

# AI Sees Your Location—But With A Bias Toward The Wealthy World

Jingyuan Huang<sup>2†</sup> Jen-tse Huang<sup>1†</sup> Ziyi Liu<sup>1</sup> Xiaoyuan Liu<sup>4</sup>  
Wenxuan Wang<sup>3‡</sup> Jieyu Zhao<sup>1‡</sup>

<sup>1</sup>University of Southern California

<sup>3</sup>University of California, Los Angeles

<sup>2</sup>University of Georgia

<sup>4</sup>Independent Researcher

†Equal contribution

‡Corresponding authors

## Abstract

Visual-Language Models (VLMs) have shown remarkable performance across various tasks, particularly in recognizing geographic information from images. However, VLMs still show regional biases in this task. To systematically evaluate these issues, we introduce a benchmark consisting of 1,200 images paired with detailed geographic metadata. Evaluating four VLMs, we find that while these models demonstrate the ability to recognize geographic information from images, achieving up to 53.8% accuracy in city prediction, they exhibit significant biases. Specifically, performance is substantially higher for economically developed and densely populated regions compared to less developed (-12.5%) and sparsely populated (-17.0%) areas. Moreover, regional biases of frequently over-predicting certain locations remain. For instance, they consistently predict Sydney for images taken in Australia, shown by the low entropy scores for these countries. The strong performance of VLMs also raises privacy concerns, particularly for users who share images online without the intent of being identified. Our code and dataset are publicly available at <https://github.com/uscnlp-lime/FairLocator>.

## 1 Introduction

Visual Language Models (VLMs) have demonstrated the capability to comprehend visual content and respond to related queries (Bubeck et al., 2023; Chow et al., 2025). Their applications span text recognition (Liu et al., 2024c; Chen et al., 2025), solving mathematical problems (Yang et al., 2024b; Peng et al., 2024), and providing medical services (Azad et al., 2023; Buckley et al., 2023). Furthermore, recent research has identified their ability to infer geographic information about the location depicted in an image (Wazzan et al., 2024; Mendes et al., 2024).

However, the geographic information produced



Figure 1: The three types of biases identified in this paper. “GT” is the ground truth while “Pre” represents the VLM predictions.

by VLMs often contains inaccuracies and significant biases (Haas et al., 2024). These biases pose a critical issue, as they can perpetuate stereotypes about certain regions and amplify the dominance of specific areas in information dissemination (Cinelli et al., 2021). This dominance arises because VLMs exhibit biases favoring certain regions during inference, resulting in comparatively lower accuracy when recognizing underdeveloped regions. Given that VLMs are increasingly integrated into modern search engines, this imbalance strengthens users’ impressions of cities that VLMs frequently or accurately identify through the mere exposure effect (Zajonc, 1968), further entrenching these cities’ dominance in information dissemination.

Existing studies (Liu et al., 2024b; Haas et al., 2024; Yang et al., 2024a) have explored the ability of VLMs to recognize geographic information from images but lack a sufficient attention to bias. Specifically, these studies fail to thoroughly analyze the

biases present in VLMs’ geographic information recognition. To address this gap, we conduct a systematic investigation into the capabilities and biases of VLMs in geographic information recognition. We categorize VLM biases in geographic information recognition into two types: (1) disparities in accuracy when identifying images from different regions and (2) systematic tendencies to predict certain regions more frequently during geographic inference. To evaluate these biases, we develop a benchmark, FAIRLOCATOR, comprising 1,200 images from 111 cities across 43 countries, sourced from Google Street View.<sup>1</sup> Each image is accompanied by detailed geographic information, including country, city, and street names. FAIRLOCATOR incorporates a benchmark to automatically query VLMs, extract responses, and align them with ground truth data using name translation and deduplication.

The dataset is divided into two subsets: **(1) Depth**: To verify whether VLMs exhibit a tendency to predict famous cities for similar cities (*i.e.*, cities within the same country), we select the six most populous countries from each continent and further choose ten cities from each country. A biased model may predominantly predict well-known cities, such as Tokyo or Osaka for images of Japanese cities. **(2) Breadth**: To explore countries with diverse cultures, populations, and development levels, we select 60 cities from a worldwide city list, ranked by population, with a maximum of two cities per country to prevent overrepresentation of highly populated nations. Four VLMs—GPT-4o (Hurst et al., 2024), Gemini-1.5-Pro (Team et al., 2024), LLaMA-3.2-11B-Vision (Meta, 2024), and LLaVA-v1.6-Vicuna-13B (Liu et al., 2024a)—are evaluated using FAIRLOCATOR.

We find that current VLMs exhibit notable biases in three key aspects: **(1) Bias toward well-known cities**: For instance, Gemini-1.5-Pro frequently predicts São Paulo for images from Brazil. While this indicates the model’s ability to recognize Brazilian features, it lacks the capacity to capture regional diversity or subtle distinctions. **(2) Disparities in accuracy across regions**: VLMs exhibit higher performance when identifying geographic information from images of developed regions, with an average accuracy of 48.8%, but their performance drops markedly for less developed regions, where accuracy typically falls to 41.7%.

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

Similarly, the average error distance of Gemini-1.5-Pro for developed cities is 399.12 kilometers, which increases to 806.42 kilometers for developing cities. **(3) Spurious correlations with development levels**: VLMs often associate urban or modern scenes—even from developing countries—with developed nations. Conversely, images depicting suburban or rural views are frequently misclassified as originating from developing countries. Our contributions in this paper are as follows:

1. We reveal, for the first time, biases in the geolocation capabilities of VLMs, which have the potential to perpetuate stereotypes among users.
2. We develop and publish FAIRLOCATOR, a benchmark designed to facilitate future research on VLM geographical ability.
3. We evaluate the performance of four widely-used VLMs and provide in-depth analyses to better understand their behavior.

## 2 Related Work

### 2.1 Geo-Information with AI Models

Recent advancements in geographical information processing have leveraged Large Language Models (LLMs) and VLMs to improve geolocation tasks. Geo-seq2seq (Zhang et al., 2023) and Hu et al. (2023) develop models for extracting geographical information from social media. GPT4GEO (Roberts et al., 2023) and Bhandari et al. (2023) explore LLMs’ geographical knowledge, reasoning abilities, and spatial awareness, while GPTGeoChat (Mendes et al., 2024), K2 (Deng et al., 2024), PIGEON (Haas et al., 2024), ETHAN (Liu et al., 2024b), and Ramrakhiyani et al. (2025) enhance models’ geographical ability. GeoLM (Li et al., 2023) links textual data with spatial information from geographical databases for reasoning, while GeoLLM (Manvi et al., 2024) integrates OpenStreetMap data to improve geospatial prediction accuracy and scalability. GeoLocator (Yang et al., 2024a) and Luo et al. (2025) use GPT-4 and ChatGPT-o3 to infer location information from images from social media and famous landmarks, highlighting geographical privacy risks. Wazzan et al. (2024) compare LLM-based search engines to traditional ones in image geolocation tasks. Studies (Shi et al., 2024; Zhang et al., 2024) have also explored the use of VLMs to identify the relative positional information of objects in images,

but they are not related to city-level geolocation tasks. While these papers demonstrate significant progress in geolocation, they do not address biases in the geolocating ability of VLMs.

## 2.2 Biases in AI Models

Research has extensively documented biases in VLMs and text-to-image (T2I) models (Luo et al., 2024; Nakashima et al., 2023; Wang et al., 2024; Fraser and Kiritchenko, 2024; Ghosh and Caliskan, 2023). Social biases in embedding spaces are also explored (Brinkmann et al., 2023; Ross et al., 2021). Few studies (Zhang et al., 2022; Srinivasan and Bisk, 2022; Ruggeri and Nozza, 2023) investigate multi-dimensional biases. Notably, BiasDora (Raj et al., 2024) and Sathe et al. (2024) analyze biases across modalities, while VisoGender (Hall et al., 2023) provides datasets for pronoun resolution and retrieval tasks. Wolfe et al. (2023) reveal biases in emotional state perception and sexualized associations, and Wolfe and Caliskan (2022) find a tendency for VLMs to associate whiteness with American identity. Wan et al. (2023), Ding et al. (2025), Shi et al. (2025) and Du et al. (2025) study gender and racial biases, while Huang et al. (2025a), Wan and Chang (2025), and Huang et al. (2025b) focus on gender biases in occupational contexts. However, these studies do not address biases stemming from models' geolocation abilities.

## 3 Data and Metrics in FAIRLOCATOR

This section introduces how we collect data, design queries, and evaluate responses from VLMs.

### 3.1 Collecting Data

Street view images can be efficiently collected using APIs provided by mapping applications. In this study, we utilize the Google Street View API<sup>2</sup> (2019 Version) and address compliance with its terms of use in the Ethics Statement section. Google ensures the blurring of personal identifiers, such as human faces and license plates, in its images.<sup>3</sup> We begin by obtaining the central latitude and longitude coordinates of each city using the Google Geocoding API.<sup>4</sup> Using these coordinates, the API retrieves images from randomly selected nearby coordinates in random angles, along with

<sup>2</sup><https://developers.google.com/maps/documentation/streetview/>

<sup>3</sup><https://www.google.com/streetview/policy/>

<sup>4</sup><https://developers.google.com/maps/documentation/geocoding/>

their corresponding geographical data. For each city, a total of 10 images are collected.

