

# GeoReasoner: Geo-localization with Reasoning in Street Views using a Large Vision-Language Model

Ling Li<sup>1</sup> Yu Ye<sup>2</sup> Bingchuan Jiang<sup>3</sup> Wei Zeng<sup>1,4</sup>

## Abstract

This work tackles the problem of geo-localization with a new paradigm using a large vision-language model (LVLM) augmented with human inference knowledge. A primary challenge here is the scarcity of data for training the LVLM - existing street-view datasets often contain numerous low-quality images lacking visual clues, and lack any reasoning inference. To address the data-quality issue, we devise a CLIP-based network to quantify the degree of street-view images being locatable, leading to the creation of a new dataset comprising highly locatable street views. To enhance reasoning inference, we integrate external knowledge obtained from real geo-localization games, tapping into valuable human inference capabilities. The data are utilized to train *GeoReasoner*, which undergoes fine-tuning through dedicated reasoning and location-tuning stages. Qualitative and quantitative evaluations illustrate that *GeoReasoner* outperforms counterpart LVLMs by more than 25% at country-level and 38% at city-level geo-localization tasks, and surpasses Street-CLIP performance while requiring fewer training resources. The data and code are available at <https://github.com/lingli1996/GeoReasoner>.

## 1. Introduction

Street-view geo-localization seeks to predict geographical locations for the given street-view images. The significance of street-view geo-localization is evident in a variety of applications, spanning social studies (Ye et al., 2019b), urban planning (Shen et al., 2018), and navigation (Chalvatzaras

<sup>1</sup>The Hong Kong University of Science and Technology (Guangzhou) <sup>2</sup>Tongji University <sup>3</sup>Strategic Support Force Information Engineering University <sup>4</sup>The Hong Kong University of Science and Technology. Correspondence to: Bingchuan Jiang <jbc021@163.com>.

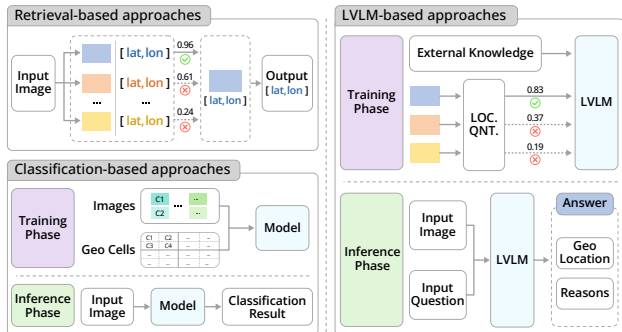


Figure 1. Different paradigms in existing and the proposed geo-localization approaches: retrieval-based (left-top), classification-based (left-bottom), and our LVL-based (right).

et al., 2022). As shown in Figure 1 (left), existing frameworks for street-view geo-localization can be mainly divided into two categories: *retrieval-based* and *classification-based*. *Retrieval-based* approaches entail identifying the most similar image within a geo-tagged image gallery and returning the corresponding geographical location (Zhu et al., 2022; Lin et al., 2022; Zhang et al., 2023b). However, the methods rely on the diversity and comprehensiveness of the geo-tagged image gallery, which can be challenging to curate. Alternatively, *classification-based* approaches partition the Earth’s surface into distinct regions and assign the input image to a specific region (Clark et al., 2023; Pramanick et al., 2022; Müller-Budack et al., 2018; Seo et al., 2018; Weyand et al., 2016). While these methods leverage shared visual features within a single region, they may neglect valuable semantic information (e.g., signboard texts) crucial for geo-localization. More importantly, these classification methods often operate as black-box models, lacking reasoning capabilities for users to interpret.

Achieving street view-based geo-localization with reasoning capability poses a considerable challenge. This study introduces a new paradigm that facilitates geo-localization with reasoning capability for street-view images, as depicted in Figure 1(right). The paradigm leverages an LVL for its excellent capability in handling multi-modal visual and textual inputs and incorporates external knowledge learned from various online games for the reasoning procedure. Specifically, we introduce the concept of *locatability* as a metric

to quantify the degree of locatability in street-view images. On this basis, we devise a CLIP-based visual-text pairing network to match large-scale Google Street View (GSV) images with 3K finely reasoned text-image pairs from online games, to tackle the challenge of the absence of a high-quality street-view dataset. The process filters through over 70K GSV images with geo-tags, all of which exhibit a high degree of locatability.

Next, we construct an LVLM model, named *GeoReasoner*, to overcome the difficulty of integrating reasoning capability in geo-localization. The training procedures of *GeoReasoner* are divided into two folds: reasoning tuning and location tuning. In the first stage, we utilize the 3K reasoned text-image pairs encapsulating human inference knowledge, to fine-tune a well-trained LVLM model with LoRA (Hu et al., 2022) for reasoning adaptation. In the second stage, we leverage the curated 70K high-locatability GSV images dataset, to further fine-tune the LVLM model with another LoRA stacked on the first one for location tuning. We assess *GeoReasoner* in terms of accuracy for both country-level (*i.e.*, predicting the country in which a street view is located) and city-level (*i.e.*, predicting the city in which a street view is located) geo-localization. The results demonstrate that *GeoReasoner* outperforms the other counterparts by more than 25% at the country-level geo-localization and 38% at the city-level geo-localization with reasoning on our test dataset. Notably, *GeoReasoner* performs slightly better than StreetCILP (Haas et al., 2023), which was trained on a substantially larger dataset of 1.1 million geo-tagged street-view images. We also evaluate *GeoReasoner* against state-of-the-art models for geo-localization using open benchmark datasets. The results show that *GeoReasoner* achieves comparable performance with only 10k Flickr images used for training. The main contributions of our work are:

- We present a new paradigm that leverages an LVLM and external knowledge of human inference for geo-localization with reasoning from street-view images.
- We introduce the concept of *locatability* and devise a CLIP-based network to quantify the degree of locatability in street-view images.
- We propose *GeoReasoner*, an LVLM that outperforms existing geo-localization models and provides detailed reasoning for the inferred results.

## 2. Related work

### 2.1. Street Views

Street views, as the realm of physical environments routinely accessed and engaged with in daily life, bear significant relevance to human perception (Ye et al., 2019b) and urban design (Shen et al., 2018). Analyses of street views

contribute to decision-making support (Ye et al., 2019a), improved understanding of urban social and economic structures (Bai et al., 2023b), and traffic asset monitoring and maintenance (Campbell et al., 2019; Li et al., 2021). This study places an emphasis on geo-localization based on street views. Specifically, drawing motivation from Zhang et al. (2018), we delineate the distribution of scene elements to quantify the degree of *locatability* in street views. Highly locatable street-view images are curated to train an LVLM that surpasses existing geo-localization models.

### 2.2. Image-based Geo-localization

Geo-localization entails determining spatial coordinates on the Earth’s surface, with broad applications in practical scenarios, including tracking individual trajectories (Cheng et al., 2022) and positioning autonomous vehicles (Chalvatzaras et al., 2022). This study focuses on image-based geo-localization, utilizing image data as input. Research on image-based geo-localization can be primarily classified into two approaches: retrieval-based (Zhu et al., 2022; Lin et al., 2022; Zhang et al., 2023b) and classification-based (Clark et al., 2023; Pramanick et al., 2022; Müller-Budack et al., 2018; Seo et al., 2018; Weyand et al., 2016).

