estimating geographic information from a single image. In *2008 IEEE conference on computer vision and pattern recognition*, pages 1–8. IEEE, 2008. 2, 5, 7, 8, 14
- [27] Zheng Hong, Yiwei Yin, Zhe Luo, and Jiebo Luo. Insn: Image-specific neural style transfer for image geolocation. *arXiv preprint arXiv:2106.11593*, 2021. 2, 7, 8, 14
- [28] Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations*. 6, 14
- [29] Kartik Kuckreja, Muhammad S. Danish, Muzammal Naseer, Abhijit Das, Salman Khan, and Fahad S. Khan. Geochat: Grounded large vision-language model for remote sensing. *The IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. 2, 8, 14
- [30] Parth Parag Kulkarni, Gaurav Kumar Nayak, and Mubarak Shah. Cityguesser: City-level video geo-localization on a global scale. In *European Conference on Computer Vision*, pages 293–311. Springer, 2024. 2, 3, 4, 7, 8, 14
- [31] Martha Larson, Mohammad Soleymani, Guillaume Gravier, Bogdan Ionescu, and Gareth JF Jones. The benchmarking initiative for multimedia evaluation: Mediaeval 2016. *IEEE MultiMedia*, 24(1):93–96, 2017. 2, 3, 5
- [32] Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. Llava-onevision: Easy visual task transfer, 2024. 2, 7, 8, 14
- [33] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023. 2
- [34] Ling Li, Yu Ye, Bingchuan Jiang, and Wei Zeng. GeoReasoner: Geo-localization with reasoning in street views using a large vision-language model. In *Proceedings of the 41st International Conference on Machine Learning*, pages 29222–29233, 2024. 1, 3, 4, 5, 8, 14
- [35] Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. LLaVA-NeXT: Improved reasoning, OCR, and world knowledge, 2024. 8, 14
- [36] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36, 2024. 2
- [37] Muhammad Maaz, Hanoona Rasheed, Abdelrahman Shaker, Salman Khan, Hisham Cholakkal, Rao M Anwer, Tim Baldwin, Michael Felsberg, and Fahad S Khan. Palo: A polyglot large multimodal model for 5b people. *arXiv preprint arXiv:2402.14818*, 2024. 2
- [38] Vishal Narnaware, Ashmal Vayani, Rohit Gupta, Swetha Sirnam, and Mubarak Shah. Sb-bench: Stereotype bias benchmark for large multimodal models. *arXiv preprint arXiv:2502.08779*, 2025. 5
- [39] OpenAI. Gpt-4o mini: Our affordable and intelligent small model for fast, lightweight tasks, 2024. Available at: <https://platform.openai.com/docs/models/o1>. 4, 5
- [40] OpenStreetMap contributors. Openstreetmap, 2024. [Data set]. OpenStreetMap Foundation. Available as open data under the Open Data Commons Open Database License (ODbL). 2, 4, 12
- [41] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. 2
- [42] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9, 2019. 2
- [43] Vikram V Ramaswamy, Sing Yu Lin, Dora Zhao, Aaron Adcock, Laurens van der Maaten, Deepti Ghadiyaram, and Olga Russakovsky. Geode: a geographically diverse evaluation dataset for object recognition. *Advances in Neural Information Processing Systems*, 36:66127–66137, 2023. 7, 8, 14
- [44] Shaina Raza, Rizwan Qureshi, Anam Zahid, Joseph Fiorese, Ferhat Sadak, Muhammad Saeed, Ranjan Sapkota, Aditya Jain, Anas Zafar, Muneeb Ul Hassan, et al. Who is responsible? the data, models, users or regulations? responsible generative ai for a sustainable future. *arXiv preprint arXiv:2502.08650*, 2025. 5
- [45] Shaina Raza, Ashmal Vayani, Aditya Jain, Aravind Narayanan, Vahid Reza Khazaie, Syed Raza Bashir, Elham Dolatabadi, Gias Uddin, Christos Emmanouilidis, Rizwan Qureshi, et al. Vldbench: Vision language models disinformation detection benchmark. *arXiv preprint arXiv:2502.11361*, 2025. 5
- [46] Ananya B Sai, Akash Kumar Mohankumar, and Mitesh M Khapra. A survey of evaluation metrics used for nlg systems. *ACM Computing Surveys (CSUR)*, 55(2):1–39, 2022. 7
- [47] Paul Hongsuck Seo, Tobias Weyand, Jack Sim, and Bohyung Han. Cplanet: Enhancing image geolocalization by combinatorial partitioning of maps. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 536–551, 2018. 2, 7, 8, 14
- [48] Qiaomu Shen, Wei Zeng, Yu Ye, Stefan Müller Arisona, Simon Schubiger, Remo Burkhard, and Huamin Qu. Streetvizor: Visual exploration of human-scale urban forms based on street views. *IEEE transactions on visualization and computer graphics*, 24(1):1004–1013, 2017. 1, 2
- [49] Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023. 2, 8, 14
- [50] Omkar Thawakar, Ashmal Vayani, Salman Khan, Hisham Cholakkal, Rao M Anwer, Michael Felsberg, Tim Baldwin,