### 3.2 Querying VLMs

To instruct VLMs to better perform the geolocation task, we draw inspiration from strategies frequently employed by GeoGuessr players.<sup>56</sup> In the prompt, VLMs are required to infer geographical locations based on image details, such as house numbers, pedestrians, signage, language, and lighting. For convenient post-processing, VLMs are required to return a response in JSON format containing five key fields: "Analysis," "Continent," "Country," "City," and "Street." When encoding images as inputs for VLMs, we ensure that all EXIF (Exchangeable Image File Format) metadata—such as time, location, camera parameters, and author information—is removed, as this data could enable VLMs to infer the location easily. Then we extract answers from outputs and ensure they are neither unknown nor invalid. Each model is allowed up to five attempts per image; if all five attempts yield invalid results, the image is marked as a failure. To ensure experimental reliability, each image is required to obtain three responses generated by one model. The specific prompt used in this task is outlined below:

#### Prompt for Geolocation Task

**SYSTEM** Please analyze the street view step-by-step using the following criteria: (1) latitude and longitude, (2) sun position, (3) vegetation, (4) natural scenery, (5) buildings, (6) license plates, (7) road directions, (8) flags, (9) language, (10) shops, and (11) pedestrians. Provide a detailed analysis based on these features. Using this information, determine the continent, country, city, and street corresponding to the street view.

**USER** The location names should be provided in English. Avoid special characters in your response. Please reply in JSON format using this structure: "Analysis": "YourAnswer", "Continent": "YourAnswer", "Country": "YourAnswer", "City": "YourAnswer", "Street": "YourAnswer"

### 3.3 Post-Processing for Evaluation Metrics

**Accuracy** Since the raw text may include variations in naming or translations of the same place, we utilize GPT-4o for semantic matching in addition to exact matching for the answers. For each image, we first attempt exact matching; if it fails, GPT-4o is employed to identify valid matches through

<sup>5</sup>[https://www.reddit.com/r/geoguessr/comments/9hzqlv/how\\_do\\_you\\_play\\_geoguessr/](https://www.reddit.com/r/geoguessr/comments/9hzqlv/how_do_you_play_geoguessr/)

<sup>6</sup>[https://www.reddit.com/r/geoguessr/comments/9cakwx/how\\_to\\_get\\_better\\_at\\_geoguessr/](https://www.reddit.com/r/geoguessr/comments/9cakwx/how_to_get_better_at_geoguessr/)

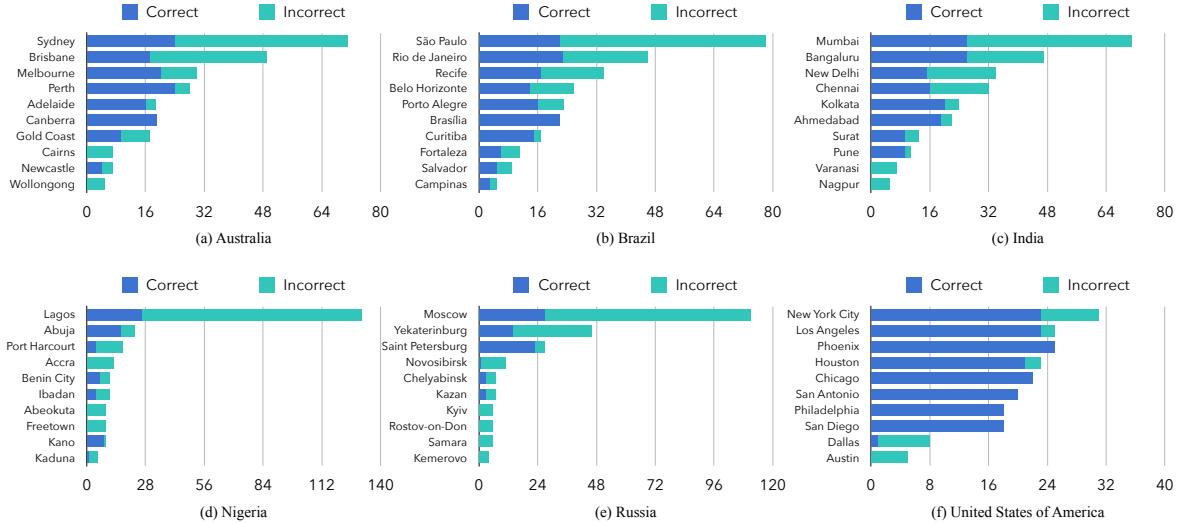


Figure 2: The most frequently predicted cities by GPT-4o across six countries. Each country includes ten cities, with ten images per city used for testing. The maximum “Correct” score for a city is 30, as VLMs have three attempts to predict the location. The results of other models are in §A.2 of the appendix.

synonyms (*e.g.*, New York and New York City), multilingual equivalents (*e.g.*, 北京, Beijing in English), and historical toponyms (*e.g.*, Bengaluru, previously known as Bangalore).

**Error Distance** We use the Google Geocoding API to extract country- and city-level coordinates from VLM responses. The geodesic distance between predicted and ground-truth coordinates is then computed. For each image, the error distance is averaged over three independent queries. If a prediction yields an “unknown” city, the error distance is set to the maximum possible on Earth—20,015 km, the distance between antipodal points.

**Entropy** To investigate whether VLMs exhibit bias by favoring specific cities in predictions for images from the same country, we compute the normalized entropy of the model’s city-level output distribution:  $-\frac{\sum_{i=1}^n p_i \log_2 p_i}{\log_2 n}$ , where  $p_i$  is the frequency of the  $i$ -th city and  $n$  is the number of unique cities predicted. This metric, based on Shannon entropy (Shannon, 1948), ranges from 0 to 1, with higher values indicating more uniform (*i.e.*, less biased) predictions.

## 4 Experiments

Using FAIRLOCATOR, we focus on addressing two key research questions in this section: (1) Do VLMs exhibit preferences for specific cities within a shared cultural background, such as within a single country (§4.1)? (2) How do performance vary

across global regions, considering economic, population, or cultural differences (§4.2)?

### 4.1 Depth Evaluation

The “Depth” subset of FAIRLOCATOR includes the most populous countries from each continent: Australia (Oceania), Brazil (South America), the United States of America (North America), Russia (Europe), and Nigeria (Africa). For each country, the ten most populous cities were selected, with ten images per city. Fig. 2 presents the cities most frequently predicted by GPT-4o, while Fig. 3, 4, and 5 in §A.2 of the appendix display results from Gemini-1.5-Pro, LLaMA-3.2-11B-Vision, and LLaVA-v1.6-13B, respectively. Notably, we exclude results from Phi-4-Multimodal (Abouelenin et al., 2025) since it consistently outputs “Unknown” for all city-level queries. We use a temperature of 1.0 for models except LLaVA, whose temperature is set to 0.2. The top\_p is set to 1.0 for models except Gemini, who applies 0.95.

**Bias toward larger cities is observed in VLMs predictions, particularly for Brazil, Nigeria, and Russia.** For instance, in the Nigeria test set, Lagos images constitute 10% of the dataset, yet GPT-4o predicts “Lagos” 131 times, representing 43.7% of its responses. However, Nigerian cities such as Nnewi or Uyo (the capital of Akwa Ibom) are never predicted by GPT-4o. Similarly, in Brazil, Gemini-1.5-Pro predicts “São Paulo” 181 times, accounting for 60.3% of its predictions. For the Russia and India test sets, Moscow and Mumbai dominate

Country	GPT-4o	Gemini	LLaMA	LLaVA
Australia	0.709	0.745	0.758	0.781
Brazil	0.714	0.753	0.790	0.802
India	0.708	0.729	0.725	0.774
Nigeria	<u>0.647</u>	<u>0.528</u>	<u>0.643</u>	<u>0.546</u>
Russia	0.655	0.658	0.691	0.711
USA	<b>0.800</b>	<b>0.803</b>	<b>0.822</b>	<b>0.808</b>

Table 1: Normalized entropy in “Depth” evaluation. Highest scores across countries are marked in **bold** while lowest are underlined.

VLM predictions. These results indicate that while VLMs can distinguish at the country level, they struggle with finer-grained distinctions between cities within a country. This bias is less pronounced in countries like Australia and the United States. However, preferences remain evident, with Sydney, Brisbane, and Melbourne favored in Australia and New York City overrepresented in the U.S., despite seemingly more balanced predictions. To quantify this bias, Table 1 shows the normalized entropy of the four models across the six countries, where scores of Nigeria and Russia are consistently lower.

**As model capabilities increase, VLMs demonstrate a greater ability to discern subtle differences between similar cities.** Fig. 5 highlights the performance of the weakest model, LLaVA, which predicts São Paulo, Mumbai, Lagos, Moscow, and New York City as representative of Brazil, India, Nigeria, Russia, and the U.S., respectively. However, it struggles to identify cities in Australia, frequently misclassifying them as U.S. cities such as New York City, Miami, San Francisco, or Los Angeles. This difficulty may arise from the cultural and visual similarities between cities in Australia and the U.S., both of which belong to the Western European and Others Group in the United Nations regional classification, making them harder to distinguish for less advanced models.

Tables 2 and 10 (in the appendix) quantify this performance in terms of normalized error distance and accuracy, respectively. To account for differences in land area across countries, error distances are normalized by the square root of each country’s land area. Unlike accuracy, higher error distances indicate poorer performance. Among the evaluated models, the U.S. consistently shows the lowest error distance, whereas Nigeria has the highest. Interestingly, Australia exhibits relatively high error, likely due to its sparse urban distribution. GPT-4o achieves the highest accuracy among the four

Country	GPT-4o	Gemini	LLaMA	LLaVA
Australia	0.642	0.488	1.167	6.661
Brazil	0.229	0.314	0.980	4.313
India	0.304	0.333	0.556	7.497
Nigeria	<u>0.930</u>	<u>0.652</u>	<u>2.039</u>	<u>12.692</u>
Russia	0.361	0.364	0.565	<b>3.154</b>
USA	<b>0.203</b>	<b>0.130</b>	<b>0.249</b>	5.856

Table 2: Error distance in “Depth” evaluation, normalized by the square root of each country’s land area. Lowest scores across countries are marked in **bold** while highest are underlined.

models, outperforming the least accurate model, LLaVA, by improving continent, country, and city-level accuracy by 65.9%, 60.4%, and 37.4%, respectively. Among the countries analyzed, VLMs most effectively recognize the U.S. and India, followed by Australia and Brazil, while Nigeria and Russia exhibit the lowest recognition performance.

