

# ReaGeo: Reasoning-Enhanced End-to-End Geocoding with LLMs

Jian Cui<sup>1,2†</sup> Zhiyuan Ren<sup>1†</sup> Desheng Weng<sup>1</sup> Yongqi Zhao<sup>1</sup>  
GONG WENBIN<sup>1</sup> Yu Lei<sup>1</sup> Zhenning Dong<sup>1,2</sup>

<sup>1</sup>Amap, Alibaba Group <sup>2</sup>Tsinghua University

<sup>1</sup>{cj87619, renzhiyuan.rzy, wengdesheng.wds, yongqi.zyq}@alibaba-inc.com

<sup>1</sup>{gongwenbin.gwb, yu.lei, zhenning.dong}@alibaba-inc.com

## Abstract

This paper proposes ReaGeo, an end-to-end geocoding framework based on large language models, designed to overcome the limitations of traditional multi-stage approaches that rely on text or vector similarity retrieval over geographic databases, including workflow complexity, error propagation, and heavy dependence on structured geographic knowledge bases. The method converts geographic coordinates into geohash sequences, reformulating the coordinate prediction task as a text generation problem, and introduces a Chain-of-Thought mechanism to enhance the model’s reasoning over spatial relationships. Furthermore, reinforcement learning with a distance-deviation-based reward is applied to optimize the generation accuracy. Comprehensive experiments show that ReaGeo can accurately handle explicit address queries in single-point predictions and effectively resolve vague relative location queries. In addition, the model demonstrates strong predictive capability for non-point geometric regions, highlighting its versatility and generalization ability in geocoding tasks.

## 1 Introduction

Geocoding converts human-readable geographic location descriptions into machine-readable coordinates that enable visualization on digital maps, navigation, or spatial positioning (Rose et al., 2004; Rushton et al., 2007; Goldberg et al., 2007; Yin et al., 2025). Geocoding demonstrates significant applications across daily life, industrial applications, and academic research. For instance, in traffic accident reporting scenarios, it is necessary to find the accurate geolocations described by the person who made the alarm call (Li et al., 2020; Cheng et al., 2024). Similarly, precise destination identification in logistics order processing benefits from such capabilities. Furthermore, people

can identify specific geolocations based on news events or social media messages (Zhang and Gellerter, 2014; Zou et al., 2020). Industrially, leading platforms such as Google Maps, Microsoft Bing Maps, and Amap (The leading digital map service in China) also use geocoding as a public capability and provide API services (Mokhtari et al., 2019). These services support diverse sectors including e-commerce, urban planning, disaster management and so forth, with global API call volumes exceeding millions daily. Consequently, optimizing localization accuracy, such as achieving sub-street-level precision, is essential for geocoding systems.

Traditional geocoding approaches are primarily retrieval-based methods built on structured geographic databases. A common strategy is to decompose an input query into geographic entities using methods such as Named Entity Recognition (NER), hierarchical address tagging, or structural parsing supported by learning-based models (Hu, 2018; Clemens, 2020; Radford, 2021; Yin et al., 2023; Li et al., 2023). After entity extraction, the parsed components are matched against large-scale geographic databases through rule-based or fuzzy text retrieval mechanisms, often combined with ranking models, spatial constraints, or similarity metrics to link queries with the most plausible geographic entities (Melo and Martins, 2017; Yin et al., 2019b,a; Hu et al., 2023; Rezaei Ghahroodi et al., 2024; Yin et al., 2025). In addition to text-matching pipelines, an alternative line of work directly embeds address strings into vector representations and performs cosine-similarity retrieval to identify the most likely geographic candidates, which helps improve robustness against spelling variations and informal expressions. Despite their widespread adoption, these pipeline-based methods face several limitations. Incomplete or ambiguous textual descriptions can disrupt entity extraction and degrade retrieval accuracy. Moreover, semantic inconsistencies between user queries and database

<sup>†</sup>Equal contribution.




Figure 1: Geohash in digital maps. Geohash encodes geographic coordinates into a short alphanumeric string via recursive range bisection and Base32 encoding.

records—arising from abbreviations, dialectal variations, informal expressions, or inconsistent naming conventions, frequently lead to retrieval failures or incorrect matches.