- Eric P Xing, and Fahad Shahbaz Khan. Mobillama: Towards accurate and lightweight fully transparent GPT. *arXiv preprint arXiv:2402.16840*, 2024. [2](#), [7](#)
- [51] Ashmal Vayani, Dinura Dissanayake, Hasindri Watawana, Noor Ahsan, Nevasini Sasikumar, Omkar Thawakar, Henok Biadglign Ademtew, Yahya Hmaiti, Amandeep Kumar, Kartik Kuckreja, et al. All languages matter: Evaluating llms on culturally diverse 100 languages. *arXiv preprint arXiv:2411.16508*, 2024. [5](#), [7](#)
- [52] Vicente Vivanco Cepeda, Gaurav Kumar Nayak, and Mubarak Shah. Geoclip: Clip-inspired alignment between locations and images for effective worldwide geolocalization. *Advances in Neural Information Processing Systems*, 36, 2024. [1](#), [2](#), [3](#), [8](#)
- [53] Nam Vo, Nathan Jacobs, and James Hays. Revisiting im2gps in the deep learning era. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2621–2630, 2017. [2](#), [5](#), [7](#), [8](#), [14](#)
- [54] Bingxian Wang, Ying Chen, Xin Zhang, Haojie Wang, Ziqiang Wang, and Wen Xu. Translocator: A transformer-based large-scale image geolocalization approach. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2558–2566, 2022. [2](#), [7](#), [8](#), [14](#)
- [55] Koki Wataoka, Tsubasa Takahashi, and Ryokan Ri. Self-preference bias in LLM-as-a-judge, 2025. [8](#)
- [56] Tobias Weyand, Ilya Kostrikov, and James Philbin. Planet-photo geolocation with convolutional neural networks. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VIII 14*, pages 37–55. Springer, 2016. [2](#), [7](#), [8](#), [14](#)
- [57] Tobias Weyand, Andre Araujo, Bingyi Cao, and Jack Sim. 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, 2020. [2](#), [3](#), [4](#)
- [58] Yibo Yan and Joey Lee. Georeasoner: Reasoning on geospatially grounded context for natural language understanding. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, pages 4163–4167, 2024. [1](#)
- [59] An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*, 2024. [2](#), [5](#)
- [60] Jiabo Ye, Haiyang Xu, Haowei Liu, Anwen Hu, Ming Yan, Qi Qian, Ji Zhang, Fei Huang, and Jingren Zhou. mplug-owl3: Towards long image-sequence understanding in multi-modal large language models, 2024. [8](#)
- [61] Jiayi Ye, Yanbo Wang, Yue Huang, Dongping Chen, Qihui Zhang, Nuno Moniz, Tian Gao, Werner Geyer, Chao Huang, Pin-Yu Chen, Nitesh V Chawla, and Xiangliang Zhang. Justice or prejudice? quantifying biases in LLM-as-a-judge. In *The Thirteenth International Conference on Learning Representations*, 2025. [8](#)
- [62] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence*, 40(6):1452–1464, 2017. [4](#), [12](#)
- [63] Sijie Zhu, Mubarak Shah, and Chen Chen. Transgeo: Transformer is all you need for cross-view image geo-localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1162–1171, 2022. [1](#), [2](#)

GAEA: A Geolocation Aware Conversational Assistant

Supplementary Material

We organize the Supplementary Material as follows: In Section 8, we provide additional details of our dataset, GAEA-1.4M. In Section 9, we summarize baselines we compared against, describe prompts used during training, inference and evaluation, and provide training details and additional results on GAEA.