The retrieval-based approach involves the sequential matching of a single image with a gallery of overhead views, each labeled with geographical coordinates, and identifying the result with the highest matching as the location. However, the utilization of this method is limited due to its requirement for additional reference datasets. The classification-based approach, exemplified by Weyand et al. (2016), involves subdividing the Earth’s surface into thousands of geographical cells and predicting the geographical unit to which an image belongs. The prediction effectiveness can be boosted with a dataset comprising millions of street views, whilst the granularity is influenced by the number of subdivided geographical cells. As such, many studies have been devoted to learning to corresponding multi-level features at different granularity (Vo et al., 2017), or multi-pair features for different tasks (Müller-Budack et al., 2018; Pramanick et al., 2022; Vivanco Cepeda et al., 2024).

We approach image-based geo-localization with a novel paradigm. Specifically, we integrate semantic visual concepts that offer locatable features (Theiner et al., 2022), and incorporate human reasoning knowledge learned from geo-localization games using an LVLM.

### 2.3. Vision-Language Models

The emergence of Large Language Models (LLMs) has significantly impacted various tasks related to natural language processing. These models exhibit remarkable performance in tasks such as text generation (Zhang et al., 2023a) and text-based question answering (Shao et al., 2023), owing to

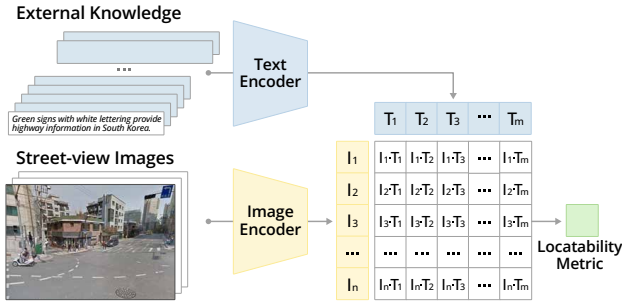


Figure 2. The locatability quantization network devises a CLIP-based visual-text pairing approach to predict the locatability metric.

their robust and versatile capabilities. As a result, research attention has shifted towards exploring prompt engineering techniques to enhance the performance of LLMs in downstream tasks (Wei et al., 2022; Yao et al., 2024; Dai et al., 2023; Xu et al., 2023; Ying et al., 2024).

Large vision-language models (LVLMs) integrate visual encoders with LLMs, exhibiting remarkable effectiveness in visual question-answering tasks (Liu et al., 2024; Bai et al., 2023a; Rao et al., 2023). This study harnesses the capabilities of LVLMs to address geo-localization of street views. However, the optimal utilization of LVLMs remains a challenging issue, particularly due to the absence of high-quality training data and a lack of reasoning capabilities. We overcome these challenges through an innovative paradigm and the thoughtful design of model architecture, contributing to a more effective utilization of LVLMs in this domain.

### 3. GeoReasoner

This section outlines our approach to addressing two challenges: 1) the absence of a high-quality street-view geo-localization dataset (discussed in Sect. 3.1), and 2) the difficulty of integrating reasoning in geo-localization (discussed in Sect. 3.2), when constructing *GeoReasoner*.

#### 3.1. Locatability-Enhanced Data Curation

Throughout the development of this work, we observed variations in the degree of locatability among different street views. For example, the images featuring textual signboards or prominent landmarks (e.g., Eiffel Tower) are easily locatable, whilst those captured in a tunnel or obscured by a wall tend to be less locatable. Refer to Figure 4 for further illustration. Simply merging all these street-view images to train an LVLM is not an optimal approach, as the inclusion of poor-quality data can adversely affect the training efficiency of updating an LVLM (Radford et al., 2021). To this end, we introduce *locatability*, a metric that quantifies the level of locatability of street-view images. We then devise a CLIP-based visual-text pairing network to produce the desired locatability metric for an input street-view image,

as shown in Figure 2. The network naturally incorporates data from two perspectives:

- **Street-View Images.** We collected street-view images from the Google Street View<sup>1</sup> (GSV). To enrich the diversity of the dataset, we first selected the top global cities according to the Globalization and World Cities Study Group and Network (GaWC) ranking. Next, we utilized the global OpenStreetMap<sup>2</sup> (OSM) geographic database to obtain the vector data of the road network in these urban areas. The road network was passed to ArcPy, a Python site package of ArcGIS, to automatically extract sampling points at 4000-meter intervals and generate a CSV table containing information about these sampling points. Subsequently, we employed the GSV API interface to compile a comprehensive dataset. This dataset encompassed street-view images captured from four distinct directions - front, back, left, and right - of each sampling point. Considering the impact of data sparsity and image similarity, we randomly selected two of the four views from each data point, denoted as  $[I_x, I_y]$ , where  $x \in (left, right)$ ,  $y \in (front, back)$ . The process has yielded a total of over 130k street-view images with geo-tags collected from 72 cities in 48 countries.
- **Textual Clues.** Textual clues often serve a pivotal role in delineating the geographical locations of street-view images. However, there is currently no readily available textual clue dataset. Two prominent games, namely GeoGuessr<sup>3</sup> and Tuxun<sup>4</sup>, which are centered around geo-localization based on street views, provide a potential solution to this gap. The communities surrounding these games possess a well-structured and efficacious collection of textual clues, utilized for geographical positioning across various countries and cities. These clues are collaboratively maintained and curated by both administrators and players. The domain knowledge lays the groundwork for identifying decisive factors and assessing the significance of each factor in street views. We acquired image-text data pairs for geo-localization from the two open-source communities. Subsequently, we utilized the BERT-based Named Entity Recognition (NER) (Kenton & Toutanova, 2019) model to clean and filter text that lacked specific geographical location information. In this way, we collected a total of over 3K textual clues that encapsulate rich geo-localization information. For instance, “houses in central Chile are more likely to have terracotta tiled roofs”. Each clue is paired with a corresponding street-view image.

<sup>1</sup><https://www.google.com/streetview>

<sup>2</sup><https://www.openstreetmap.org>

<sup>3</sup><https://www.geoguessr.com>

<sup>4</sup><https://tuxun.fun>

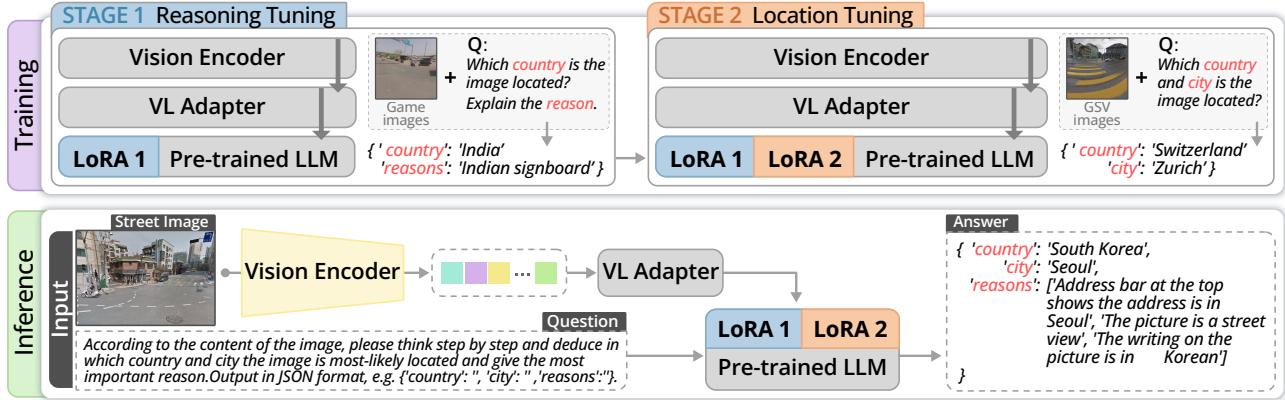


Figure 3. The architecture of *GeoReasoner* consists of three modules: *Vision Encoder*, *VL Adapter* and *Pre-trained LLM*. The model undergoes a two-fold supervised fine-tuning process: reasoning tuning and location tuning, to enable geo-localization with reasoning.