To address the challenges of traditional pipeline-based geocoding, an innovative end-to-end method based on large language models (LLMs) is adopted, which eliminates the complex and error-prone multi-stage process relying on geographic knowledge bases by directly learning the joint distribution of query texts and geographic spaces. The query text is input to a generative model that autoregressively decodes a geohash sequence, with the center point of the area represented by the geohash approximated as the final coordinates. Geohash, a spatial indexing method, models the Earth’s surface as a two-dimensional plane and recursively partitions it into smaller sub-blocks, with each sub-block assigned the same code within a certain range of latitudes and longitudes (Clemens, 2020), as shown in Figure 1. By transforming coordinate prediction into a generative task, alignment between text and geographic modalities is achieved, allowing the powerful language understanding and knowledge memory capabilities of LLMs to be fully leveraged for geocoding applications.

In industrial settings, digital map services such as Google Maps and Amap generate massive volumes of user query and click-through logs daily, which are used to construct an end-to-end training dataset that forms the foundation for large-scale model training. And in this study, Chain-of-Thought (CoT) (Wei et al., 2022) and reinforcement learning are employed to enhance the model’s spatial reasoning and overall prediction performance. Comprehensive experiments demonstrate the superiority of this end-to-end approach, achieving accurate predictions for both single-point locations and non-point geometries.

The main contributions of this paper are summa-

rized as follows:

- ReaGeo, the first end-to-end geocoding model based on LLMs is proposed, offering a unified framework that directly maps textual queries to geohash sequences and thereby eliminates the need for complex and error-prone multi-stage pipelines in geocoding systems.
- A reinforcement learning reward function is developed for geocoding, based on the distance deviation between predicted results and ground truth, which strengthens the model’s spatial reasoning capability and improves its ability to resolve vague or implicitly expressed relative location descriptions.
- The proposed ReaGeo achieves accurate predictions for both single-point locations, including explicit address descriptions and vague relative positional instructions, and non-point geometric shapes, such as line-shaped roads and polygonal administrative areas, demonstrating strong versatility and generalization in complex geospatial tasks.

## 2 Related Work

### 2.1 Geocoding

In the field of geocoding, previous research mainly focuses on traditional pipeline-based text retrieval approaches (Goldberg et al., 2007; Hu et al., 2023; Yin et al., 2025). These methods typically rely on NER, entity linking, and database queries to perform geocoding (Melo and Martins, 2017; Hu, 2018; Mokhtari et al., 2019; Radford, 2021; Rezaei Ghahroodi et al., 2024). Although these traditional approaches complete geocoding to a certain extent, they face several challenges, such as complex pipelines, dependence on large volumes of manually annotated data, and limited ability to handle ambiguous or incomplete geographic descriptions.

With the rapid development of deep learning and reinforcement learning, the precision and efficiency of geocoding have been significantly improved in recent years. These emerging techniques bring new opportunities to the geocoding field, enabling models to better understand and process complex geographic information. For instance, deep learning automatically extracts salient features from textual and spatial data, while reinforcement learning enhances model performance on complex tasks by

optimizing decision-making behaviors (Yin et al., 2019a, 2023; Li et al., 2023; Rezaei Ghahroodi et al., 2024).

In related research domains, reverse geocoding, which converts coordinates into textual descriptions, also receives increasing attention. Moreover, multi-modal localization studies, such as AddressCLIP (Xu et al., 2024), MixVPR (Ali-Bey et al., 2023), G3 (Jia et al., 2024), and other works (Berton et al., 2022; Durgam et al., 2024; Wilson et al., 2024), explore the integration of visual information for geo-localization. These studies highlight promising directions for future research in leveraging textual, spatial, and visual information to achieve more accurate and robust geocoding.

## 2.2 Large Language Models and Reinforcement Learning

In recent years, LLMs have made significant strides in natural language processing. The introduction of Transformer-based (Vaswani et al., 2017) models, which leverage the self-attention mechanism, dramatically improves machine translation performance and reduces training time. Subsequent architectures such as BERT (Devlin et al., 2019) and GPT (Radford et al., 2018) reshape model pre-training.