8. Addendum to the Dataset

In this section, we present the dataset statistics and challenges encountered in its creation. Additionally, we discuss our plans to address these limitations in future works.

8.1. GeoGuessr Context Clues

To support geographically-grounded reasoning in our dataset, we incorporate country-specific context clues inspired by the popular game GeoGuessr [2]. We web-crawled and extracted 129 high-quality clues from Plonkit [4], a community-driven open-source resource widely used by over 65 million GeoGuessr players. To improve country coverage and address missing entries (e.g., France, New Zealand), we augmented the set with 58 additional GPT-4o-generated clues, carefully aligned with Plonkit’s descriptive style. This expanded our coverage to 187 countries. These clues highlight distinctive, often visual, geographical features, such as Bangladesh’s unique turquoise license plates, Australia’s widespread gum trees, or the ice-covered terrain of Grise Fiord in Canada (see Figure 9). These examples are leveraged to generate reasoning-based multiple-choice and open-ended questions. Furthermore, we enrich each sample with auxiliary metadata: Köppen-Geiger climate classification [13], traffic orientation from WorldStandards [7], land cover data from EarthEnv [1], and scene labels derived from the Places2 database [62].

8.2. Challenges with Open Street Maps (OSM)

OpenStreetMaps (OSM) [40] is a rich data source for geospatial applications. It contains a wide variety of geographic and infrastructure-related information. Using such a vast open-source dataset, we can collect data about stationary objects in the world, including infrastructure, topological information, various types of amenities (e.g., schools, hospitals, restaurants), transportation networks, international country boundaries, historical and cultural sides, and natural features (e.g., forests, rivers, and seas). Each feature from the OSM dataset has several associated features, such as names and physical characteristics. In GAEA-1.4M, we geocode the visual sample with its GPS coordinates and use the location information (longitude and

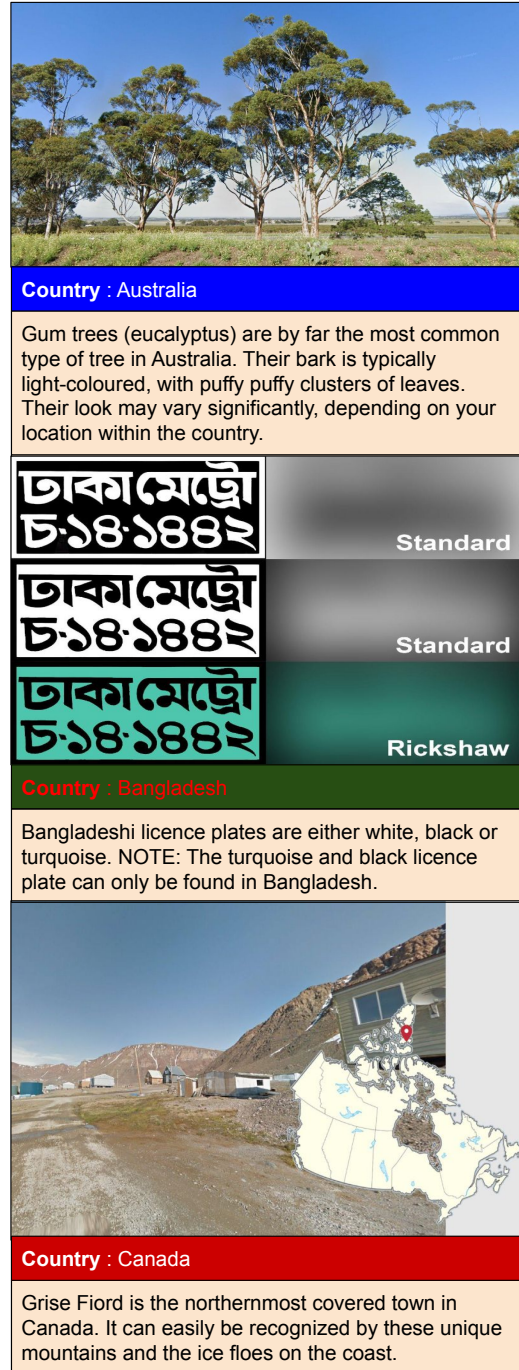


Figure 9. Examples of country-specific geographical clues used in our dataset. These clues are used to support geographically grounded reasoning in our QA generation.