With the GSV images and textual clues, our subsequent task is to filter GSV images with a high degree of locatability, for the purpose of training an LVLM. To achieve this, we design a CLIP-based visual-text pairing network. As depicted in Figure 2, the GSV images undergo processing by an image encoder that deduces the image attributes.

Here, we first use MaskFormer (Cheng et al., 2021) to predict segmentation masks for various classes in GSV images, such as buildings, sky, and vehicles. We then compute an  $n$ -length vector  $\mathbf{I}_{seg}$ , which quantifies the area ratio of each mask class, where  $n$  represents the number of classes. Subsequently, we utilize Sentence-BERT (Reimers & Gurevych, 2019) to measure the similarity between textual clues and semantic segmentation labels, yielding an  $m \times n$  matrix  $M$ , where  $m$  is the number of textual clues. After that, we normalize  $M$  using min-max normalization, and set values lower than the threshold to zero, resulting in another  $m \times n$  matrix  $\hat{M}$ . We reduce  $\hat{M}$  to an  $n$ -length vector by calculating the mean across its rows, and then normalize it to obtain  $\mathbf{w}_{loc}$ . This vector represents the importance of each semantic segmentation label for geo-localization. With the segmentation mask ratio  $\mathbf{I}_{seg}$  and the corresponding weight  $\mathbf{w}_{loc}$ , the locatability metric of a GSV image is computed through the multiplication and accumulation of the respective values, as follows:

$$locatability(\mathbf{I}_{seg}, \mathbf{w}_{loc}) = \sum_{k=1}^n \mathbf{I}_{seg}(k) \cdot \mathbf{w}_{loc}^k \quad (1)$$

where  $\mathbf{I}_{seg}(k)$  denotes pixel ratio of the  $k$ -th class in the segmentation mask  $\mathbf{I}_{seg}$ .

A higher *locatability* value indicates a higher degree of visual clues exhibited in a GSV image for geo-localization, while a lower value suggests the opposite. Empirically, we selected a threshold value of 0.4 for filtering locatable GSV images. This resulted in over 70k highly locatable images with geo-tags passing to the next stage for training an LVLM.

### 3.2. Geo-localization with Reasoning

While many models (e.g., Clark et al. (2023); Pramanick et al. (2022); Müller-Budack et al. (2018); Seo et al. (2018); Weyand et al. (2016)) exist for image-based geo-localization, these models typically predict locations without providing the inference process. This introduces several limitations: First, the models operate as black boxes without providing insights, making it challenging for users to interpret. This obstacle impedes further refinement of the geo-localization model. More importantly, studies have demonstrated that integrating the reasoning process can enhance the capabilities of LLMs (Qiao et al., 2023). Therefore, our objective is to construct an LVLM for image-based geo-localization with reasoning capability.

**Model Architecture.** Figure 3 illustrates the architecture of the proposed model *GeoReasoner*, which is based on Qwen-VL (Bai et al., 2023a). *GeoReasoner* consists of three modules: *Vision Encoder*, *Vision-Language (VL) Adapter* and *Pre-trained LLM*. Specifically, the *Vision Encoder* module employs the Vision Transformer (ViT) (Dosovitskiy et al., 2021) architecture. The input street-view images are resized to a specific resolution and then divided into a set of image patches. To refine image patches into sequential representations compatible with an LLM, the *VL Adapter* is introduced. In the *VL Adapter*, the sequence of visual features is initially condensed to a fixed length to address efficiency challenges posed by the substantial number of visual feature sequences. Subsequently, the processed visual features are integrated with the LLM using cross-attention mechanisms. Following this, the compressed visual feature sequence and text sequence are passed to the *Pre-trained LLM* module, which functions as a decoder for generating the answer.

**Supervised Fine-tuning.** The overall model undergoes a staged pre-training process that is divided into two folds: reasoning tuning and location tuning. In the first stage, our objective is to enhance the model’s reasoning capability



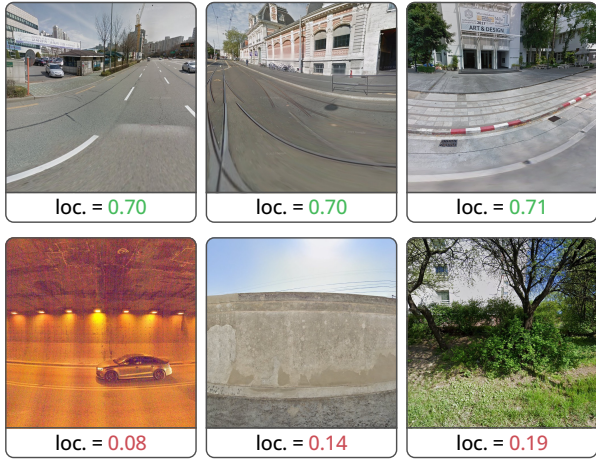


Figure 4. Locatability examples. Top row: the street views are highly locatable by signboards, architectural styles, and landmarks. Bottom row: no visual clues for locating the street views.

by utilizing textual clues paired with street-view images collected from geo-localization games. The input street-view image & question, and the output answer are formatted as prompts in the following manner:

```
[INPUT] {image : IMAGE,
          question : 'which country is the image located?
                    Explain the reason'}
[OUTPUT] {country : 'South Korea',
          reason : 'Korean signboard'}
```

Here, we can only provide reasoning at the country level due to the granularity exhibited in the image-text pairs. Nevertheless, this reasoning procedure is sufficient to facilitate the second stage of location tuning. Next, we integrate the prior knowledge of country information with highly locatable GSV images with geo-tags to infer the fine-grained city-level location information. We utilize a similar prompt format as in the first stage but without a reasoning requirement. Both stages are fine-tuned from the pre-trained Qwen-VL with LoRA, which contributes to the overall performance improvement of Qwen-VL in both the reasoning and location tuning stages, allowing the model to better capture complex relationships within the image-text pairs.

## 4. Experiments

We conduct a series of experiments to evaluate the effectiveness of the locatability-enhanced geo-localization dataset (Sect. 4.1) and the model *GeoReasoner* for geo-localization with reasoning (Sect. 4.2).