GPT-3 (Brown et al., 2020), a model with 175 billion parameters, achieves notable few-shot learning performance. Techniques like Instruction Fine-Tuning (IFT), Supervised Fine-Tuning (SFT), and Reinforcement Learning from Human Feedback (RLHF) are widely adopted to improve instruction following and alignment with human preferences (Ouyang et al., 2022; Bai et al., 2022; Achiam et al., 2023). Reinforcement learning methods treat an LLM as a policy and optimize task-specific rewards via policy-gradient approaches. Proximal Policy Optimization (PPO) (Schulman et al., 2017) becomes a de-facto choice by constraining updates within a clipped trust region to keep the language distribution close to a reference model while maximizing reward. Extensions such as Group Relative Policy Optimization (GRPO) (Shao et al., 2024) compute advantages relative to intra-group baselines to improve sample efficiency and fine-tuning stability for billion-parameter agents. This line of work spawns powerful open-source models (e.g., LLaMA (Touvron et al., 2023), Qwen (Team et al., 2024; Yang et al., 2024, 2025), DeepSeek (Bi et al., 2024)) that serve as foundations for downstream tasks and for reinforcement learning-based

fine-tuning. Reinforcement learning thus plays an increasingly central role in enhancing LLMs’ instruction-following and complex reasoning abilities.

## 3 Methods

### 3.1 Preliminaries

Unlike traditional vector-retrieval-based geocoding methods illustrated in Figure 2(a), the proposed end-to-end approach, ReaGeo, learns the joint distribution between query texts and geographic spaces, thereby inferring the locations that the queries indicate. Formally, this joint distribution constitutes a multi-modal alignment problem, but we convert the geographic space into text through geohash transformation, thus simplifying the multi-modal learning cost. A large language model is adopted as the backbone and is post-trained to complete the alignment of query texts and geohash strings. During inference, the query is fed into ReaGeo, from which the most probable geohash is decoded, and the coordinates of the centroid of the represented region are taken as the final predicted coordinates. This study adopts a geohash Base32 encoding of length 9, which offers a precision of approximately 2.4 meters, meeting the accuracy requirements of typical search scenarios. The framework of ReaGeo is illustrated in Figure 2(b), and its geocoding process can be formulated as follows:

$$Y_{\text{geohash}} = F_{\text{ReaGeo}}(X_{\text{query}}), \quad (1)$$

$$y_{\text{pred}} \triangleq (\phi, \lambda) = G_{\text{decode}}(Y_{\text{geohash}}) \quad (2)$$

where  $X_{\text{query}}$  denotes the input address query,  $F_{\text{ReaGeo}}$  represents the end-to-end mapping performed using GeoLLM, and  $Y_{\text{geohash}}$  denotes the output geohash. Subsequently, the geohash decoding  $G_{\text{decode}}$  is applied to  $Y_{\text{geohash}}$  to obtain the final predicted coordinates  $y_{\text{pred}}$ , in which  $\phi$  is the latitude and  $\lambda$  is the longitude, respectively.

The proposed ReaGeo consists of two training stages. In the first stage, the CoT method is introduced to enhance the model’s reasoning capability, thereby capturing the associations among neighboring addresses and improving prediction performance on geographically ambiguous queries. The second stage builds upon the first by employing reinforcement learning with deviation distance as an effective reward signal to further refine positional prediction accuracy.