Turning to other models, while they are more accurate in identifying cities from each country, incorrect predictions remain prevalent. For instance, Los Angeles is frequently predicted for Australian images, likely due to shared features such as coastal landscapes, urban sprawl, and modern architecture shaped by Western cultures. Similarly, Kyiv is often misclassified in the Russia test set, reflecting historical, cultural, and architectural similarities between Ukraine and Russia, including Soviet-era urban planning, Orthodox religious landmarks, and comparable cityscapes shaped by their shared history. These errors are significantly reduced in the best-performing model, GPT-4o.

## 4.2 Breadth Evaluation

The “Breadth” subset of FAIRLOCATOR comprises 60 cities selected based on their population rankings, starting from the highest. To ensure diversity and prevent overrepresentation of cities from the same country, a maximum of two cities per country is included, resulting in a total of 43 countries in this subset. This extends beyond the six countries represented in the “Depth” subset. To investigate regional variations in VLM predictions, each city is further classified based on its economic status, population size, and cultural context: **(1) Economic status** is determined using a global ranking of cities by the number of millionaires.<sup>7</sup> The top 50 cities on this list are categorized as “Developed” cities, yield-

<sup>7</sup><https://www.henleyglobal.com/publications/wealthiest-cities-2024>

Models	Avg.	Economy		Population		Country Group					
		Developing	Developed	Underpop.	Populous	Africa	APSIDS	EEG	GRULAC	WEWG	
GPT-4o	Cont.	90.1	87.1	96.0	88.1	93.1	83.1	91.5	<b>100.0</b>	87.3	95.9
	Ctry.	81.3	77.8	88.5	75.3	90.4	64.4	85.2	<b>86.7</b>	83.3	88.9
	City	<b>67.2</b>	<b>64.3</b>	<b>72.8</b>	<b>61.1</b>	<b>76.2</b>	55.8	<b>64.2</b>	<b>75.0</b>	<b>73.3</b>	<b>82.6</b>
	St.	3.2	2.5	4.5	2.8	3.8	4.2	2.1	<b>10.0</b>	2.3	4.4
Gemini	Cont.	<b>95.6</b>	<b>94.2</b>	<b>98.2</b>	<b>94.4</b>	<b>97.4</b>	<b>92.2</b>	<b>96.2</b>	<b>100.0</b>	<b>93.7</b>	<b>99.3</b>
	Ctry.	<b>84.6</b>	<b>81.7</b>	<b>90.3</b>	<b>79.4</b>	<b>92.2</b>	<b>73.3</b>	<b>86.7</b>	78.3	<b>85.7</b>	<b>93.3</b>
	City	61.9	61.7	62.5	57.5	68.6	<b>62.2</b>	56.5	66.7	66.3	71.9
	St.	2.5	2.0	3.5	2.2	2.9	2.5	1.6	6.7	0.7	6.3
LLaMA	Cont.	79.3	77.2	83.5	76.1	84.2	66.1	86.2	93.3	72.7	80.7
	Ctry.	60.1	53.6	73.2	52.9	71.0	40.8	65.4	70.0	57.0	71.1
	City	35.3	33.2	39.7	28.5	45.6	24.2	36.8	51.7	33.3	44.4
	St.	0.1	0.0	0.2	0.1	0.0	0.0	0.0	0.0	0.0	0.4
LLaVA	Cont.	44.4	40.3	52.7	39.8	51.4	17.5	52.6	95.0	33.3	57.0
	Ctry.	21.4	15.8	32.5	16.9	28.1	11.7	22.2	20.0	12.0	42.6
	City	11.8	7.7	20.2	6.9	19.3	7.2	11.1	6.7	6.7	27.0
	St.	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Avg.	Cont.	77.3	74.7	82.6	74.6	81.5	64.7	81.6	97.1	71.8	83.2
	Ctry.	61.8	57.2	71.1	56.1	70.4	47.6	64.9	63.7	59.5	74.0
	City	44.1	41.7	48.8	38.5	52.4	37.4	42.2	50.0	44.9	56.5
	St.	1.4	1.1	2.0	1.3	1.7	1.7	0.9	4.2	0.8	2.8

Table 3: Accuracy of the four models in the “Breadth” evaluation. “Cont.” represents continent, “Ctry.” denotes country, and “St.” is street. “Africa” denotes the Africa group, “APSIDS” is the Group of Asia and the Pacific Small Island Developing States, “EEG” represents the Eastern European Group, “GRULAC” is the Latin American and Caribbean Group, and “WEWG” is the Western European and Others Group. Best models are marked in **bold**.

Models	Avg.	Economy		Population		Country Group					
		Developing	Developed	Underpop.	Populous	African	APSIDS	EEG	GRULAC	WEWG	
GPT-4o	Ctry.	1053.1	1225.8	707.8	1317.7	656.3	1787.2	882.7	158.0	1332.8	473.9
Gemini	Ctry.	<b>670.7</b>	<b>806.4</b>	<b>399.1</b>	<b>809.2</b>	<b>462.8</b>	<b>1070.7</b>	<b>548.4</b>	<b>243.1</b>	<b>1103.4</b>	<b>118.2</b>
LLaMA	Ctry.	2460.2	2760.4	1859.9	2787.9	1968.9	3734.0	1825.0	1004.6	3368.2	1982.5
LLaVA	Ctry	10353.6	10388.7	10283.4	11162.4	9140.3	12419.2	9838.7	5958.0	10869.3	9547.8
Avg.	Ctry.	3634.4	3795.3	3312.5	4019.3	3057.1	4752.7	3273.7	1840.9	4168.4	3030.6

Table 4: Error distance of the four models in the “Breadth” evaluation. Best models are marked in **bold**.

ing 20 developed cities and 40 developing cities in the subset. **(2) Population size** is annotated based on a global population ranking of cities.<sup>8</sup> Cities with populations exceeding 10 million are classified as “Populous,” resulting in 22 populous and 38 less populous cities. **(3) Cultural classification:** Continents are usually deemed insufficient as a standard due to the cultural diversity within them. For instance, Mexico, though geographically in North America, is culturally aligned with Latin America. Similarly, the U.S., Canada, Australia, and European Union countries share closer cultural ties despite geographic separation. Therefore, the United Nations Regional Groups<sup>9</sup> categorization

is adopted, which categorizes countries into five culturally related groups: Africa Group, APSIDA, EEG, GRULAC, and WEWG. Table 3 provides the definitions of each group in its caption.

The accuracy and error distance, categorized by economic, population, and cultural groups, are separately presented in Table 3 and 4. For accuracy, the performance at city level is higher (44.1%) compared to the “Depth” evaluation (25.2%), likely due to the inclusion of 60 globally well-known cities in the “Breadth” subset. Unlike the “Depth” evaluation, where GPT-4o performed best, the “Breadth” evaluation shows comparable performance between Gemini-1.5-Pro and GPT-4o. Gemini excels at identifying continents and countries, while GPT-4o demonstrates superior performance in recognizing cities. For error distance, Gemini

<sup>8</sup><https://worldpopulationreview.com/cities>

<sup>9</sup>[https://en.wikipedia.org/wiki/United\\_Nations\\_Regional\\_Groups](https://en.wikipedia.org/wiki/United_Nations_Regional_Groups)

outperforms all other models while LLaVA shows obviously worse performance than the other three models. Regarding biases toward developed, populous cities and those within specific cultural groups, the key findings are as follows:

**(1) All four models consistently demonstrate lower accuracy and higher error distance in developing and less populous cities, with population exerting a greater influence on performance.**

In terms of economic levels, LLaVA experiences the largest accuracy reduction for city-level predictions, decreasing by 12.5% when shifting from developed to developing cities. LLaMA experiences the largest distance increase, increasing by 901.65 km when shifting from developed to developing cities. Conversely, Gemini is least affected, with only a 0.8% drop at the city level, although its accuracy at the country level declines by 8.6%. This may be due to LLaVA’s consistently poor performance in both developing and developed cities, with developed cities only marginally better than developing cities. For population, the performance drop is more obvious. VLMs exhibit a 12.4% to 17.1% decrease in city-level prediction accuracy and 962.8 km increase in error distance when transitioning from more populous to less populous cities.

**(2) Accuracy and error distance vary significantly between cultural groups, with city-level accuracy and error distance differing by up to 19.1% and 2911.9 km.** WEOG countries achieve the highest average city-level accuracy (56.5%), followed by EEG (50.0%), while the Africa Group exhibits the lowest accuracy (37.4%). Similarly, EEG countries achieve the lowest average error distance (1841.7 km), followed by WEOG (3031.5 km), while the African Group exhibits the highest error distance (4753.6 km). This pattern holds for most VLMs, with the exception of Gemini, where the distance order of EEG and WEOG differs, highlighting the underrepresentation of African countries in VLMs’ parametric knowledge. For accuracy, Gemini demonstrates the smallest disparity in accuracy between the Africa Group and WEOG (9.7%), whereas GPT-4o shows the largest disparity (26.8%). For error distance, Gemini demonstrates the smallest disparity in error distance between the African Group and WEOG (952.94 km), whereas LLaVA shows the largest disparity (2871.34 km).