### 4.1. Experiments on Locatability-Enhanced Dataset

#### 4.1.1. QUALITATIVE COMPARISON

Figure 4 presents examples of the predicted locatability degrees of different street-view images by our locatability quantization network. The top row showcases street views




Figure 5. The relationship between building proportion and the degree of locatability in street views. The locatability metric peaks when the building proportion is approximately 0.2.

distinguished by prominent localizable attributes. The left image features the Korean language on a signboard, the middle image captures the distinctive *Art Nouveau* architectural style commonly found in Switzerland, and the right image shows an art & design museum in India. In contrast, street views in the bottom row display lower locatability degrees. The left image resembles a tunnel, lacking additional discernible information for accurate localization. Similarly, the middle image is occluded by a wall, and the right image faces common vegetation that is available worldwide.

For the proposed locatability metric in Equation 1, we also evaluated the relationship between building proportion and the degree of locatability of street views. The results are shown in Figure 5. The locatability metric slightly increases as the building proportion ranges from 0 to 0.2, but decreases as the building proportion continues to increase. The results indicate that buildings are not the sole determinant of locatability. As the proportion of buildings increases, the street-view images transition from panoramic to close-up views, leading to reduced information availability and consequently diminishing the degree of locatability.

The qualitative analysis indicates the effectiveness of the locatability quantization network in predicting locatability degrees of street-view images. Furthermore, the prediction aligns with human inference knowledge harvested from real geo-localization games, providing the ground truths for fine-




Figure 6. Quantitative comparison of country- and city-level geo-localization accuracy by different models trained on mixed datasets with varying proportions of high locatable GSV images.

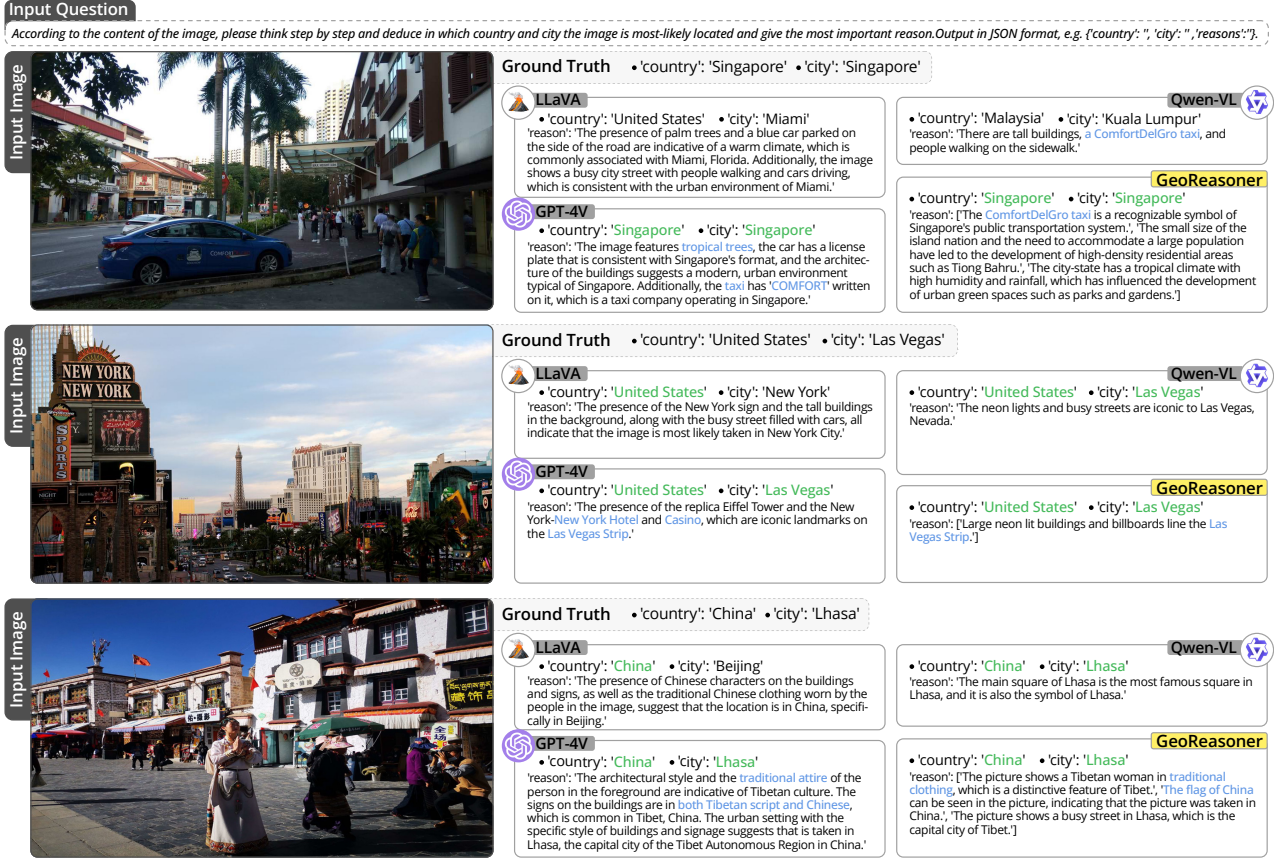


Figure 7. Examples of LVLM-based approaches in geo-localization with reasoning. Prediction results matching the ground truth are highlighted in green, while reasons offering valid information are marked in blue.

tuning the reasoning component in *GeoReasoner*.

#### 4.1.2. QUANTITATIVE COMPARISON

We conducted quantitative experiments to investigate the importance of using high-locatability GSV images in training the location component in *GeoReasoner*. Various datasets were prepared, featuring different proportions of high-locatability GSV images, ranging from 0% (only low-locatability GSV images) to 100% (only high-locatability GSV images). To ensure fairness, each experimental group retained consistent 10K GSV images, with only the proportion of high-locatability images varying. Subsequently, models were trained for each dataset, and their accuracy in country- and city-level geo-localization was evaluated on a randomly sampled set of 1K GSV images.

The experimental results are presented in Figure 6. Overall, the results reveal that as the proportion of high-locatability GSV images in the training dataset increases, the performance of the fine-tuned location component improves in both country- and city-level geo-localization. Specifically, the country- and city-level geo-localization accuracy increases from 0.63 & 0.47 for 0% high-locatability GSV images, to 0.72 & 0.51 for 100% high-locatability GSV

images. Notably, the experiments only utilize 10K GSV images instead of all the curated 70K high-locatability GSV images due to training complexity. Nevertheless, the results demonstrate that high-locatability GSV images offer more meaningful insights and less extraneous noise, making them highly valuable in the geo-localization task.

## 4.2. Experiments on Geo-localization with Reasoning

### 4.2.1. QUALITATIVE COMPARISON WITH SOTA

To assess the efficacy of *GeoReasoner* in terms of geo-localization with reasoning, we conduct a qualitative comparison with state-of-the-art LVLM-based approaches, including LLaVA (Liu et al., 2024), Qwen-VL (Qwen-7B) (Bai et al., 2023a), and GPT-4V (Achiam et al., 2023). In the experimental phase, we presented the same input street-view images, reasoning process, and result formats to these models. Specifically, a consistent prompt is used, as below:

According to the content of the image, please think step by step and deduce in which country and city the image is most likely located and offer possible explanations. Output in JSON format, e.g., {'country': '', 'city': '', 'reasons': ''}.