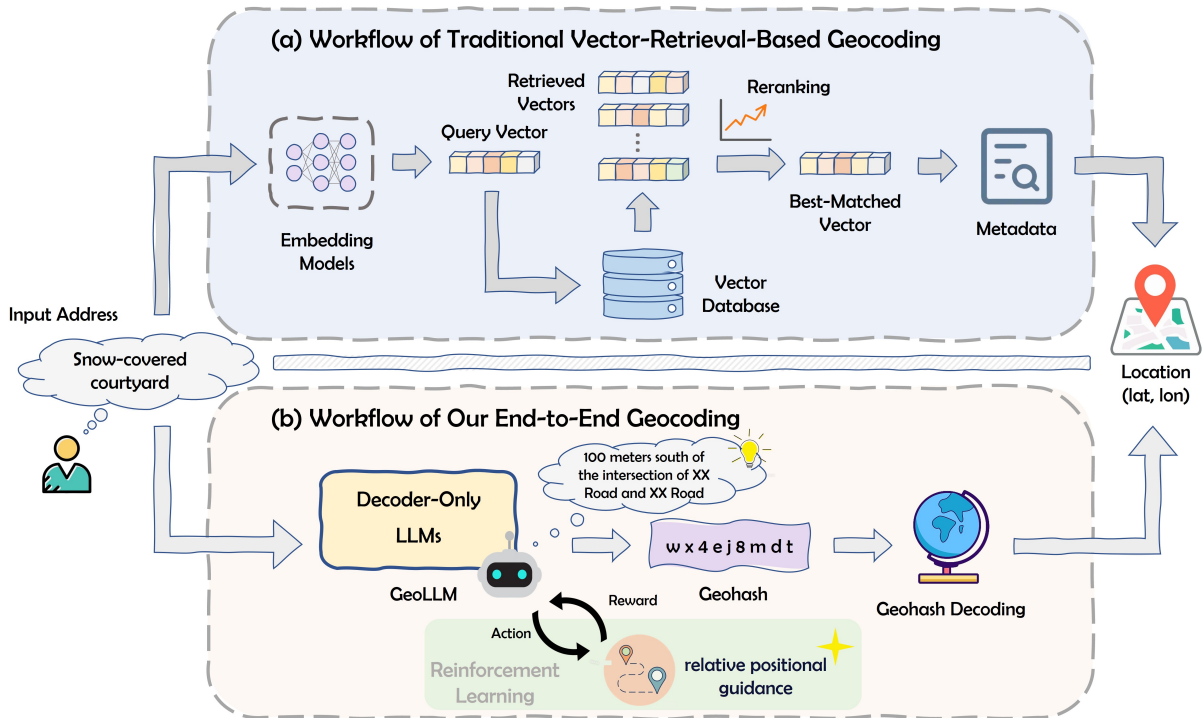


Figure 2: The workflow of traditional geocoding methods and our end-to-end geocoding method. (a) Workflow of traditional vector-retrieval-based geocoding: The input address is first transformed into a query vector through an embedding model. This vector is then used to retrieve a set of candidate vectors from a large vector database, followed by a reranking stage to identify the best-matched vector. Finally, associated metadata is retrieved, from which the corresponding geographic coordinates are extracted. (b) Workflow of our end-to-end geocoding: The input query is fed directly into ReaGeo, a decoder-only LLM which is supervised fine-tuned with the CoT method, enabling step-by-step spatial reasoning. To further enhance the model’s capability in resolving relative location descriptions, reinforcement learning with relative positional guidance is employed. ReaGeo generates the geohash sequence in an autoregressive manner, which is subsequently decoded into geographic coordinates.

### 3.2 Decoder-Only Architecture

The core objective of our geocoding task is to directly map natural language queries (such as addresses or location descriptions) into geohash in a generative manner. The decoder-only architecture, by leveraging autoregressive generation, enables end-to-end production of coordinate sequences from textual inputs without relying on complex multi-stage pipeline processes, thereby reducing intermediate error propagation and minimizing dependence on extensive, well-curated geographic knowledge bases.

In this study, Qwen2.5-3B is adopted, which achieves a well-balanced trade-off among language capability, inference accuracy, speed, and resource consumption. This model can effectively handle high-precision parsing of both structured and unstructured address descriptions, while supporting low-latency, high-concurrency deployment in both cloud-based and on-premises production environments.

### 3.3 CoT

To strengthen the association between toponymic references and their formal address representations, thereby improving prediction performance for incomplete or colloquial location descriptions, the CoT method is introduced. Specifically, for a POI address, the approach uses the address itself as the endpoint and its adjacent addresses as the starting points to calculate relative positional relationships. These relative position associations serve as intermediate reasoning steps, extending the model’s output to guide it in learning neighborhood knowledge to compensate for the lack of input information before generating the corresponding geohash.

To alleviate memory pressure during inference while maintaining compatibility with the original geohash output format, the CoT framework is combined with Qwen’s thinking-mode tagging mechanism. Specifically, special tags `<thinking>` and `</thinking>` are introduced in the training samples to explicitly demarcate intermediate reasoning

text, and the thinking mode can be disabled during inference. An example of the data format is as follows:

```
"input": "Snow-covered courtyard
<thinking>"
"output": "100 meters south of the
intersection of XX Road and XX Road, XX
District </thinking> w x 4 e j 8 m d t"
```

### 3.4 Reinforcement Learning

As an efficient and stable policy optimization algorithm, GRPO is widely adopted in various generative tasks. Unlike value-function-based methods such as PPO, GRPO does not require training a separate value model. Instead, it calculates the group relative advantage by normalizing the rewards of a set of candidate outputs generated for the same input query, and then updates the policy based on this relative comparison, thereby significantly reducing training complexity and instability. Based on the above strengths, this study employs GRPO to directly optimize geocoding precision. This approach enables the model to learn fine-grained spatial relationships, especially for understanding precise directional offsets within a query. In this phase, the training data incorporates random offset descriptions, while maintaining the direct query-geohash format. An example of the data format is as follows:

```
"input": "50 meters north of No.385 XX
Road, XX Town, XX District, XX"
"output": "w x 4 s f d v p y"
```