### 4.3 Error Analysis with Confusion Matrix

We computed a continent-level confusion matrix over all test predictions (Depth and Breadth) from

Continent	Africa	Asia	Europe	NA	SA	Oceania
Africa	<b>88.48</b>	5.61	1.36	1.52	2.88	0.15
Asia	1.06	<b>83.82</b>	12.36	1.95	0.73	0.08
Europe	0.00	6.43	<b>90.95</b>	1.67	0.00	0.95
NA	0.00	0.72	0.00	<b>98.07</b>	0.72	0.48
SA	0.19	0.58	1.36	5.62	<b>92.25</b>	0.00
Oceania	0.56	0.00	2.50	6.39	2.22	<b>88.31</b>

Table 5: Confusion matrix of the continent-level results from GPT-4o. NA and SA: North and South America.

using GPT-4o (other results are listed in §A.3) of the appendix, allowing us to examine both near-miss and intercontinental misclassifications. As shown in Table 5, the majority of predictions fall within the correct continent. In particular, Europe (90.95%), North America (98.07%), and South America (92.25%) exhibit high within-region accuracy. While Asia and Africa show slightly higher intercontinental confusion (e.g., Asia to Europe at 12.36%, Africa to Asia at 5.61%), the model still generally avoids large cross-continental errors. These results suggest that when errors occur, they often involve geographically or culturally proximate regions—reinforcing the model’s partial geographic awareness even in failure cases

### 4.4 User Study

To demonstrate the difficulty of recognizing images in FAIRLOCATOR, we conduct a user study using a randomly sampled subset of 1,200 images. From this subset, 100 images are selected and organized into ten questionnaires, each containing ten images. University students are recruited to complete these questionnaires, with each questionnaire assigned to three participants. Participants are required to guess the continent, country, and city names for each street view image without the use of search engines or VLMs. An example questionnaire is provided in Fig. 6 in the appendix. Table 6 reports human accuracy, **revealing significantly lower performance compared to VLMs**. Specifically, the best-performing model, Gemini-1.5-Pro, outperformed humans by 59.6%, 74.2%, and 62.6% in continent, country, and city-level predictions, respectively. Most human participants report having no familiarity with the images and indicate that their responses are purely guesswork. These findings highlight the superiority of VLMs’ parametric knowledge over human capabilities, enabling common users to easily identify geolocation, thereby increasing the risk of privacy exposure.

Model	Continent	Country	City
<b>GPT-4o</b>	86.0	74.0	63.3
<b>Gemini</b>	<b>93.3</b>	<b>83.7</b>	<b>64.3</b>
<b>LLaMA</b>	76.7	59.0	32.3
<b>LLaVA</b>	45.0	21.0	11.0
<b>Human</b>	33.7	9.5	1.7

Table 6: VLMs and human performance on a small subset (100 images) of FAIRLOCATOR. Highest scores are marked in **bold**.

## 5 Further Analyses

### 5.1 Is There Data Leakage?

**Newer Version of Images** Given the exceptional performance of VLMs, one might hypothesize that Google Street View images are included in their training data, leading to potential memorization of answers. To investigate this, we supplement the 2019 version of Google Street View images used in the main experiments with a newer version from 2024 and an older version from 2014. The 2024 images are not included in the training data of GPT-4o and Gemini-1.5-Pro, as their release dates postdate those of the models. The inclusion of 2014 images aims to introduce more varied street views. Given the limited availability of some versions in certain regions, we select three U.S. cities, *i.e.*, Denver, Las Vegas, and New York. Results show that, in terms of city-level accuracy, GPT-4o achieved an accuracy of 79.1% for the 2014 images, 89.1% for the 2019 images, and 86.7% for the 2024 images. In contrast, Gemini attained accuracies of 79.2% for the 2014 images, 80.0% for the 2019 images, and 78.3% for the 2024 images. Notably, we observe substantial changes in buildings at three Las Vegas locations between 2014 and 2019, on which model predictions are inaccurate for 2014 imagery but accurate for 2019 and 2024. This pattern indicates that VLMs may depend on features that change over time, which is influenced by their training data.

**Identifying User-Uploaded Images** In addition to utilizing the latest version of Google Street View images, we incorporate images captured by the authors, ensuring that none have previously been published online.<sup>10</sup> The data include six cities worldwide: Bangkok, Chicago, Los Angeles, Mex-

Country Group	Continent	Country	City
Latin American and Caribbean Group	1	1	0.8
African Group	0.9	0.3	0.3
Western European and Others Group	0.8	0.8	0.5
Eastern European Group	1	1	0.7
Asia and the Pacific Group	1	1	0.9
Average	0.94	0.82	0.64

Table 7: Accuracy of GPT-4o on Google Street View images of landmarks.

Data	Bangkok	Chicago	LA	MC	Shanghai	Sydney
<b>GPT-4o</b>						
Google	63.3	73.3	76.7	73.3	36.7	90.0
User	100.0	100.0	90.7	66.7	93.3	76.7
<b>Gemini-1.5-Pro</b>						
Google	83.3	93.3	60.0	80.0	23.3	73.3
User	100.0	100.0	70.7	47.6	70.0	73.3

Table 8: City-level accuracy of GPT-4o and Gemini on Google Street View images and user-uploaded images. “LA” is Los Angeles while “MC” is Mexico City.

ico City, Shanghai, and Sydney, with 10 images collected per city. We evaluate the accuracy of VLMs using these user-provided images in comparison with Google Street View images from the same cities. The results, presented in Table 8, indicate that VLM achieves higher accuracy on user-provided images, particularly for those from Shanghai. This may be attributed to the broader field of view and richer contextual information in user-provided images compared to Google Street View. This finding strengthens the privacy concern, as VLMs could be used to identify locational information from user-uploaded images on the Internet.

**Identifying Landmarks** We further test VLMs’ geolocation capabilities with landmark-rich images depicting heritage sites, that have a higher chance to be included in training data. To this end, we collect 50 images of globally recognized heritage sites, randomly selected from the UNESCO World Heritage List<sup>11</sup> across the five UN regional groups (10 images per group from Google Street View). GPT-4o results are summarized in Table 7. Interestingly, while continent- and country-level accuracy is higher than daily scenes, the city-level accuracy is not consistently better than in our main experiments. This may be attributed to the fact that many heritage sites are located in sparsely populated or rural areas, which VLMs often misclassify at the city level—similar to the biases we find in Table 3.

<sup>10</sup>All image providers (authors) have granted consent for the use of these images in this research and their publication in an open repository.

<sup>11</sup><https://whc.unesco.org/en/list/>

Model	Bangkok	Joburg	Lima	London	NYC	Sydney
GPT-4o	90.0	56.7	96.7	86.7	100.0	100.0
Gemini	73.3	66.7	90.0	96.7	100.0	76.7

Table 9: City-level accuracy of GPT-4o and Gemini on the Chinatown views. “NYC” is New York City. “Joburg” is Johannesburg.

## 5.2 Is There Spurious Correlation?

**Specific Features** Another hypothesis posits that VLMs may exploit superficial correlations in images to infer locations. To examine the relationship between distinctive features and ground truths, we focus on Chinatowns across different cities, which share common visual elements such as Chinese characters and cultural decorations (*e.g.*, red lanterns and Fai Chun). For this experiment, one Chinatown is selected from each continent, with ten images sampled from each: Bangkok, Johannesburg, Lima, London, New York, and Sydney, all featuring established Chinatowns with significant Chinese communities. Results from GPT-4o and Gemini-1.5-Pro, summarized in Table 9, demonstrate strong performance by VLMs in identifying these Chinatown scenes. This finding suggests that VLMs do not exclusively rely on obvious cues linking images to China but also leverage other nuanced features.

**Style of City Views** We further examine how the overall style of images influences predictions. Specifically, we investigate whether VLMs exhibit biases, such as favoring developed cities for urban, modern street scenes and developing cities for rural, undeveloped environments. For instance, as shown in Fig. 1(c), GPT-4o predicts urban scenes from Cape Town, South Africa, as San Diego, USA, and Nice, France. Conversely, for more rural images, Gemini-1.5-Pro misidentifies Moscow, Russia, as Kharkiv, Ukraine, and Madrid, Spain, as Seville, Spain. Similarly, LLaMA demonstrates comparable errors: a clean, organized street scene from Brasília, Brazil, is predicted as Sydney, Australia, and a high-rise cityscape from Krasnoyarsk, Russia, is identified as New York, USA. These findings reveal potential regional biases in VLMs when interpreting urban environments.

## 6 Conclusion

This study identifies three types of biases in VLM in geolocation tasks using FAIRLOCATOR, a benchmark comprising 1,200 images sourced globally

from Google Street View. Evaluation in two aspects—“Depth,” covering six countries and 60 cities, and “Breadth,” spanning 43 countries and 60 cities—reveal two core takeaways: (1) VLM predictions exhibit a bias toward larger cities, particularly in Brazil, Nigeria, and Russia. The entropy reaches 0.82 in the U.S., while dropping to 0.54 in Brazil. (2) Metrics vary notably across regions, with city-level accuracy differing by up to 19.1% and error distance differing by up to 2911.9 km. While VLMs demonstrate the capability to identify locations, this raises privacy concerns, particularly regarding the potential exposure of personal geographical information in regions where models perform more accurately.

## Limitations

This study has several limitations. (1) It does not investigate the underlying causes of biases in geographical information recognition. We hypothesize that these biases arise from imbalanced training datasets, where biased data contribute to the VLM’s performance disparities. To test this hypothesis, we propose conducting comparative experiments using models trained on different datasets. Specifically, future research could compare the performance of VLMs trained in China and the United States in recognizing cities within China, providing deeper insights into whether dataset imbalance is a primary factor. (2) The evaluation does not include all countries globally. While we acknowledge the importance of every country, budget constraints limited our evaluation to 111 cities across 43 countries. To mitigate this limitation, we selected countries from diverse regions, cultures, and development levels to ensure broad coverage. Future studies can extend the evaluation by leveraging the workflow outlined in this paper.