Table 1. Comparison of Precision, Recall and F1 scores in country-level and city-level geo-localization. \* represents the model trained on high-locatability GSV images.

| Model                                 | Country   |         |               | City      |         |               |
|---------------------------------------|-----------|---------|---------------|-----------|---------|---------------|
|                                       | Accuracy↑ | Recall↑ | F1↑           | Accuracy↑ | Recall↑ | F1↑           |
| StreetCLIP (Haas et al., 2023)        | 0.7943    | 1.00    | 0.8854        | 0.7457    | 1.00    | 0.8543        |
| LLaVA (Liu et al., 2024)              | 0.4029    | 1.00    | 0.5744        | 0.2400    | 1.00    | 0.3871        |
| Qwen-VL (Qwen-7B) (Bai et al., 2023a) | 0.5829    | 0.95    | 0.7225        | 0.3743    | 0.89    | 0.5270        |
| GPT-4V (Achiam et al., 2023)          | 0.8917    | 0.34    | 0.4923        | 0.5083    | 0.31    | 0.3851        |
| ViT* (Dosovitskiy et al., 2021)       | 0.7100    | 1.00    | 0.8304        | 0.6762    | 1.00    | 0.8068        |
| GeoReasoner*                          | 0.8237    | 1.00    | <b>0.9033</b> | 0.7521    | 1.00    | <b>0.8585</b> |

Figure 7 illustrates the inference results of counterpart models and *GeoReasoner* on three diverse street views from different countries and cities—namely, Singapore-Singapore (top), United States-Las Vegas (middle) and China-Lhasa (bottom). Overall, *GeoReasoner* not only outperforms existing models in the accuracy of country or city-level predictions but also provides coherent explanations with insightful reasoning for the inference results.

In Figure 7 (top), *GeoReasoner* identifies the word ‘COMFORT’ on the taxi in the image. Drawing from prior knowledge, ‘the ComfortDelGro taxi is a distinctive symbol of Singapore’s public transportation system’ in the text-image pairs, the model deduces that the area is likely to be in Singapore. GPT-4V predicts the same geo-location with accurate reasoning, yet the other two models fail, either due to not recognizing the taxi by LLaVA or making an incorrect inference about the city by Qwen-VL.

Figure 7 (middle) presents a scene of the Las Vegas Strip. A conspicuous ‘NEW YORK’ sign is prominently visible in the upper-left corner of the image. This sign causes the reasoning error in the task performed by LLaVA. Although Qwen-VL generates accurate predictions of Las Vegas-United States, the most essential factor, i.e., ‘Las Vegas Strip’, is not considered in the reasoning process. In contrast, both *GeoReasoner* and GPT-4V provide the correct geo-location along with accurate inference.

Based on the depiction of Chinese characters and traditional clothing in Figure 7 (right), all models make accurate predictions regarding the country, identifying it as China. However, LLaVA makes an incorrect prediction of the city, specifying Beijing. In contrast, the other models successfully predict the city as Lhasa, providing sensible and justifiable reasons for their inferences.

#### 4.2.2. QUANTITATIVE COMPARISON WITH SOTA

We further conduct quantitative experiments to compare with counterparts LVLMs. In addition, we choose StreetCLIP (Haas et al., 2023) as the state-of-the-art classification-based approach and omit retrieval-based approaches relying on a geo-tagged image gallery that is not available. It is

important to clarify that, for the LVLm-based approaches, obtaining corresponding and relevant answers is not guaranteed at all times. Therefore, we included *Recall* rate to measure the proportion of effective answers within the large language models. When calculating the *Accuracy* rate, only the accuracy of these effective answers is taken into account. We additionally compute *F1* values, taking into consideration both *Accuracy* and *Recall* metrics.

Table 1 presents the prediction results by the counterparts and *GeoReasoner*. Overall, *GeoReasoner* outperforms all the counterparts, particularly those LVLm-based approaches. Taking the best performed Qwen-VL for example, *GeoReasoner* outperforms it 25.02% on country-level geo-localization and 38.61% on city-level geo-localization, in terms of F1 value. Surprisingly, the recall performance of GPT-4V for the geo-localization task was notably low. Most of the responses were mainly: ‘I’m sorry, I can’t provide assistance with that request.’ or ‘I’m sorry, but I am unable to provide the exact location, such as the country and city, for the image you have provided. My capabilities do not include analyzing specific details to determine the geographical location of the image content.’

We speculate that GPT-4V has undergone extensive measures to ensure the model’s security and privacy, which may contribute to its reluctance or denial of recognition in the task of geo-localization.

In comparison to StreetCLIP that is specialized in geo-localization, *GeoReasoner* demonstrates only a slight superiority. Nevertheless, it’s important to note that StreetCLIP was trained on a significantly larger dataset of over 1.1 million street-view images, while our *GeoReasoner* was trained with only 70K street views. For ViT trained on the same data, *GeoReasoner* still exhibits superior geolocation capabilities. Moreover, *GeoReasoner* offers reasoning capability, providing added value for various downstream tasks.

#### 4.2.3. ABLATION EXPERIMENTS

To assess the contributions of the location tuning and reasoning tuning components in *GeoReasoner*, we design several ablation experiments using the Qwen-VL (Bai et al., 2023a)



Table 2. Results of the ablation experiments using baseline Qwen-VL (Qwen-7B), GeoReasoner w/o location tuning, GeoReasoner w/o reasoning tuning, and the full GeoReasoner models.

| Model                            | Training  |          | Performance   |                 |               |               |              |               |
|----------------------------------|-----------|----------|---------------|-----------------|---------------|---------------|--------------|---------------|
|                                  | Reasoning | Location | Accuracy↑     | Country Recall↑ | F1↑           | Accuracy↑     | City Recall↑ | F1↑           |
| Qwen-VL (Qwen-7B)                | -         | -        | 0.5829        | 0.95            | 0.7225        | 0.3743        | 0.89         | 0.5270        |
| GeoReasoner w/o location tuning  | ✓         | ×        | 0.6971        | 1.00            | 0.8215        | 0.4114        | 0.99         | 0.5813        |
| GeoReasoner w/o reasoning tuning | ×         | ✓        | 0.7803        | 1.00            | 0.8766        | 0.7029        | 1.00         | 0.8255        |
| GeoReasoner                      | ✓         | ✓        | <b>0.8237</b> | 1.00            | <b>0.9033</b> | <b>0.7521</b> | 1.00         | <b>0.8584</b> |

Table 3. Comparison results on Im2GPS dataset. The top five rows are derived from the results reported in the paper, while the last four rows are from retesting on the filtered Im2GPS dataset, which includes only highly locatable data.

| Model        | Dataset w/ Filter |      | Street 1km | City 25km | Country 750km |
|--------------|-------------------|------|------------|-----------|---------------|
|              | Train             | Test |            |           |               |
| PlaNet       | ×                 | ×    | 0.08       | 0.25      | 0.54          |
| CPlaNet      | ×                 | ×    | 0.17       | 0.37      | 0.62          |
| ISNs         | ×                 | ×    | 0.17       | 0.43      | 0.67          |
| Translocator | ×                 | ×    | 0.20       | 0.48      | 0.76          |
| GeoDecoder   | ×                 | ×    | 0.22       | 0.50      | 0.80          |
| ISNs         | ×                 | ✓    | 0.25       | 0.43      | 0.78          |
| GeoCLIP      | ×                 | ✓    | 0.25       | 0.49      | 0.87          |
| GeoReasoner  | ×                 | ✓    | 0.10       | 0.41      | 0.82          |
| GeoReasoner  | ✓                 | ✓    | 0.13       | 0.44      | 0.86          |

pre-trained model as the baseline. Next, we integrated the Qwen-VL pre-trained model with LoRA1 (*GeoReasoner* without location tuning) and LoRA2 (*GeoReasoner* without reasoning tuning). The last experiment involved the full *GeoReasoner* model, including both the location tuning and reasoning tuning components. The same prompts were utilized for all these models, as in the previous experiments.