**Reward function design.** For the geocoding task, a task-specific reward function based on geospatial distance deviation is innovatively designed as follows:

$$R(y_{\text{pred}}, y_{\text{true}}) = \frac{\sqrt{T} - \sqrt{D(y_{\text{pred}}, y_{\text{true}})}}{S} \quad (3)$$

where  $y_{\text{pred}}$  and  $y_{\text{true}}$  are the predicted and ground-truth coordinates,  $T$  is a hyperparameter set to 100,  $D(\cdot)$  computes the WGS-84 ellipsoid-based geodesic distance between two coordinates in meters, and  $S$  is a normalization factor set to 1000.

This reward increases as the predicted coordinates approach the true location, encouraging precise spatial inference during training.

## 4 Experiments

In this section, the dataset and evaluation metrics are first introduced, followed by an evaluation of

single-point prediction performance on both clear and ambiguous queries, ablation studies to assess the contribution of key components, and finally case studies visualizing predictions for point and non-point geographic regions.

### 4.1 Datasets

A total of 239,918 data samples from both rural and urban areas of Beijing is obtained via Amap search engine logs and the Amap points of interest (POI) database, and partitioned them into training and test sets. All samples are unique, and for each location in the test set, the corresponding neighboring area is covered within the training set. From the perspective of data distribution and scenario characteristics, there are significant differences between rural and urban areas. Rural areas generally have looser address structures, higher address ambiguity, and sparser infrastructure annotations. In contrast, urban areas have more regular address systems, a high density of POI, and a high degree of road grid formation. These differences directly affect the modeling difficulty and performance of various methods.

For each data sample, two types of input address descriptions are constructed, which are further categorized into two classes of data, as follows:

**Base Data.** Retrieval queries are constructed from three textual sources: search engine queries, POI names, and POI addresses. For search engine queries, the coordinates of the user’s final clicked POI are used as the prediction target. For POI names and addresses, the corresponding POI coordinates serve as the prediction target. The training set of the base data is primarily used in ReaGeo’s SFT phase.

**Anchor Offset Data.** Relative location descriptions are commonly used in real-world scenarios such as traffic accident reports and daily communication, making accurate interpretation and geocoding of them highly valuable. To this end, based on the base data, the anchor offset data is constructed, in which phrases such as “[distance] + meters + [direction] + of” are added to the beginning of each description. The direction is randomly selected from the four cardinal directions (east, south, west, and north), and the distance is uniformly sampled from the range of 30 to 500 meters. For example, “200 meters south of the intersection of XX Road and XX Road”. Meanwhile, the corresponding coordinates are shifted by the distance indicated in its description to generate new labels. The training

Table 1: Details of Experimental Data Partitioning

Categories	Data	Training	Test
Rural	Base	109050	11995
	Anchor Offset	109050	11995
Urban	Base	106878	11995
	Anchor Offset	106878	11995
Total	Base	215928	23990
	Anchor Offset	215928	23990

set of the anchor offset data is primarily used in ReaGeo’s reinforcement learning phase.

The detailed dataset split is presented in Table 1.

## 4.2 Implementation Details

In the SFT phase incorporating CoT, a full-parameter fine-tuning of Qwen2.5-3B is conducted on 2 NVIDIA H20 GPUs, in a distributed data-parallel setup. The batch size is set to 128 per GPU and the number of epochs to 5. The AdamW optimizer is employed with an initial learning rate of  $5 \times 10^{-5}$ .

Reinforcement learning builds upon the SFT-trained checkpoint enhanced with CoT. During GRPO phase, the actor model is updated with a batch size of 512 prompts per optimization step, subdivided into micro-batches of 128 to fit memory constraints. Rollout generation is performed with a batch size of 512 prompts, split into micro-batches of 128, and each prompt samples 8 candidate responses to estimate advantages within a group. The number of training epochs in this phase is set to 3, and the actor learning rate is initialized at  $1 \times 10^{-6}$ .

## 4.3 Evaluation Protocols

Two metrics are used to evaluate the model’s performance, as detailed below:

The first is the average deviation distance (ADD), which measures spatial accuracy by calculating the WGS-84 ellipsoid-based geodesic distance between the predicted coordinates and the ground-truth coordinates, and then averaging the results across all samples. A lower ADD value indicates higher geocoding precision.