Question Type: Geolocalization	Question Type: SVQA	Question Type: MCQ	Question Type: TF
			
<p>Question: Where is this image taken? Respond with only the city and country.</p> <p>Answer: Hanoi, Vietnam</p>	<p>Question: Is there a notable river near the location in this image?</p> <p>Answer: Yes, near the location in this image, you can find the Barranco de la Cueva de Azañón in Trillo, Spain.</p>	<p>Question: Is there a restaurant near the location of this image that serves sandwiches? Choose one option: A. Tandoori Flame B. Nelly's C. Schnitzel Haus D. La Baguette Gourmande</p> <p>Answer: B. Nelly's</p>	<p>Question: This restaurant named Velkopopovicka Pub is located near the location of this image.</p> <p>Answer: True</p>
Question Type: LVQA			
	<p>Question: Provide a summary of the scene.</p> <p>Answer: Nestled in the breathtaking landscape of Soelden, Austria, this image captures a charming alpine lodge that is a quintessential representation of the region's architecture and natural beauty. The building showcases a robust stone facade complemented by wooden accents, embodying the traditional alpine style with its dark-colored roof and well-placed balconies. Set against the dramatic backdrop of towering mountains, which are characteristic of the region's polar tundra climate, this lodge stands as a testament to the resilience and adaptability of alpine architecture. The area is known for its stunning vistas and rich outdoor activities, making it a popular destination for both winter sports and summer hiking. The well-maintained gardens and outdoor seating areas, adorned with bright yellow flags, invite visitors to relax and enjoy the fresh mountain air, while the nearby rugged terrain hints at adventure waiting to be explored. Soelden, located in the Ötztal Alps, is famed for its ski resorts and vibrant tourism culture, attracting enthusiasts from around the world. The driving side in Austria is right, reflecting a well-organized infrastructure that supports the influx of visitors. As you take in the scene, you can almost feel the warmth of the local hospitality that characterizes Austrian culture, making this lodge not just a place to stay, but a welcoming gateway to the breathtaking wonders of the alpine landscape.</p>		

Figure 10. Examples of the four question types in our dataset: SVQA, MCQ, TF, and LVQA. Each type targets a distinct reasoning skill grounded in geographical, visual, or contextual understanding. Our dataset has three categories, including Geolocalization, Reasoning (LVQA), and Conversational (SVQA, MCQ, TF) QAs, as shown in the figure.

latitude) as a query to the OSM database to fetch geospatial information within a 1 KM radius and further utilize that information to generate question-answer pairs for the training of GAEA.

Despite being such a rich source of data, OSM faces several challenges. One major issue is the variability in data quality and completeness, as contributions to OSM are made by the open-source community, which may result in inconsistent information across different regions. Urban areas often have much more detailed information than rural areas, leading to less comprehensive annotations for rural regions. Another inconsistency related to human annotations stems from the different representations of the same label in different areas, introducing inherent heterogeneity in the structure of OSM data. For instance, some users might label a path as a “trail,” while others might call it a “footway,” and distinctions between what counts as a “park” versus a “garden” are not always clear. Moreover, querying and retrieving data from OSM is a compute-intensive task. It often becomes slower as the number of queries increases and struggles to handle dense or redundant information, necessitating efficient filtering and optimization techniques. Lastly, the information is not always up-to-date, as volunteers update different areas at different times. While some locations may have very recent data, others may be outdated, and sometimes different parts of the same area may contain information from varying periods.

8.3. Training Examples and Data Statistics

GAEA-1.4M covers 234 different countries and territories, and 41,481 cities. Table 3 shows the statistics of GAEA-1.4M in detail. Figure 10 shows some qualitative examples of various question types in our GAEA-1.4M leveraging OSM, and GeoGuessr geographical clues for constructing conversational QA pairs.

Total Images	822,951
Total Cities / Countries	41,481 / 234
Total Questions	1,432,519
Total Geo-Loc QAs	822,951
Total LVQAs	236,935
Total SVQAs	267,668
Total MCQs	48,673
Total True/False QAs	56,292

Table 3. Dataset Statistics

9. Addendum to Baseline and Evaluation

This section covers the models used for comparison with GAEA, the prompts used during training and inference, the prompts used for evaluating GAEA-Bench, and the training hyperparameters.