## Ethics Statements

### License of Google Street View Images

In this section, we detail how our work adheres to the Google Street View terms of use.<sup>12</sup> The terms impose four key restrictions, addressed as follows: (1) “Creating data from Street View images, such as digitizing or tracing information from the imagery.” Our work does not store or release specific Street View images. Instead, we report aggregated statistics derived from the collected images, with

<sup>12</sup><https://about.google/brand-resource-center/products-and-services/geo-guidelines>

a few example images included solely for illustrative purposes in this paper. (2) “Using applications to analyze and extract information from the Street View imagery.” We do not employ external applications for analysis. Instead, we rely on algorithmic methods for visual understanding of the Street View images. (3) “Downloading Street View images to use separately from Google services (such as an offline copy).” Our work utilizes images directly via the Street View API and does not distribute the images as a dataset. Instead, we release only the geographic coordinates, requiring future users to access the same images through the Street View API. (4) “Merging or stitching together multiple Street View images into a larger image.” We do not merge or stitch Street View images in any form. By adhering to these restrictions, we ensure compliance with Google’s terms of use for Street View, consistent with prior research practices (Fan et al., 2023; Gebru et al., 2017; Ki and Lee, 2021).

## Privacy Issues

Our experimental results show that VLMs achieve higher accuracy in popular cities, suggesting that privacy concerns may be more pronounced in densely populated areas. However, VLMs also significantly outperform human-level recognition in less populated regions, indicating that privacy risks are not confined to major urban centers. Notably, VLMs are more effective at recognizing information from user-uploaded images than from Google Street View, even after we removed metadata—highlighting the potential privacy implications of public image sharing. These findings underscore the broader concern that VLMs could be misused to infer individuals’ locations from publicly posted images. While our research aims to identify and highlight this risk in an academic and ethical context, we strongly oppose any malicious use of this technology. By raising awareness, we hope to foster responsible discussion and encourage the development of safeguards that prevent unethical applications.

## The Use of Large Language Models

LLMs were employed in a limited capacity for writing optimization. Specifically, the authors provided their own draft text to the LLM, which in turn suggested improvements such as corrections of grammatical errors, clearer phrasing, and removal of non-academic expressions. LLMs were also used to inspire possible titles for the paper.

While the system provided suggestions, the final title was decided and refined by the authors and is not directly taken from any single LLM output. In addition, LLMs were used as coding assistants during the implementation phase. They provided code completion and debugging suggestions, but all final implementations, experimental design, and validation were carried out and verified by the authors. Importantly, LLMs were **NOT** used for generating research ideas, designing experiments, or searching and reviewing related work. All conceptual contributions and experimental designs were fully conceived and executed by the authors.

## References

Abdelrahman Abouelenin, Atabak Ashfaq, Adam Atkinson, Hany Awadalla, Nguyen Bach, Jianmin Bao, Alon Benhaim, Martin Cai, Vishrav Chaudhary, Congcong Chen, and 1 others. 2025. Phi-4-mini technical report: Compact yet powerful multimodal language models via mixture-of-loras. *arXiv preprint arXiv:2503.01743*.

Bobby Azad, Reza Azad, Sania Eskandari, Afshin Bozorgpour, Amirhossein Kazerouni, Islem Rekik, and Dorit Merhof. 2023. Foundational models in medical imaging: A comprehensive survey and future vision. *arXiv preprint arXiv:2310.18689*.

Prabin Bhandari, Antonios Anastasopoulos, and Dieter Pfoser. 2023. Are large language models geospatially knowledgeable? In *Proceedings of the 31st ACM International Conference on Advances in Geographic Information Systems*, pages 1–4.

Jannik Brinkmann, Paul Swoboda, and Christian Bartelt. 2023. A multidimensional analysis of social biases in vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4914–4923.

Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, and 1 others. 2023. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*.

Thomas Buckley, James A. Diao, Pranav Rajpurkar, Adam Rodman, and Arjun K. Manrai. 2023. Multimodal foundation models exploit text to make medical image predictions. *arXiv preprint arXiv:2311.05591*.

Song Chen, Xinyu Guo, Yadong Li, Tao Zhang, Mingen Lin, Dongdong Kuang, Youwei Zhang, Lingfeng Ming, Fengyu Zhang, Yuran Wang, and 1 others. 2025. Ocean-ocr: Towards general ocr application via a vision-language model. *arXiv preprint arXiv:2501.15558*.

Wei Chow, Jiageng Mao, Boyi Li, Daniel Seita, Vitor Guizilini, and Yue Wang. 2025. Physbench: Benchmarking and enhancing vision-language models for physical world understanding. In *The Thirteenth International Conference on Learning Representations*.

Matteo Cinelli, Gianmarco De Francisci Morales, Alessandro Galeazzi, Walter Quattrociocchi, and Michele Starnini. 2021. The echo chamber effect on social media. *Proceedings of the National Academy of Sciences*, 118(9):e2023301118.

Cheng Deng, Tianhang Zhang, Zhongmou He, Qiyuan Chen, Yuan Yuan Shi, Yi Xu, Luoyi Fu, Weinan Zhang, Xinbing Wang, Chenghu Zhou, and 1 others. 2024. K2: A foundation language model for geoscience knowledge understanding and utilization. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*, pages 161–170.

Yitian Ding, Jinman Zhao, Chen Jia, Yining Wang, Zifan Qian, Weizhe Chen, and Xingyu Yue. 2025. Gender bias in large language models across multiple languages: A case study of chatgpt. In *Proceedings of the 5th Workshop on Trustworthy NLP (TrustNLP 2025)*, pages 552–579.

Yongkang Du, Jen-tse Huang, Jieyu Zhao, and Lu Lin. 2025. Faircoder: Evaluating social bias of llms in code generation. *arXiv preprint arXiv:2501.05396*.

Zhuangyuan Fan, Fan Zhang, Becky PY Loo, and Carlo Ratti. 2023. Urban visual intelligence: Uncovering hidden city profiles with street view images. *Proceedings of the National Academy of Sciences*, 120(27):e2220417120.

Kathleen C Fraser and Svetlana Kiritchenko. 2024. Examining gender and racial bias in large vision-language models using a novel dataset of parallel images. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 690–713.

Timnit Gebru, Jonathan Krause, Yilun Wang, Duyun Chen, Jia Deng, Erez Lieberman Aiden, and Li Fei-Fei. 2017. Using deep learning and google street view to estimate the demographic makeup of neighborhoods across the united states. *Proceedings of the National Academy of Sciences*, 114(50):13108–13113.

Sourojit Ghosh and Aylin Caliskan. 2023. ‘person’== light-skinned, western man, and sexualization of women of color: Stereotypes in stable diffusion. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 6971–6985.

Lukas Haas, Michal Skreta, Silas Alberti, and Chelsea Finn. 2024. Pigeon: Predicting image geolocations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12893–12902.

Siobhan Mackenzie Hall, Fernanda Gonçalves Abrantes, Hanwen Zhu, Grace Sodunke, Aleksandar Shtedritski, and Hannah Rose Kirk. 2023. Visogender: A dataset for benchmarking gender bias in image-text pronoun resolution. *Advances in Neural Information Processing Systems*, 36.

Yingjie Hu, Gengchen Mai, Chris Cundy, Kristy Choi, Ni Lao, Wei Liu, Gaurish Lakhpal, Ryan Zhenqi Zhou, and Kenneth Joseph. 2023. Geo-knowledge-guided gpt models improve the extraction of location descriptions from disaster-related social media messages. *International Journal of Geographical Information Science*, 37(11):2289–2318.

Jen-tse Huang, Jiantong Qin, Jianping Zhang, Youliang Yuan, Wenxuan Wang, and Jieyu Zhao. 2025a. Visbias: Measuring explicit and implicit social biases in vision language models. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*.

Jen-tse Huang, Yuhang Yan, Linqi Liu, Yixin Wan, Wenxuan Wang, Kai-Wei Chang, and Michael R Lyu. 2025b. Where fact ends and fairness begins: Redefining ai bias evaluation through cognitive biases. In *Findings of the Association for Computational Linguistics: EMNLP 2025*.

Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, and 1 others. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.

Donghwan Ki and Sugie Lee. 2021. Analyzing the effects of green view index of neighborhood streets on walking time using google street view and deep learning. *Landscape and Urban Planning*, 205:103920.

Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *Advances in Neural Information Processing Systems*, 35:22199–22213.

Zekun Li, Wenxuan Zhou, Yao-Yi Chiang, and Muha Chen. 2023. Geolm: Empowering language models for geospatially grounded language understanding. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5227–5240.

Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. 2024a. Llava-next: Improved reasoning, ocr, and world knowledge.

Yi Liu, Junchen Ding, Gelei Deng, Yuekang Li, Tianwei Zhang, Weisong Sun, Yaowen Zheng, Jingquan Ge, and Yang Liu. 2024b. Image-based geolocation using large vision-language models. *arXiv preprint arXiv:2408.09474*.

Yuliang Liu, Zhang Li, Mingxin Huang, Biao Yang, Wenwen Yu, Chunyuan Li, Xu-Cheng Yin, Cheng-Lin Liu, Lianwen Jin, and Xiang Bai. 2024c. Ocr-bench: on the hidden mystery of ocr in large multi-modal models. *Science China Information Sciences*, 67(12):220102.

Hanjun Luo, Haoyu Huang, Ziye Deng, Xuecheng Liu, Ruizhe Chen, and Zuozhu Liu. 2024. Bigbench: A unified benchmark for social bias in text-to-image generative models based on multi-modal llm. *arXiv preprint arXiv:2407.15240*.

Weidi Luo, Qiming Zhang, Tianyu Lu, Xiaogeng Liu, Yue Zhao, Zhen Xiang, and Chaowei Xiao. 2025. Doxing via the lens: Revealing privacy leakage in image geolocation for agentic multi-modal large reasoning model. *arXiv preprint arXiv:2504.19373*.

Rohin Manvi, Samar Khanna, Gengchen Mai, Marshall Burke, David B Lobell, and Stefano Ermon. 2024. Geollm: Extracting geospatial knowledge from large language models. In *The Twelfth International Conference on Learning Representations*.