The second is a distance-threshold-based accuracy metric (accuracy under distance thresholds,  $\text{Acc}@k$ ), which evaluates the model’s localization capability, where  $k$  denotes the maximum allowable spatial deviation. Four thresholds are selected, namely  $k = 100, 200, \text{ and } 500$  meters, represented

as  $\text{Acc}@100$ ,  $\text{Acc}@200$ , and  $\text{Acc}@500$ , respectively. These metrics indicate the proportion of samples whose predicted coordinates fall within the corresponding distance from the ground-truth coordinates, thereby assessing the model’s localization accuracy under varying spatial precision requirements.

## 4.4 Quantitative Evaluation

Quantitative experiments are conducted on rural and urban areas separately using both the base data and the anchor offset data, and the core task is to predict single-point coordinates.

Seven comparison methods are selected, comprising traditional and end-to-end pipelines: 1) NER based on Bert-Base (Devlin et al., 2019) combined with Edit Distance, denoted as **NER+Levenshtein**; 2) Recalling the single most relevant address using vector-based retrieval with BGE-Base (Xiao et al., 2024), denoted as **Vector-Retrieval@Top1**; 3) Recalling the top five most relevant addresses using vector-based retrieval with BGE-Base, followed by reranking with BGE-Reranker-Base (Xiao et al., 2024), denoted as **Vector-Retrieval@Top5+Reranker**; 4) **Qwen3-Max** (Yang et al., 2025), the Qwen series model with the best overall performance to date, without any fine-tuning; 5) the geocoding API of Baidu Maps (Baidu Maps, 2025), denoted as **Baidu Maps**; and 6) the geocoding API of Tencent Maps (Tencent Maps, 2025), denoted as **Tencent Maps**.

### 4.4.1 Performance on Base Data

The experimental results are shown in Table 2. In both rural and urban areas, NER+Levenshtein exhibits significant limitations, primarily due to the complex structure of geographic addresses and its inability to capture semantic equivalence. Vector-retrieval-based methods are relatively stable and achieve decent performance, yet they are still constrained by the coverage of the retrieval database and the accuracy of ranking. Qwen3-Max predicts geographic coordinates based on its existing knowledge, but even though the search range is constrained, the results remain highly unsatisfactory. Baidu Maps and Tencent Maps are limited by their matching mechanisms and data coverage, resulting in only mediocre performance when handling ambiguous or non-standard addresses.

ReaGeo outperforms others, achieving the lowest ADD and the highest  $\text{Acc}@k$  on both rural and urban data. Its generalized understanding allows it

Table 2: Quantitative Results of Each Method

Data	Methods	Rural				Urban			
		ADD	Acc @100	Acc @200	Acc @500	ADD	Acc @100	Acc @200	Acc @500
		(m)	(%)	(%)	(%)	(m)	(%)	(%)	(%)
Base Data	NER+Levenshtein	5016.7	36.2	49.2	62.4	2946.4	57.9	69.3	77.1
	Vector-Retrival@Top1	924.1	51.6	67.7	83.4	199.1	75.3	88.7	95.6
	Vector-Retrival@Top5+Reranker	1044.4	51.4	67.2	82.4	162.6	75.6	88.7	95.7
	Qwen3-Max	6548.1	0.0	1.7	8.5	1055.0	0.0	5.1	25.4
	Baidu Maps	2017.5	36.0	50.3	69.4	347.7	66.4	79.8	90.1
	Tencent Maps	1689.2	32.8	50.9	71.9	204.1	69.9	84.9	93.7
	ReaGeo (Ours)	<b>756.6</b>	<b>58.1</b>	<b>73.5</b>	<b>85.8</b>	<b>119.6</b>	<b>81.9</b>	<b>92.4</b>	<b>97.2</b>
Anchor Offset Data	NER+Levenshtein	9422.5	4.3	12.0	38.8	6791.2	7.3	19.4	56.3
	Vector-Retrival@Top1	1083.5	9.8	26.5	76.9	398.4	12.7	32.6	91.2
	Vector-Retrival@Top5+Reranker	1465.7	9.0	23.8	70.2	378.1	12.6	32.2	89.5
	Qwen3-Max	9207.6	0.0	0.0	5.1	1163.1	1.7	5.1	30.5
	Baidu Maps	2140.4	14.6	29.7	64.6	493.5	22.1	42.2	86.2
	Tencent Maps	1792.0	7.3	20.3	64.7	400.0	11.8	31.2	88.5
	ReaGeo (Ours)	<b>840.9</b>	<b>23.7</b>	<b>50.6</b>	<b>84.1</b>	<b>208.7</b>	<b>41.5</b>	<b>71.5</b>	<b>96.3</b>