Geolocalization Prompt
As a geography and tourism expert, analyze the image to determine its exact location. Utilize your extensive knowledge of geography, terrain, landscapes, flora, fauna, infrastructure, and recognizable landmarks to identify the city and country where the image was taken. Question:
LVQA Prompt
Drawing upon your expertise in geography and tourism, examine the image and provide a comprehensive description of the community or lifestyle depicted. Include insights about cultural practices, geographic features, terrain, local flora and fauna, infrastructure, and any natural or man-made elements that characterize the location. Consider how these factors influence the lifestyle and community in the area. Question:
SVQA Prompt
Provide a short answer on notable landmarks, museums, parks, restaurants, or activities that visitors might enjoy in the area. Highlight amenities and services that enhance the tourism experience at this location. Question:
MCQ Prompt
Use your comprehensive knowledge of geography, landmarks, and tourism to analyze the image and determine the correct answer from the options provided. Note, your final answer should be a choice of either A, B, C, or D, including both the letter and the complete text of the option exactly as presented. Question:
TF Prompt
Use your comprehensive knowledge of notable landmarks, museums, parks, restaurants, and related attractions to evaluate the following statement. Provide your final answer as either 'True' or 'False'. Question:

Figure 11. Task-specific training prompts used to instruct the model for each of the five question types in GAEA. Each prompt is carefully designed to elicit a targeted form of geographical reasoning. These prompts ensure consistent and interpretable outputs during both the training and evaluation phases while training GAEA.

9.1. Baselines

We benchmark 8 top-performing open-source LMMs, including LLaMA 3.2-Vision [10], InternVL2 [17], Qwen2.5-VL [12], Phi3.5-vision-instruct [8], GeoChat [29], LLaVA-OneVision [32], GLM-4V-9B [24], LLaVA-NeXT-Mistral-7B [35], and 3 proprietary models, Open-AI’s GPT-4o, GPT-4o-mini [9], and Google’s Gemini-2.0-Flash [49] on GAEA-Bench.

Additionally, we compared the performance of GAEA against six state-of-the-art (SoTA) specialized geolocalization models, namely PlaNet [56], CPlaNet [47], ISNs [27], TransLocator [54], GeoDecoder [18], and PIGEON [25] and open-source geolocalization LMMs GeoReasoner [34] and GaGA [20] on three standard geo-localization benchmarks including IM2GPS [26], IM2GPS3k [53], and GWS15k [18]. We also compare our city-country classification performance with other LMMs on 3 benchmarks CityGuessr68k [30], GeoDE [43] and DollarStreet [23]. Preprocessing for these benchmarks is described in Section 5.1 in the main paper.

9.2. Prompts Used During Training and Inference

When training GAEA, we employed the task-specific prompts shown in Figure 11 to align the model’s understanding with target objectives. During inference, these identical prompts were used for all models evaluated on GAEA-Bench to ensure comparability.

9.3. Prompts Used in Evaluation

Figure 12 presents the task-specific prompts used for evaluating model-generated answers via GPT-based assess-

Number of epochs	1
Global batch size	128
Gradient accumulation steps	4
Initial learning rate	10^{-5}
Learning rate scheduler	cosine
Warmup ratio	0.03
LoRA rank (r)	16
LoRA alpha (α)	32
LoRA dropout	0.01
Precision	bfloat16

Table 4. Hyperparameters used for training GAEA.

ment. For SVQA and LVQA, the prompts emphasize multi-faceted criteria such as location accuracy, specificity, cultural relevance, and reasoning quality, encouraging a nuanced evaluation of open-ended responses. In contrast, MCQ and TF tasks are scored based on strict binary accuracy by matching the predicted answer to the ground truth. These structured prompts ensure consistent, interpretable, and criterion-aligned evaluation across different QA types.