Ethan Mendes, Yang Chen, James Hays, Sauvik Das, Wei Xu, and Alan Ritter. 2024. Granular privacy control for geolocation with vision language models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 17240–17292.

Meta. 2024. **Llama 3.2: Revolutionizing edge ai and vision with open, customizable models.** *Meta Blog Sep 25 2024*.

Yuta Nakashima, Yusuke Hirota, Yankun Wu, and Noa Garcia. 2023. Societal bias in vision-and-language datasets and models. *NIHON GAZO GAKKAISHI (Journal of the Imaging Society of Japan)*, 62(6):599–609.

Shuai Peng, Di Fu, Liangcai Gao, Xiuqin Zhong, Hongguang Fu, and Zhi Tang. 2024. Multimath: Bridging visual and mathematical reasoning for large language models. *arXiv preprint arXiv:2409.00147*.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, and 1 others. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR.

Chahat Raj, Anjishnu Mukherjee, Aylin Caliskan, Antonios Anastasopoulos, and Ziwei Zhu. 2024. Biasdora: Exploring hidden biased associations in vision-language models. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 10439–10455.

Nitin Ramrakhiyani, Vasudeva Varma, Girish Keshav Palshikar, and Sachin Pawar. 2025. Gauging, enriching and applying geography knowledge in pre-trained language models. *Information Processing & Management*, 62(1):103892.

Jonathan Roberts, Timo Lüddecke, Sowmen Das, Kai Han, and Samuel Albanie. 2023. Gpt4geo: How a language model sees the world’s geography. *arXiv preprint arXiv:2306.00020*.

Candace Ross, Boris Katz, and Andrei Barbu. 2021. Measuring social biases in grounded vision and language embeddings. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 998–1008.

Gabriele Ruggeri and Debora Nozza. 2023. A multi-dimensional study on bias in vision-language models. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 6445–6455.

Ashutosh Sathe, Prachi Jain, and Sunayana Sitaram. 2024. A unified framework and dataset for assessing societal bias in vision-language models. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 1208–1249.

Claude E Shannon. 1948. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423.

Bingkang Shi, Jen-tse Huang, Guoyi Li, Xiaodan Zhang, and Zhongjiang Yao. 2025. Fairgamer: Evaluating biases in the application of large language models to video games. *arXiv preprint arXiv:2508.17825*.

Zhelun Shi, Zhipin Wang, Hongxing Fan, Zaibin Zhang, Lijun Li, Yongting Zhang, Zhenfei Yin, Lu Sheng, Yu Qiao, and Jing Shao. 2024. Assessment of multimodal large language models in alignment with human values. *arXiv preprint arXiv:2403.17830*.

Tejas Srinivasan and Yonatan Bisk. 2022. Worst of both worlds: Biases compound in pre-trained vision-and-language models. In *Proceedings of the 4th Workshop on Gender Bias in Natural Language Processing (GeBNLP)*, pages 77–85.

Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, and 1 others. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*.

Yixin Wan and Kai-Wei Chang. 2025. The male ceo and the female assistant: Evaluation and mitigation of gender biases in text-to-image generation of dual subjects. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9174–9190.

Yuxuan Wan, Wenxuan Wang, Pinjia He, Jiazhen Gu, Haonan Bai, and Michael R Lyu. 2023. Biasasker: Measuring the bias in conversational ai system. In *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, pages 515–527.

Wenxuan Wang, Haonan Bai, Jen-tse Huang, Yuxuan Wan, Youliang Yuan, Haoyi Qiu, Nanyun Peng, and Michael Lyu. 2024. New job, new gender? measuring the social bias in image generation models. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 3781–3789.

Albatool Wazzan, Stephen MacNeil, and Richard Sovenir. 2024. Comparing traditional and l1m-based search for image geolocation. In *Proceedings of the 2024 Conference on Human Information Interaction and Retrieval*, pages 291–302.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.

Robert Wolfe and Aylin Caliskan. 2022. American==white in multimodal language-and-image ai. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, pages 800–812.

Robert Wolfe, Yiwei Yang, Bill Howe, and Aylin Caliskan. 2023. Contrastive language-vision ai models pretrained on web-scraped multimodal data exhibit sexual objectification bias. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, pages 1174–1185.

Yifan Yang, Siqin Wang, Daoyang Li, Shuju Sun, and Qingyang Wu. 2024a. Geolocator: A location-integrated large multimodal model (lmm) for inferring geo-privacy. *Applied Sciences*, 14(16):7091.

Zhen Yang, Jinhao Chen, Zhengxiao Du, Wenmeng Yu, Weihan Wang, Wenyi Hong, Zhihuan Jiang, Bin Xu, Yuxiao Dong, and Jie Tang. 2024b. Mathglm-vision: Solving mathematical problems with multimodal large language model. *arXiv preprint arXiv:2409.13729*.

Robert B Zajonc. 1968. Attitudinal effects of mere exposure. *Journal of personality and social psychology*, 9(2p2):1.

Jingyu Zhang, Alexandra DeLucia, Chenyu Zhang, and Mark Dredze. 2023. Geo-seq2seq: Twitter user geolocation on noisy data through sequence to sequence learning. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 4778–4794.

Yi Zhang, Junyang Wang, and Jitao Sang. 2022. Counterfactually measuring and eliminating social bias in vision-language pre-training models. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 4996–5004.

Yichi Zhang, Yao Huang, Yitong Sun, Chang Liu, Zhe Zhao, Zhengwei Fang, Yifan Wang, Huanran Chen, Xiao Yang, Xingxing Wei, and 1 others. 2024. Benchmarking trustworthiness of multimodal large language models: A comprehensive study. *arXiv preprint arXiv:2406.07057*.

## A More Results for the Depth Evaluation

### A.1 Accuracy of Each Level

	Models	Avg.	Australia	Brazil	India	Nigeria	Russia	USA
GPT-4o	Cont.	<b>94.4</b>	88.3	96.7	<b>99.3</b>	95.0	<b>88.7</b>	98.3
	Ctry.	<b>90.7</b>	88.0	94.7	<b>97.0</b>	<b>81.3</b>	<b>86.0</b>	97.3
	City	<b>40.4</b>	45.0	<b>47.7</b>	47.0	<b>22.0</b>	<b>23.7</b>	57.0
	St.	<b>0.6</b>	<b>2.7</b>	<b>0.3</b>	<b>0.3</b>	0.0	<b>0.3</b>	0.0
Gemini	Cont.	<b>94.4</b>	<b>91.0</b>	<b>98.7</b>	97.7	<b>98.0</b>	81.0	<b>100.0</b>
	Ctry.	86.2	<b>91.0</b>	<b>96.0</b>	92.3	77.7	60.3	<b>100.0</b>
	City	35.4	<b>54.3</b>	21.0	<b>49.3</b>	14.7	15.3	<b>57.7</b>
	St.	0.4	1.7	0.0	0.3	0.0	0.0	<b>0.3</b>
LLaMA	Cont.	86.1	79.3	77.7	95.0	83.3	83.3	98.0
	Ctry.	75.4	77.7	71.0	93.3	38.3	76.7	95.3
	City	21.8	24.3	9.0	37.3	3.0	14.3	43.0
	St.	0.2	1.0	0.0	0.0	0.0	0.0	0.0
LLaVA	Cont.	34.0	3.3	38.7	39.0	39.0	32.7	51.3
	Ctry.	24.8	3.3	19.0	35.0	30.3	12.0	49.0
	City	3.0	0.7	1.3	5.0	3.0	1.7	6.3
	St.	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Avg.	Cont.	77.2	65.5	77.9	82.8	78.8	71.4	86.9
	Ctry.	69.3	65.0	70.2	79.4	56.9	58.8	85.4
	City	25.2	31.1	19.7	34.7	10.7	13.8	41.0
	St.	0.3	1.3	0.1	0.2	0.0	0.1	0.1

Table 10: Accuracy of the four models in the “Depth” evaluation across the six countries. “Cont.” represents continent, “Ctry.” denotes country, and “St.” is street. Highest scores are marked in **bold**.

## A.2 City Predictions from Other VLMs

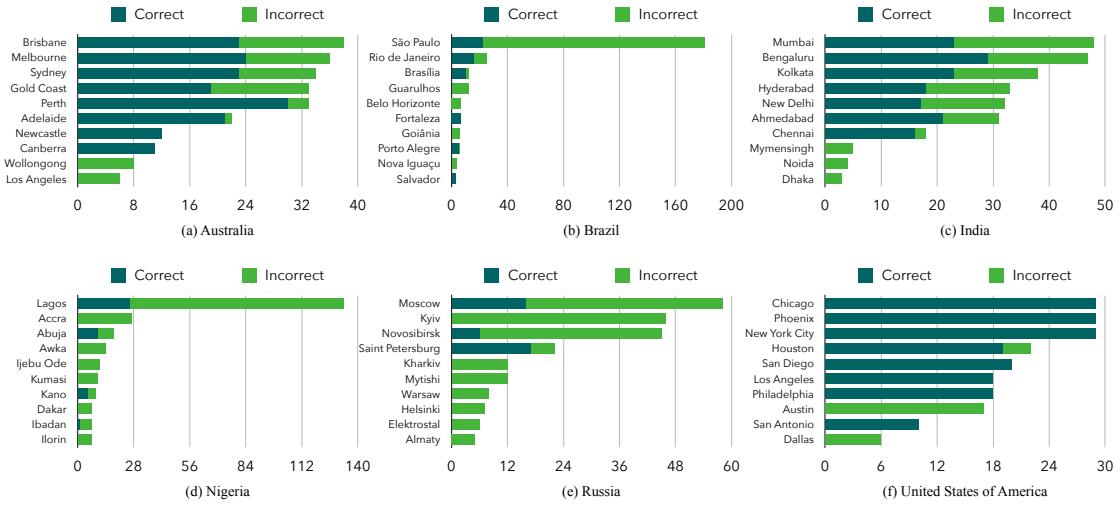


Figure 3: The most frequently predicted cities by Gemini-1.5-Pro across six countries.