to use numerous address annotations for accurate alignment and reason about semantic equivalence through pre-trained geographical knowledge. Furthermore, It leverages LLMs’ long-tail information understanding to handle vague addresses and strong context dependence, and its joint modeling of geographical prior and dynamic reasoning enables reasoning about implicit spatial relationships.

#### 4.4.2 Performance on Anchor Offset Data

As presented in Table 2, all the methods exhibit varying degrees of performance degradation on the anchor offset data. The ADD metric degrades particularly severely for NER+Levenshtein and Qwen3-Max. For the former, the randomly added positional bias descriptions in the addresses increase the edit distance, making it more likely to miss the optimal match. For the latter, such random positional bias descriptions introduce greater confusion into the model’s reasoning process. As for the other baseline methods, their performance degradation falls within expected ranges, and their overall performance remains suboptimal.

Across both rural and urban scenarios in the anchor offset data, ReaGeo demonstrates remarkable robustness and precision. It still outperforms other methods due to its understanding of complex geographical semantics and innovative training strategies that target core geocoding challenges like ad-

dress ambiguity, spatial relationship reasoning, and distribution shift.

To further demonstrate the advantages of the proposed method, experiments are conducted on a second dataset, which will be publicly released, with details provided in Appendix A.

#### 4.5 Ablation Studies

The effectiveness of the proposed method is validated from four aspects.

**Analysis of Key Components.** Here, the base data and the anchor offset data are combined. Table 3 presents the results of the ablation study on key components. Removing the reinforcement learning component leads to a significant degradation in performance metrics, highlighting its critical role in enhancing geocoding precision. Eliminating the CoT mechanism also leads to degraded geocoding performance and an increased number of instances in which the model produces invalid outputs (referred to as error count, EC). Additionally, the thinking delimiters within the CoT mechanism offer a minor improvement in the proportion of high-precision samples.

**Impact of Model Size.** Experiments are also conducted on the combined base and anchor offset data here. Figure 3 compares the geocoding performance of Qwen with different model sizes. As the model size increases, localization accuracy

Table 3: Ablation Study of Key Components

Methods	ADD (m)	Acc@100 (%)	Acc@200 (%)	Acc@500 (%)	EC
w/o CoT	475.5	48.7	69.3	90.0	80
w/o thinking delimiters	<b>474.2</b>	50.7	71.7	90.6	21
w/o GRPO	599.2	40.6	57.0	86.1	26
full	481.2	<b>51.3</b>	<b>72.0</b>	<b>90.9</b>	<b>20</b>

Table 4: Comparison of Performance under Different Output Formats in SFT

Methods	ADD (m)	Acc @100 (%)	Acc @200 (%)	Acc @500 (%)
Coordinates	440.1	68.7	81.8	90.6
Geohash	<b>427.7</b>	<b>68.8</b>	<b>82.5</b>	<b>91.5</b>

Table 5: Comparison of Performance between Cardinal and Unseen Intercardinal Directions

Methods	ADD (m)	Acc @100 (%)	Acc @200 (%)	Acc @500 (%)
Cardinal	524.8	32.6	61.1	90.2
Intercardinal	562.9	22.4	50.5	89.4

continuously improves, yet the average inference time also grows. The Qwen2.5-3B strikes a good balance between accuracy and efficiency, making it the preferred choice for practical applications.

**Comparison of Output Format.** On the base data, the Qwen2.5-3B model is evaluated under SFT, comparing direct generation of latitude–longitude coordinates with geohash sequences. Table 4 demonstrates that geohash discretizes continuous geographic space into a sequential format, which better aligns with the text generation mechanism of language models and supports progressive precision control.