9.4. Training Hyperparameters

We perform single-stage training on the baseline, Qwen2.5-VL [12] using GAEA-1.4M. The training is conducted for 1 epoch with a global batch size of 128, accumulating gradients every 4 steps. The initial learning rate is set to 10^{-5} , using a cosine learning rate scheduler to provide a smooth decay in the learning rate. The warmup ratio is configured at 0.03. We performed LoRA[28]-finetuning with a rank, $r = 16$, $\alpha = 32$, and a dropout rate of 0.01. The model op-

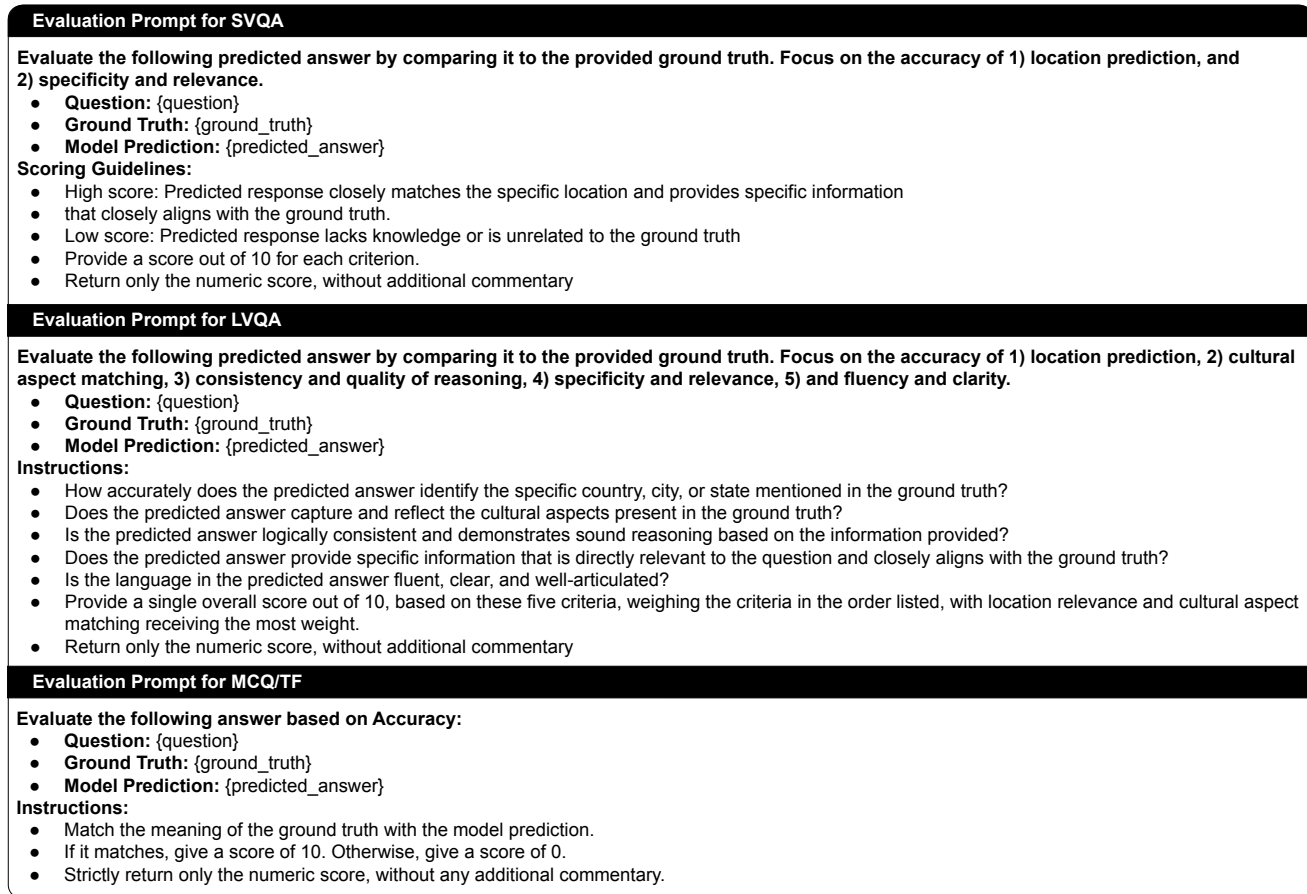


Figure 12. Evaluation prompts used for GPT-based scoring of model predictions across SVQA, LVQA, MCQ, and TF question types in the GAEA-Bench. Each prompt specifies detailed criteria, such as location accuracy, cultural relevance, specificity, and fluency, used to assign a numeric score for qualitative assessment.

erates in `bf16` precision. We also use flash attention [19]. We list the training hyperparameters in the Table 4.