Table 2 presents the quantitative results in terms of *accuracy*, *recall*, and *F1*. Overall, the results indicate that both the location tuning and reasoning tuning components improve the model performance. Specifically, the location tuning component is essential for geo-localization, as *GeoReasoner* w/o reasoning tuning (row 3) achieves much higher accuracy than *GeoReasoner* w/o location tuning (row 2), especially for fine-grained city-level prediction. This result further strengthens the evidence that high-locatability GSV images are essential for geo-localization. The reasoning tuning component also plays a significant role in the performance improvement, as evidenced by the superior performance of the full *GeoReasoner* (row 4).

#### 4.2.4. GENERALIZABILITY EVALUATION

To further assess the generalizability of *Georeasoner* in geo-localization, we conduct additional testing on open Flickr image datasets of Im2GPS (Hays & Efros, 2008) and

Table 4. Comparison results on Im2GPS3k dataset. The top six rows are derived from the results reported in the paper, while the last four rows are from retesting on the filtered Im2GPS3k dataset, which includes only highly locatable data.

| Model        | Dataset w/ Filter |      | Street 1km | City 25km | Country 750km |
|--------------|-------------------|------|------------|-----------|---------------|
|              | Train             | Test |            |           |               |
| PlaNet       | ×                 | ×    | 0.09       | 0.25      | 0.48          |
| CPlaNet      | ×                 | ×    | 0.10       | 0.27      | 0.49          |
| ISNs         | ×                 | ×    | 0.11       | 0.28      | 0.50          |
| Translocator | ×                 | ×    | 0.12       | 0.31      | 0.59          |
| GeoDecoder   | ×                 | ×    | 0.13       | 0.34      | 0.61          |
| GeoCLIP      | ×                 | ×    | 0.14       | 0.34      | 0.70          |
| ISNs         | ×                 | ✓    | 0.10       | 0.29      | 0.59          |
| GeoCLIP      | ×                 | ✓    | 0.12       | 0.38      | 0.83          |
| GeoReasoner  | ×                 | ✓    | 0.09       | 0.35      | 0.74          |
| GeoReasoner  | ✓                 | ✓    | 0.10       | 0.38      | 0.83          |

Im2GPS3k (Vo et al., 2017). Here we use only 10k Flickr images for fine-tuning *Georeasoner*. Since *Georeasoner* predicts city names rather than GPS coordinates, we first convert the predicted city names generated by *Georeasoner* into the GPS coordinates of their respective city centers, then measure the distance between these predicted coordinates and ground-truth locations.

Table 3 and Table 4 present the performance comparison of *Georeasoner* with PlaNet (Weyand et al., 2016), CPlaNet (Seo et al., 2018), ISNs (Müller-Budack et al., 2018), Translocator (Pramanick et al., 2022), GeoDecoder (Clark et al., 2023), and GeoCLIP (Vivanco Cepeda et al., 2024) on Im2GPS and Im2GPS3k datasets, respectively. The results demonstrate that fine-tuning *GeoReasoner* using highly locatable images significantly improves prediction accuracy for street, city, and country levels (row 8 vs. row 9 in Table 3, and row 9 vs. row 10 in Table 4). Remarkably, despite being fine-tuned solely on a smaller number of Flickr images, *GeoReasoner* achieves results comparable to ISNs and GeoCLIP trained on millions of Flickr images, particularly in terms of city- and country-level accuracy. Besides, *GeoReasoner* trained on the filtered, highly locatable Flickr images also show improvements in the city- and country-level geo-localization, demonstrating



the generalizability of our proposed *locatability* module.

## 5. Discussion

### The significance of high-locatability street-view images.

We observe a significant performance improvement when *GeoReasoner* is trained upon high-locatability street-view images. Such images often contain explicit visual clues such as stylized architecture, traffic signs, and landmarks, providing the model with richer contextual information. Therefore, increasing the quality of the training dataset enhances the model’s geo-localization performance. Additionally, the quantity of high-locatability images is vital, as the model trained with 70K images (as in Sect. 4.2.2) achieves significantly higher accuracy than the one trained with 10K images (Sect. 4.1.2). In balancing the quality and quantity of the training dataset, we empirically applied a threshold of 0.4 to differentiate between highly and less localizable street views. Setting the threshold too high (e.g., 0.7) can lead to a notable decrease in the amount of high-locatability images, whilst a lower threshold (e.g., 0.1) may bring in introduce low-quality images.

**The necessity of reasoning process.** The introduction of the reasoning component successfully elevated *GeoReasoner*’s performance in the geo-localization task. This signifies that LVLM can adeptly capture intricate relationships among image features, location clues, and geo-locations in the training process. Implemented an innovative solution to empower the reasoning capability within *GeoReasoner* by leveraging human inference knowledge extracted from geo-localization games. Despite the relatively small dataset, a noticeable improvement in performance has been achieved. In the future, we plan to expand the reasoning dataset by diversifying the influencing clues. For instance, the current textual clues are absent of landscape information, which could provide invaluable insights for geo-localization. We will collaborate with domain experts such as urban planners and geographers to address these limitations.

**Failure cases.** *GeoReasoner* comprehends architectural style as a pivotal factor in geo-localization. However, the model can be misled by the learned significance of architectural style. Figure 8 presents a street view of the Eiffel Tower in Paris, France (left), and replicas of the Eiffel Tower in New York, USA (middle) and in Hangzhou, China (right). *GeoReasoner* fails to distinguish between them, predicting all instances as located in Paris, France. This misclassification is not unique to *GeoReasoner* but also extends to other LVLMs like GPT-4V. Consequently, it underscores the necessity for LVLM-based methods to delve deeper into knowledge for more sophisticated geo-localization capabilities. Once again, it is imperative to collaborate with domain experts and enhance the visual clues and reasoning procedure comprehensively to tackle this issue.

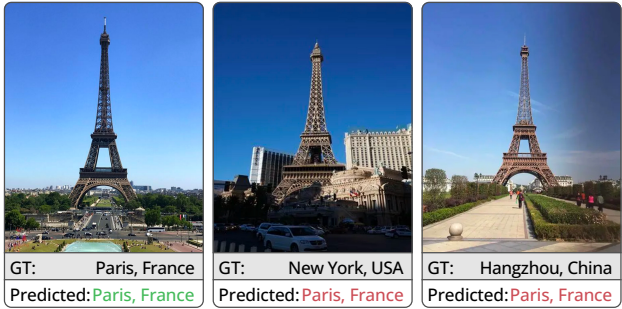


Figure 8. *GeoReasoner* fails to distinguish the Eiffel Tower and its replicas in New York, USA, and Hangzhou, China.