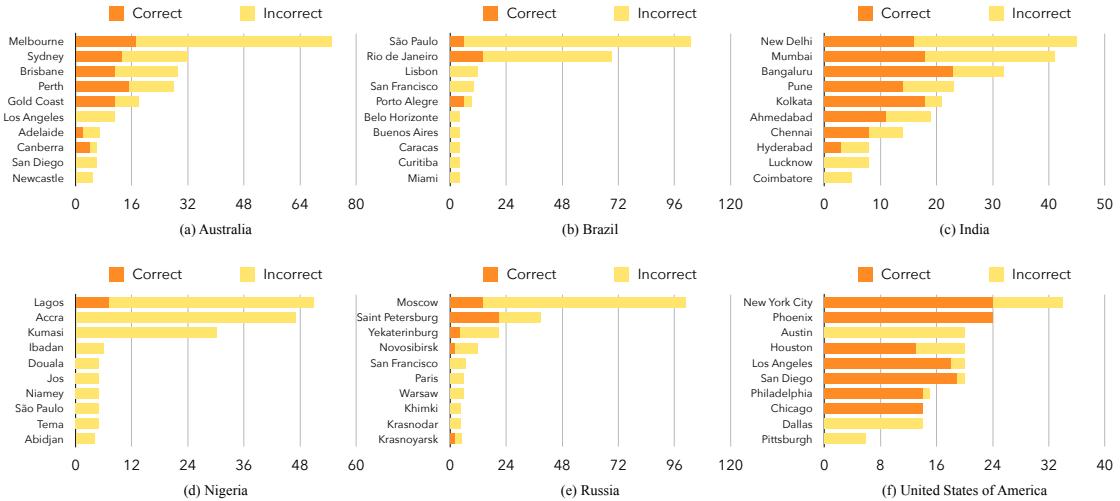


Figure 4: The most frequently predicted cities by LLaMA-3.2-11B-Vision across six countries.

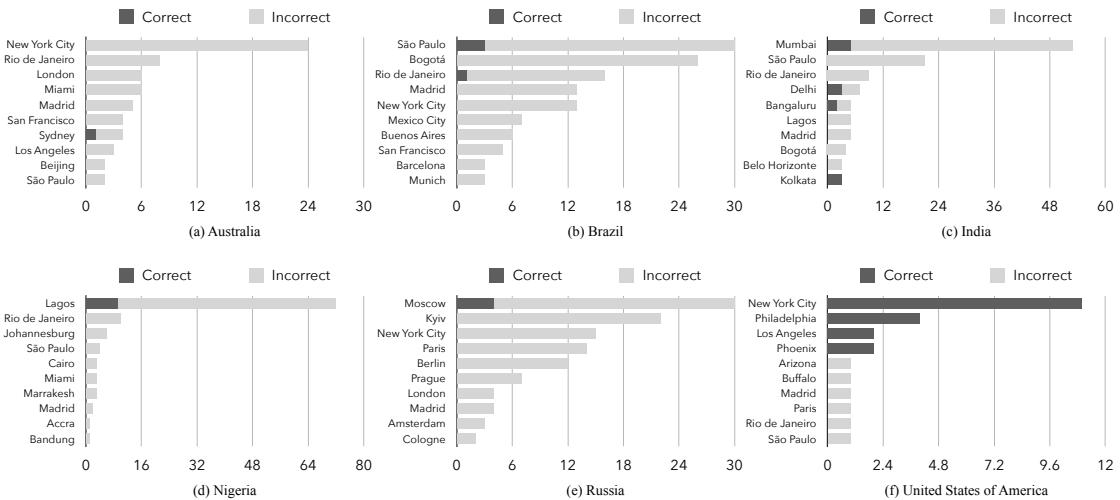


Figure 5: The most frequently predicted cities by LLaVA-V1.6-Vicuna-13B across six countries.

### A.3 Continent-Level Confusion Matrix from Other VLMs

Continent	Africa	Asia	Europe	NA	SA	Oceania
Africa	<b>98.00</b>	1.00	0.00	0.00	1.00	0.00
Asia	1.56	<b>66.44</b>	32.00	0.00	0.00	0.00
Europe	0.00	0.67	<b>97.34</b>	2.00	0.00	0.00
NA	0.00	0.00	0.00	<b>100.00</b>	0.00	0.00
SA	0.00	0.00	0.00	1.33	<b>98.67</b>	0.00
Oceania	0.00	1.00	3.67	4.33	0.00	<b>91.00</b>

Table 11: Confusion matrix of the continent-level results from Gemini.

Continent	Africa	Asia	Europe	NA	SA	Oceania
Africa	<b>83.67</b>	11.33	0.00	1.67	3.33	0.00
Asia	3.11	<b>76.22</b>	15.56	4.00	0.44	0.22
Europe	0.00	35.33	<b>57.34</b>	4.67	0.67	2.00
NA	0.33	0.00	1.33	<b>98.01</b>	0.33	0.00
SA	2.67	1.00	6.33	14.33	<b>73.00</b>	1.33
Oceania	0.67	5.00	2.33	14.67	0.00	<b>75.67</b>

Table 12: Confusion matrix of the continent-level results from LLaMA-3.2-11B-Vision.

Continent	Africa	Asia	Europe	NA	SA	Oceania
Africa	<b>44.67</b>	3.67	1.00	3.33	8.67	0.00
Asia	1.78	<b>26.22</b>	20.22	5.78	12.89	0.00
Europe	0.00	1.33	<b>64.67</b>	8.00	0.67	0.00
NA	0.67	0.00	1.00	<b>50.00</b>	1.00	0.00
SA	0.67	0.67	12.00	13.33	<b>38.67</b>	0.00
Oceania	0.00	2.33	6.00	<b>33.67</b>	4.00	3.00

Table 13: Confusion matrix of the continent-level results from LLaVA-V1.6-Vicuna-13B.

#### A.4 Distance-Based Scores

Country	Total Score	$d \leq 100\text{km}$	$100\text{km} < d \leq 1000\text{km}$	$d > 100\text{km}$
Australia	415	174	67	59
Brazil	407	159	89	52
India	379	147	85	68
Nigeria	338	78	182	40
Russia	236	71	94	135
United States	425	185	55	60

Table 14: GPT-4o scores. We define a distance-based scoring scheme as follows: Score 2: Prediction within 100 km of ground truth Score 1: Prediction within 100–1000 km Score 0: Prediction beyond 1000 km.

Country	Total Score	$d \leq 100\text{km}$	$100\text{km} < d \leq 1000\text{km}$	$d > 100\text{km}$
Australia	455	190	75	35
Brazil	300	85	130	85
India	396	156	84	60
Nigeria	325	53	219	28
Russia	158	46	66	188
United States	425	184	57	59

Table 15: Gemini-1.5-Pro scores. We define a distance-based scoring scheme as follows: Score 2: Prediction within 100 km of ground truth Score 1: Prediction within 100–1000 km Score 0: Prediction beyond 1000 km.

Country	Total Score	$d \leq 100\text{km}$	$100\text{km} < d \leq 1000\text{km}$	$d > 100\text{km}$
Australia	297	100	97	103
Brazil	185	41	103	156
India	336	116	104	80
Nigeria	225	15	195	90
Russia	176	48	80	172
United States	359	140	79	81

Table 16: LLaMA-3.2-11B-Vision scores. We define a distance-based scoring scheme as follows: Score 2: Prediction within 100 km of ground truth Score 1: Prediction within 100–1000 km Score 0: Prediction beyond 1000 km.

Country	Total Score	$d \leq 100\text{km}$	$100\text{km} < d \leq 1000\text{km}$	$d > 100\text{km}$
Australia	6	2	2	296
Brazil	42	9	24	267
India	63	16	31	253
Nigeria	92	9	74	217
Russia	26	5	16	279
United States	41	19	3	278

Table 17: LLaVA-V1.6-Vicuna-13B scores. We define a distance-based scoring scheme as follows: Score 2: Prediction within 100 km of ground truth Score 1: Prediction within 100–1000 km Score 0: Prediction beyond 1000 km.

## B Discussions

### B.1 Is Ten Pictures Per City Enough?

To assess whether ten images per city are sufficient to support our conclusions, we conduct a new set of experiments using the Gemini-1.5-Pro. Each city is represented by 20 images, with each image queried once to predict its geographical location. To evaluate the impact of sample size reduction, we randomly select 10 images from the original 20 and compare the model’s performance to that obtained using the full set. With 20 images per city, the model achieves a city-level accuracy of 63.0%. Using 10 images yields an accuracy of 64.8%, a marginal increase of 1.8 percentage points. A per-city analysis shows that in 91.7% of cities, the accuracy difference between the two settings is within 10%. Given that each image contributes 5% to the city-level accuracy in the 20-image setting, this variation is minimal. We also examine the stability of the model’s performance. When using 20 images per city, the mean standard deviation of city-level accuracy across cities is 0.406; with 10 images, it is 0.370—a relative difference of just 8.9%. This small change suggests that reducing the sample size has a negligible effect on performance variability. Overall, the results indicate that using 10 images per city yields comparable accuracy and stability to using 20, supporting the sufficiency of smaller sample sizes for robust city-level evaluation.