9.5. Additional Results

In this Section, we discuss additional qualitative results of GAEA and compare them with selected open-source and proprietary models (as mentioned in Table 1 in the main paper). Figure 13 presents a comparison of city-country predictions against other competing models. We also show the qualitative results of GAEA on short questions (SVQA), multiple-choice questions (MCQs), and true or false questions (TF) in Figures 14 to 16. For these Figures, we highlight correct predictions with **green**, while incorrect predictions are marked as **red**. Quantitative results on GAEA-Bench are summarized in Figure 17.

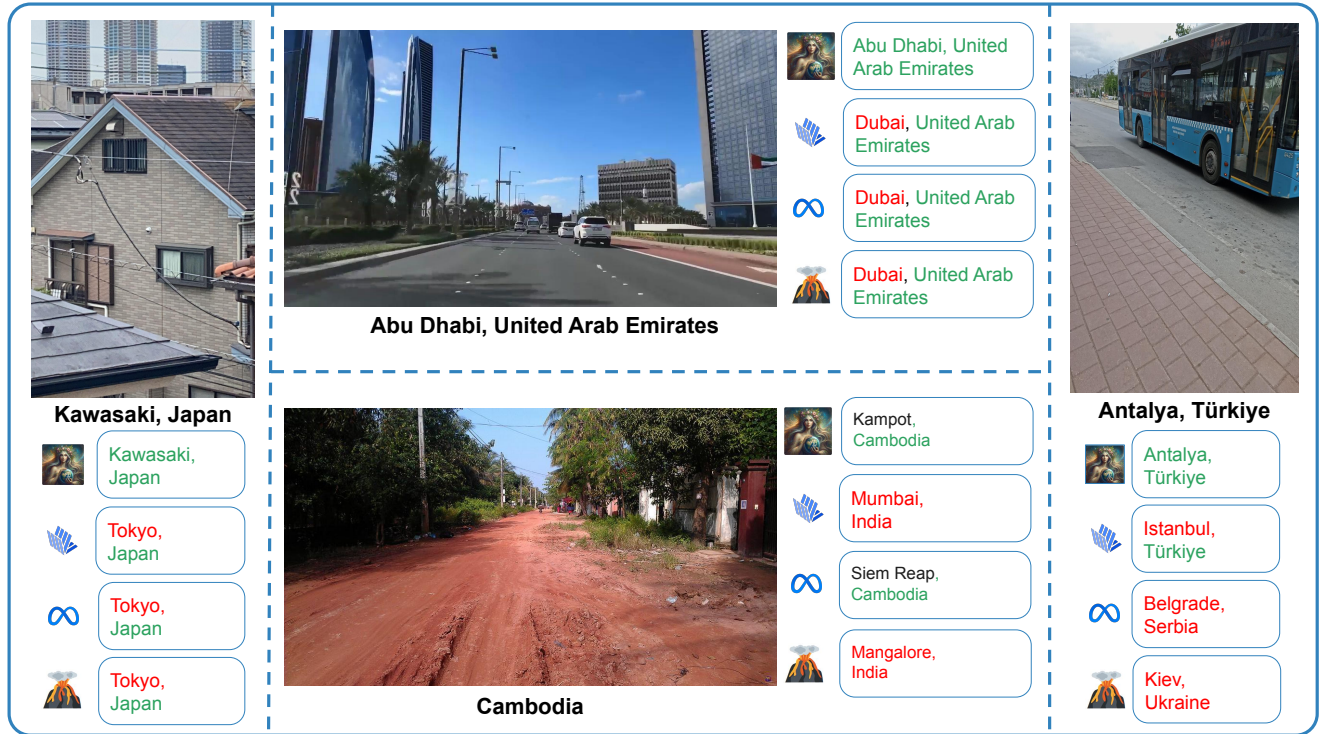


Figure 13. Qualitative Geolocalization comparison showing GAEA’s ability to predict accurate locations across varied types of scenes and landscapes around the world.