## 6. Conclusion

In this paper, we present a new paradigm that integrates a large vision-language (LVLM) model with human inference knowledge for street view geo-localization with reasoning. We introduce the concept of *locatability* and devise a CLIP-based network to quantify the degree of locatability in street-view images, facilitating the selection of high-quality data. We design an LVLM-based model named *GeoReasoner*, which harnesses external knowledge of human inference from real geo-localization games and curated high-quality data to enhance the performance of geo-localization tasks with reasoning capabilities. The model undergoes two-stage fine-tuning, namely *reasoning tuning* and *location tuning*. The *reasoning tuning* stage aims to acquire potential linkage between coarse-grained geographical locations (i.e., country) and the associated positioning reasons. In *location tuning* stage, we employ the curated high-quality data to further refine the model in fine-grained geo-localization (i.e., city) learning. Extensive experiments prove that *GeoReasoner* outperforms previous models qualitatively and quantitatively.

## Acknowledgements

We would like to thank Yao Zhou and Wenqi Shao for their insightful discussions and Ziyao Gao for her assistance in drawing the figures in this paper. We also extend our gratitude to the anonymous reviewers for their valuable comments. This work is partially supported by the National Natural Science Foundation of China (62172398, 42171456, 52078343).

## Impact Statement

*GeoReasoner* advances image-based geo-localization technologies that are pivotal for many applications such as autonomous navigation. The pipeline of constructing the dataset featuring high-locatability street views proves highly beneficial across multiple scenarios, such as urban studies, culture studies, and digital humanities, all of which are in-

creasingly reliant on the analysis of high-quality street-view data.

The proposed paradigm represents the fusion of LVLM with human inference knowledge, which has implications for the advancement of artificial intelligence (AI) that is more aligned with human cognition. The synergy can lead to the creation of AI that is not only more effective in complex inference tasks but also more understandable and relatable to human users. As AI becomes more pervasive in daily life, the importance of designing systems that are both transparent and capable of complex reasoning cannot be overstated.

## References

- Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Alteschmidt, J., Altman, S., Anadkat, S., et al. GPT-4 Technical Report. *arXiv preprint arXiv:2303.08774*, 2023.
- Bai, J., Bai, S., Yang, S., Wang, S., Tan, S., Wang, P., Lin, J., Zhou, C., and Zhou, J. Qwen-VL: A Versatile Vision-Language Model for Understanding, Localization, Text Reading, and Beyond. *arXiv preprint arXiv:2308.12966*, 2023a.
- Bai, Y., Shang, C., Li, Y., Shen, L., Jin, S., and Shen, Q. Transport Object Detection in Street View Imagery Using Decomposed Convolutional Neural Networks. *Mathematics*, 11(18):3839, 2023b.
- Campbell, A., Both, A., and Sun, Q. C. Detecting and mapping traffic signs from Google Street View images using deep learning and GIS. *Computers, Environment and Urban Systems*, 77:101350, 2019.
- Chalvatzaras, A., Pratikakis, I., and Amanatiadis, A. A. A Survey on Map-Based Localization Techniques for Autonomous Vehicles. *IEEE Transactions on Intelligent Vehicles*, 8(2):1574–1596, 2022.
- Cheng, B., Schwing, A., and Kirillov, A. Per-Pixel Classification is Not All You Need for Semantic Segmentation. *Advances in Neural Information Processing Systems*, 34:17864–17875, 2021.
- Cheng, W., Wen, R., Huang, H., Miao, W., and Wang, C. OPTDP: Towards optimal personalized trajectory differential privacy for trajectory data publishing. *Neurocomputing*, 472:201–211, 2022.
- Clark, B., Kerrigan, A., Kulkarni, P. P., Cepeda, V. V., and Shah, M. 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*, pp. 23182–23190, 2023.
- Dai, D., Sun, Y., Dong, L., Hao, Y., Ma, S., Sui, Z., and Wei, F. Why Can GPT Learn In-Context? Language Models Implicitly Perform Gradient Descent as Meta-Optimizers. In *Findings of the Association for Computational Linguistics: ACL 2023*, pp. 4005–4019, 2023.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *Proceedings of International Conference on Learning Representations*, 2021.
- Haas, L., Alberti, S., and Skreta, M. Learning Generalized Zero-Shot Learners for Open-Domain Image Geolocalization. *arXiv preprint arXiv:2302.00275*, 2023.
- Hays, J. and Efros, A. A. IM2GPS: estimating geographic information from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1–8. IEEE, 2008.
- Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., and Chen, W. LoRA: Low-Rank Adaptation of Large Language Models. In *Proceedings of International Conference on Learning Representations*, 2022.
- Kenton, J. D. M.-W. C. and Toutanova, L. K. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186, 2019.
- Li, Y., Wu, C., Li, L., Liu, Y., and Zhu, J. Caption Generation From Road Images for Traffic Scene Modeling. *IEEE Transactions on Intelligent Transportation Systems*, 23(7):7805–7816, 2021.
- Lin, J., Zheng, Z., Zhong, Z., Luo, Z., Li, S., Yang, Y., and Sebe, N. Joint Representation Learning and Key-point Detection for Cross-View Geo-Localization. *IEEE Transactions on Image Processing*, 31:3780–3792, 2022.
- Liu, H., Li, C., Wu, Q., and Lee, Y. J. Visual Instruction Tuning. *Advances in Neural Information Processing Systems*, 36, 2024.
- Müller-Budack, E., Pustu-Iren, K., and Ewerth, R. Geolocation Estimation of Photos using a Hierarchical Model and Scene Classification. In *Proceedings of the European Conference on Computer Vision*, pp. 563–579, 2018.
- Pramanick, S., Nowara, E. M., Gleason, J., Castillo, C. D., and Chellappa, R. Where in the World is this Image? Transformer-based Geo-localization in the Wild. In *Proceedings of the European Conference on Computer Vision*, pp. 196–215, 2022.