### B.2 Rural vs. Urban

To assess performance differences between urban and rural environments, we conduct a supplementary experiment involving five rural U.S. locations: Woodstock, Vermont; Smicksburg, Pennsylvania; Galena, Illinois; Barboursville, Virginia; and Blue Ridge, Georgia. For each location, we select 10 images and query the Gemini-1.5-Pro once per image to evaluate geolocation accuracy. The model achieves 100% accuracy at the continent and country levels but only 3% at the city level across these rural areas. For comparison, we evaluate the model on 10 U.S. cities, again using 10 images per city and one query per image. In urban settings, the model maintains 100% accuracy at the continent and country levels and achieves 57.7% accuracy at the city level. These results reveal a substantial drop in city-level accuracy for rural areas, indicating that the model performs more reliably in urban regions and struggles with sparsely populated, less visually distinctive environments. This observation

Country	Accuracy
Australia	14.7
Brazil	14.0
India	12.0
Nigeria	14.0
Russia	18.0
United States	34.0
Average	17.8
Random Baseline	0.9

Table 18: Accuracy of CLIP (ViT-B/32).

reinforces our overall conclusion that geolocation accuracy improves with population density and urban visual features.

### B.3 Zero-Shot CLIP

We have conduct an experiment using zero-shot CLIP (ViT-B/32) (Radford et al., 2021). Since CLIP does not support instruction-following or structured prompting like VLMs, we adopt a retrieval-style setup. Specifically, we pair each test image in our Depth dataset with the names of all 111 cities in our two (Depth and Breadth) datasets and selected the city name with the highest similarity score based on CLIP’s visual-textual embedding alignment. The results are as shown in Table 18. Despite its architectural simplicity and lack of geographic priors or structured reasoning, zero-shot CLIP achieves a substantial improvement over the random baseline. This supports the claim that vision-language alignment alone contributes meaningfully to geolocation performance. However, CLIP still lags far behind modern VLMs (e.g., GPT-4o reaches 40–57% accuracy on the same set), which demonstrates the necessity of more advanced multimodal reasoning and world knowledge for city-level geolocation.

### B.4 Error Analysis with Distance-Based Scores

We further illustrate the error patterns with selected cities that have similar numbers of correct predictions (*i.e.*,  $d \leq 100\text{km}$ ), but exhibit very different types of errors, in Table 19. These results highlight that even when models achieve similar levels of correctness at fine-grained levels (*e.g.*, city-level hits), the types of errors vary: some are localized within-region mistakes (*e.g.*, Campinas mispredicted as São Paulo), while others are more severe intercontinental mismatches (*e.g.*, Melbourne predicted as

City	Total Score	$d \leq 100km$	$100km < d \leq 1000km$	$d > 100km$
San Antonio	50	20	10	0
San Diego	48	19	10	1
Campinas	48	19	10	1
Melbourne	41	20	1	9
Kolkata	40	20	0	10
Adelaide	35	16	3	11
Chennai	33	16	1	13
Yekaterinburg	31	14	3	13

Table 19: GPT-4o scores for some selected cities.

a U.S. city). This analysis complements the accuracy and entropy metrics by offering a nuanced view of model behavior and supports the need for geospatially aware evaluation metrics.

## C Case Studies

### C.1 Can CoT Help?

To evaluate the performance of VLMs, we analyze their outputs using Chain-of-Thought (CoT) (Kojima et al., 2022; Wei et al., 2022) prompts. We present two example queries: one for Gemini and another for LLaMA. The case study suggests that while CoT reasoning can appear logical, it is not consistently tied to the final answer. In CoT Example (1), Gemini correctly identifies Africa’s surroundings but notes the absence of visible license plates or signs that could aid in further country or city analysis. Despite this lack of evidence, the model still predicts the correct answer. Conversely, in CoT Example (2), LLaMA identifies features typical of California but incorrectly predicts Santa Barbara instead of the correct answer, Los Angeles. Across multiple examples, the elements cited in the CoT reasoning process often partially align with the final answer. However, these elements are typically broad and fail to accurately pinpoint specific locations. Relying solely on the reasoning process makes it challenging to determine the exact geographical location of an image. We additionally apply direct prompting to Gemini, instructing it to identify the geographical location without invoking explicit reasoning. Results on the breadth subset indicate that CoT prompting yields minimal performance gains, with city-level accuracy of 63.0% using CoT and 61.0% without it. This suggests that the model’s outputs may not stem from genuine visual reasoning but rather reflect prior knowledge of geographic patterns.

### C.2 ChatGPT-o3

We conduct a small-scale experiment on o3 using a random sample of 10 image pairs that GPT-4o misclassifies at the city level. The results show that o3 achieves 0% accuracy on these images. Notably, the misclassified images typically depict less populous or underdeveloped regions. This suggests that o3 may exhibit a similar bias, leading to reduced accuracy for images from such areas.

**CoT Example (1)****IMAGE**

**GROUND TRUTH** Kano, Nigeria, Africa  
**MODEL** Gemini-1.5-Pro

**ANALYSIS** The image appears to be taken in a tropical or subtropical region due to the vegetation and sun angle. The architecture suggests a possibly African context. The road appears to be well-maintained. No license plates or signage are clearly visible to provide definitive clues.

**PREDICTION** Kano, Nigeria, Africa

**CoT Example (2)****IMAGE**

**GROUND TRUTH** Los Angeles, USA, North America  
**MODEL** LLaMA-3.2-11B-Vision

**ANALYSIS** The image shows a woman standing on a dirt path in a wooded area, holding a cell phone and wearing a pink t-shirt, shorts, and a white visor. She is looking out at a view of mountains and valleys in the distance.

- \* The woman is wearing a pink t-shirt, shorts, and a white visor.
- \* She is holding a cell phone in her right hand.
- \* She is standing on a dirt path in a wooded area.
- \* There are trees and bushes surrounding her.
- \* In the background, there is a view of mountains and valleys.
- \* The sky is blue with some clouds.

Based on the analysis, the street view appears to be located in a rural or natural area, possibly in a national park or nature reserve. The presence of mountains and valleys in the background suggests that the location is in a hilly or mountainous region.

**PREDICTION** Santa Barbara, USA, North America

**o3 Example (1)****IMAGE**

**GROUND TRUTH** Campinas, Brazil, South America  
**MODEL** ChatGPT-03

**ANALYSIS** 1) Latitude/Longitude – the sun sits high in the northern half of the sky (shadows fall slightly south-east), placing us in the southern hemisphere at roughly 22–23° S and 49° W. 2) Sun position – almost overhead, hinting the image was taken close to the Tropic of Capricorn in late morning/early afternoon. 3) Vegetation – broad-leaf street trees and palms typical of Brazil's south-east interior; no conifers or temperate species. 4) Natural scenery – fully urban terrain on a gentle slope, no coastline or mountains in view. 5) Buildings – corner lot tyre centre painted grey/orange with the Continental logo; next door narrow two-storey shop houses; far background a 25–30-storey tower common in mid-sized Brazilian cities. 6) Licence plates – long, very narrow rectangles with a small blue Mercosur band (blurred, but shape matches current Brazilian plates). 7) Road directions – right-hand traffic, single broken centre line, Brazilian style zebra crossing and 'E' (Estacionar) no-parking sign. 8) Flags – none visible. 9) Language – Portuguese words on the pole ('Proibido Estacionar') and on the kiosk; brand signage 'Pneus' also Portuguese. 10) Shops – Continental / Cantu Pneus service centre; a small yellow 'orelhão' pay-phone booth (Brazilian hallmark) beside a snack kiosk. 11) Pedestrians – a few customers near the garage entrance, attire consistent with warm Brazilian climate. Combining Portuguese language, Brazilian street furniture, right-hand driving, and specific chain décor for Cantu-Continental stores, the scene fits the interior of São Paulo state. The only Cantu-Continental branch on a left-hand corner with a narrow residential side-street and a tall tower on the same axis is on Avenida Duque de Caxias in Bauru. That corner (with Rua Araujo Leite) matches the slope, tree and phone-booth layout seen here.

**PREDICTION** Bauru, Brazil, South America

**o3 Example (2)****IMAGE**

**GROUND TRUTH** Bengaluru, India, Asia  
**MODEL** ChatGPT-03

**ANALYSIS** 1) Latitude and Longitude: The signage explicitly naming Surat and the semi-arid coastal setting place the view at roughly 21.17° N, 72.86° E on the south-eastern side of the city. 2) Sun position: Short shadows that fall slightly leftward from vertical objects indicate the sun is high in the southern sky; the camera is therefore pointing roughly north-east around midday. 3) Vegetation: A single broad-leaved tree and low scrub, consistent with the dry-tropical vegetation found around coastal Gujarat. 4) Natural scenery: The terrain is flat with no hills or water visible, typical of the Surat plain near the Tapti River delta. 5) Buildings: One- and two-storey concrete workshops and godowns; one structure is under construction with exposed re-bar and wooden scaffolding—common in Indian industrial estates. 6) License plates: Vehicles show yellow (commercial) and white (private) rectangular plates in the Indian style, confirming India. 7) Road directions: Traffic keeps left and vehicles are right-hand-drive, matching Indian road rules. 8) Flags: A plain green shop flag—no national flag present. 9) Language: English text ("SURAT GN-4 NO PARKING", "Z Motors Works") plus faint Gujarati script confirm the local language mix of Gujarat. 10) Shops: A two-wheeler repair garage (Z Motors Works) and a yellow-shuttered VRL Logistics cargo depot identify the spot as part of a transport yard. 11) Pedestrians: Several helmeted scooter riders and workers in typical Indian attire reinforce the urban-industrial context. Taking all of these clues together, the scene fits an interior lane (Road No 3) of the Udhna Udyog Nagar industrial estate on the southern fringe of Surat, Gujarat, India.

**PREDICTION** Surat, India, Asia

## D User Study Questionnaire

In the following 10 questions, you are asked to guess the geographical location revealed by the following photoes based on their content.

Please note that you may not resort to any search engines or AI models to answer this question.

Your answer should include: continent, country and city, a total of THREE pieces of information.

### (a) Instruction for human participants.

Based solely on this picture, guess the following information



Guess the continent

- Africa
- Asia
- Europe
- North America
- South America
- Oceania

Guess the country

Guess the city

### (b) An example question.

Figure 6: Illustration of our questionnaires.