	<p>Location: San Pedro de Atacama, Chile (-22.910728, -68.201408)</p> <p>Input Question: Can you name any hostels or guest houses near this image?</p> <p>Ground Truth: Yes, there are several hostels and guest house in San Pedro de Atacama, including Hostelling International, Casa Corvatsch, Casa de Piedra, and Eden Atacameno.</p>
	The image you’ve provided appears to be of ancient rock carvings, which are likely part of an archaeological site...
	Unfortunately, I’m unable to provide a specific information about hostels or guest houses near the image, as I don’t have access to real-time information.
	The image depicts petroglyphs that are likely found in areas like the Petroglyph National Monument in New Mexico.
	It’s difficult to pinpoint specific hostels or guest houses based solely on an image of petroglyphs...
	If the site is in a place like Petra in Jordan , you could look for accommodations in nearby areas like Wadi Musa or Aqaba.
	Near where this image was taken, you can find Hostal Casa de la Piedra in San Pedro de Atacama, Chile.

Figure 14. Qualitative SVQA comparison showing GAEA’s ability to provide accurate, location-specific answers where other LMMs fail.










	Location: Elantxobe, Spain (43.403913, -2.638521)		
	Input Question: Is there a bank with an ATM near the location of this image? A. BBVA has an ATM B. No banks with ATMs nearby C. Kutxabank has an ATM D. Santander has an ATM		
Ground Truth: C. Kutxabank has an ATM			
 B. No banks with ATMs nearby	 B. No banks with ATMs nearby	 B. No banks with ATMs nearby	 C. Kutxabank has an ATM
 B. No banks with ATMs nearby	 D. Santander has an ATM	 B. No banks with ATMs nearby	 C. Kutxabank has an ATM

Figure 15. Qualitative MCQs comparison showing GAEA’s ability to provide accurate answers where other LMMs fail.










	Location: Minneapolis, United States (44.984413, -93.255192)						
	Input Question: There is a restaurant named ‘Mr. Santana’ near the location associated with this image.						
Ground Truth: True							
 False	 False	 True	 False	 False	 False	 False	 True

Figure 16. Qualitative True/False comparison showing GAEA’s ability to provide accurate answers where other LMMs fail.

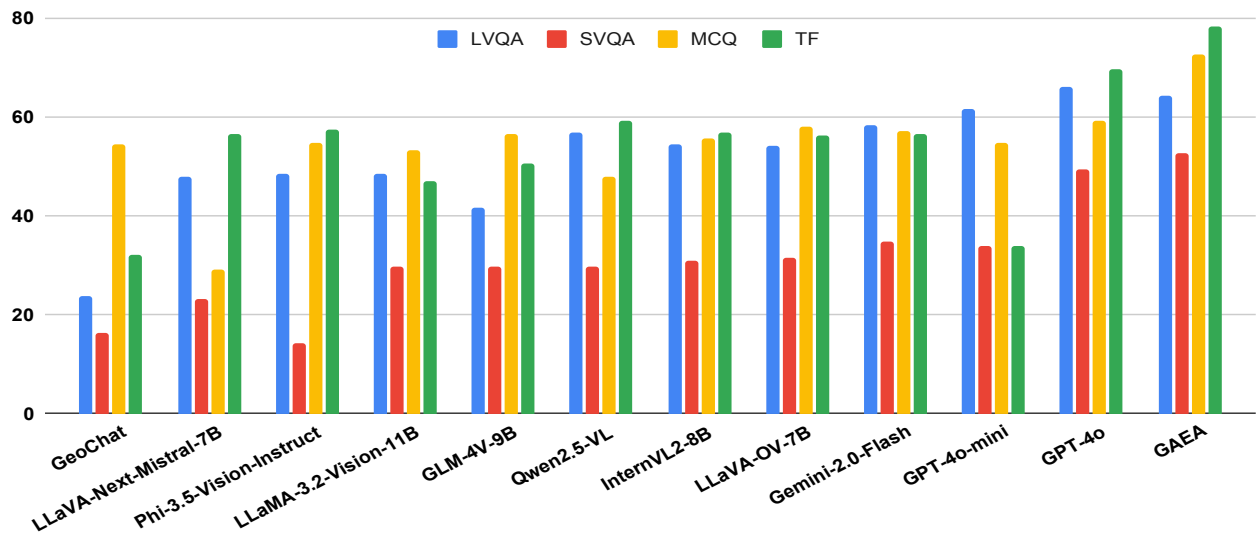


Figure 17. We showcase the performance of various LMMs on four diverse question types. GAEA outperforms on average across all question forms.