- Qiao, S., Ou, Y., Zhang, N., Chen, X., Yao, Y., Deng, S., Tan, C., Huang, F., and Chen, H. Reasoning with Language Model Prompting: A Survey. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pp. 5368–5393, 2023.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al. Learning Transferable Visual Models From Natural Language Supervision. In *Proceedings of International Conference on Machine Learning*, pp. 8748–8763. PMLR, 2021.
- Rao, J., Shan, Z., Liu, L., Zhou, Y., and Yang, Y. Retrieval-based Knowledge Augmented Vision Language Pre-training. In *Proceedings of the 31st ACM International Conference on Multimedia*, pp. 5399–5409, 2023.
- Reimers, N. and Gurevych, I. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pp. 3980–3990, 2019.
- Seo, P. H., Weyand, T., Sim, J., and Han, B. CPlaNNet: Enhancing Image Geolocation by Combinatorial Partitioning of Maps. In *Proceedings of the European Conference on Computer Vision*, pp. 536–551, 2018.
- Shao, Z., Yu, Z., Wang, M., and Yu, J. Prompting Large Language Models with Answer Heuristics for Knowledge-based Visual Question Answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14974–14983, 2023.
- Shen, Q., Zeng, W., Ye, Y., Mueller Arisona, S., Schubiger, S., Burkhard, R., and Qu, H. StreetVizor: Visual Exploration of Human-Scale Urban Forms Based on Street Views. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):1004–1013, 2018.
- Theiner, J., Müller-Budack, E., and Ewerth, R. Interpretable Semantic Photo Geolocation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 750–760, 2022.
- Vivanco Cepeda, V., Nayak, G. K., and Shah, M. GeoCLIP: Clip-Inspired Alignment between Locations and Images for Effective Worldwide Geo-localization. *Advances in Neural Information Processing Systems*, 36, 2024.
- Vo, N., Jacobs, N., and Hays, J. Revisiting IM2GPS in the Deep Learning Era. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2621–2630, 2017.
- Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., Le, Q. V., Zhou, D., et al. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. *Advances in Neural Information Processing Systems*, 35: 24824–24837, 2022.
- Weyand, T., Kostrikov, I., and Philbin, J. PlaNet - Photo Geolocation with Convolutional Neural Networks. In *Proceedings of the European Conference on Computer Vision*, pp. 37–55, 2016.
- Xu, P., Shao, W., Zhang, K., Gao, P., Liu, S., Lei, M., Meng, F., Huang, S., Qiao, Y., and Luo, P. LVLm-eHub: A Comprehensive Evaluation Benchmark for Large Vision-Language Models. *arXiv preprint arXiv:2306.09265*, 2023.
- Yao, S., Yu, D., Zhao, J., Shafran, I., Griffiths, T., Cao, Y., and Narasimhan, K. Tree of Thoughts: Deliberate Problem Solving with Large Language Models. *Advances in Neural Information Processing Systems*, 36, 2024.
- Ye, Y., Richards, D., Lu, Y., Song, X., Zhuang, Y., Zeng, W., and Zhong, T. Measuring daily accessed street greenery: A human-scale approach for informing better urban planning practices. *Landscape and Urban Planning*, 191: 103434, 2019a.
- Ye, Y., Zeng, W., Shen, Q., Zhang, X., and Lu, Y. The visual quality of streets: A human-centred continuous measurement based on machine learning algorithms and street view images. *Environment and Planning B: Urban Analytics and City Science*, 46(8):1439–1457, 2019b.
- Ying, K., Meng, F., Wang, J., Li, Z., Lin, H., Yang, Y., Zhang, H., Zhang, W., Lin, Y., Liu, S., et al. MMT-Bench: A Comprehensive Multimodal Benchmark for Evaluating Large Vision-Language Models Towards Multitask AGI. *arXiv preprint arXiv:2404.16006*, 2024.
- Zhang, F., Zhang, D., Liu, Y., and Lin, H. Representing place locales using scene elements. *Computers, Environment and Urban Systems*, 71:153–164, 2018.
- Zhang, H., Song, H., Li, S., Zhou, M., and Song, D. A Survey of Controllable Text Generation Using Transformer-based Pre-trained Language Models. *ACM Computing Surveys*, 56(3):1–37, 2023a.
- Zhang, X., Li, X., Sultani, W., Zhou, Y., and Wshah, S. Cross-View Geo-Localization via Learning Disentangled Geometric Layout Correspondence. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 3480–3488, 2023b.
- Zhu, S., Shah, M., and Chen, C. 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*, pp. 1162–1171, 2022.



## A. Implementation Details

Table 5 and Table 6 present the hyper-parameter settings and training details for the models. We conducted training and testing on Nvidia A800 (80G), with CUDA 12.1, PyTorch 2.0.0, and Transformers 4.33.0.

Table 5. The hyper-parameter settings of the proposed *GeoReasoner*.

| Hyper Params     | Value  |
|------------------|--------|
| Learning Rate    | 1e-5   |
| Total Batch Size | 64     |
| Weight Decay     | 0.1    |
| Warmup Ratio     | 0.01   |
| Optimizer        | AdamW  |
| Adam Beta1       | 0.9    |
| Adam Beta2       | 0.95   |
| LR Scheduler     | cosine |
| Model Max Length | 2048   |

Table 6. The training details of the proposed *GeoReasoner*.

|                  | Training Speed | Inference Latency | Num of Params |         | Flops |
|------------------|----------------|-------------------|---------------|---------|-------|
|                  |                |                   | Base Model    | LoRA    |       |
| LoRA1 (reason)   | 0.41 sample/s  | 1.560s            | 9.6B          | 112.19M | 71.9B |
| LoRA2 (location) | 0.63 sample/s  | 0.894s            | 9.6B          | 112.19M | 71.9B |

## B. Additional Qualitative Results

Additionally, we present the results of the *GeoReasoner* on alternative street-view images, depicted in Figure 9. Each street view image is annotated with the ground truth geographic location, along with the inference results from *GeoReasoner*. It can provide geographical predictions accompanied by reasonable explanations.

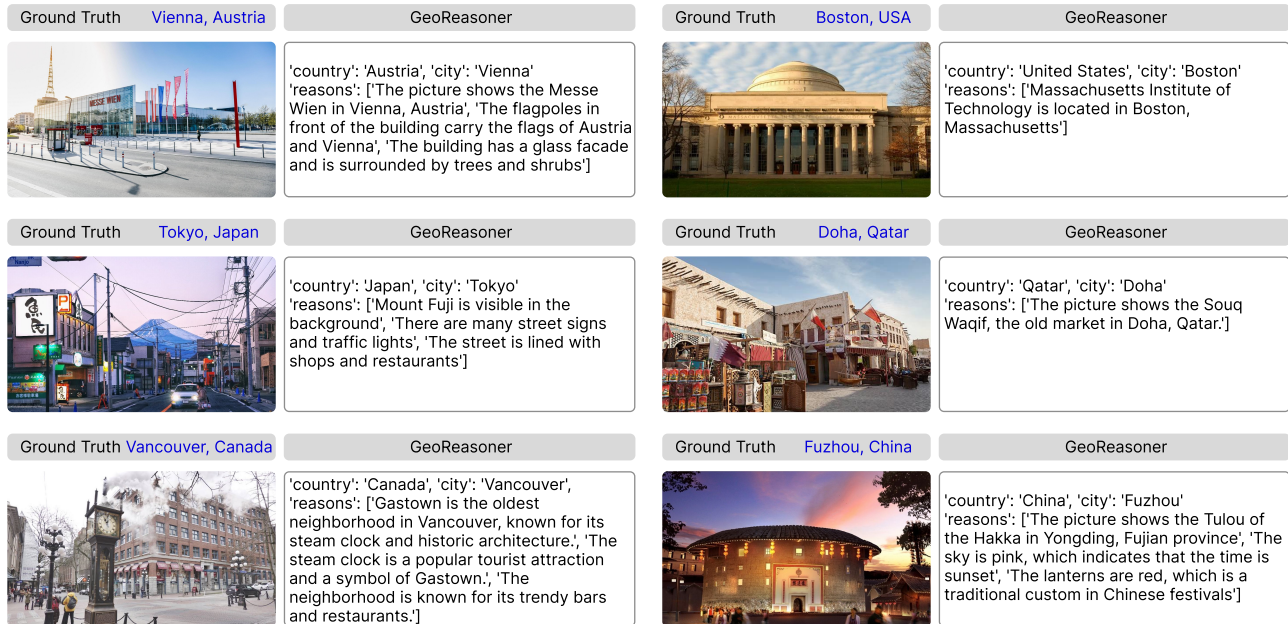


Figure 9. Additional qualitative results from the proposed *GeoReasoner*.