**Generalization to Intercardinal Directions.** An anchor-offset test dataset with four intercardinal directions (northeast, southeast, northwest, and southwest) is constructed, and the performance of the proposed method on it is compared with that on the previously used four cardinal directions (east, west, south, and north), as shown in Table 5. It can be observed that the model maintains relatively high accuracy even on unseen intercardinal directions, demonstrating a certain level of generalization ca-

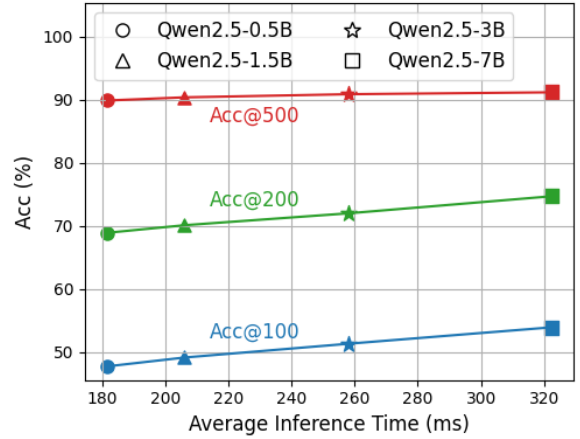


Figure 3: Comparison of Qwen with different model sizes. Both inference accuracy and latency on an NVIDIA H20 GPU increase as model size grows.

pability.

## 4.6 Case Studies

In the case studies, concrete examples are used to demonstrate that ReaGeo, while performing single-location prediction, shows strong inference capability for directional and distance indicators. Additionally, it can estimate the approximate extent of non-point geometries such as lines and polygons.

### 4.6.1 Single-Point Prediction

To evaluate ReaGeo’s ability to handle relative location descriptions, it’s asked to predict the locations of points at specified offsets in the four cardinal directions relative to a given landmark (e.g., Wudaokou Subway Station, denoted as point P). The prediction results, visualized in Figure 4, show that ReaGeo can effectively estimate the relative positional shifts of these points around the landmark within an acceptable error range, demonstrating notable practical value.

### 4.6.2 Non-Point Geometry Prediction

To enable non-point region prediction, beam search with a beam width of 50 is applied to each POI. Multiple candidate coordinate points are generated



Figure 4: Visualization of single-point location prediction. The markers 2–5 correspond to predicted positions guided by relative-position cues around landmark 1.

through this process. These points are then visualized on a digital map to intuitively show the spatial distribution of the predictions. This strategy offers a more comprehensive depiction of predictive uncertainty and visualizes the spatial patterns of the results clearly. The following renders the predicted line and polygon POIs.

**Line POI Prediction.** As illustrated by examples in Figure 5, the heat map visualization shows rendering points distributed along both roads (red lines), revealing distinctive distribution patterns. Although some rendering points deviate from the road centerlines and outlines are incomplete, they still generally follow the roads, providing an initial view of the heat distribution along the roads.

**Polygon POI Prediction.** As illustrated by examples in Figure 6, the heat map rendering shows the distribution of activity within each district. The red outlines mark the administrative boundaries of the respective areas, while the heat map visualizes the actual concentration of points. It can be seen that the heat map coverage accounts for roughly 70% of each district. Upon closer inspection, the heat rendering is mainly concentrated in the core zones, while other parts, although still within the administrative boundaries, are not directly highlighted.

## 5 Conclusions

This study introduces ReaGeo, an end-to-end geocoding framework that utilizes LLMs, CoT and reinforcement learning to directly convert textual location descriptions into geohash. The proposed method simplifies the traditional multi-stage pipeline and demonstrates robust performance in




Figure 5: Visualization of line POI prediction. Coordinate points are rendered along road segments.



Figure 6: Visualization of polygon POI prediction. The heat map shows a polygon area with its administrative boundary, with high-intensity rendering concentrated in the core zone and surrounding regions less highlighted.

handling both explicit addresses and vague relative location queries. It also demonstrates a promising capability to predict non-point geographic regions, including lines and polygons, indicating significant potential for handling complex geospatial data types. Subsequent work will focus on multi-modal geocoding by integrating visual data.

## Limitations

Although the proposed end-to-end geocoding framework achieves a higher upper bound of accuracy, it still has notable limitations. The current experiments are limited to a single city, and achieving localization over a larger geographic area requires a larger language model, which comes at the cost of increased computational resource consumption. In future work, the scalability of the model across hundreds of cities will be investigated, and the trade-off between backbone model size and geographic coverage breadth will be examined. Additionally, in the proposed LLM-based geocoding system, disentangling memorization from reasoning remains an open problem that warrants further investigation.

## Acknowledgements

This research was supported by the National Natural Science Foundation of China (NSFC) under Grants 72571008.

## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Amar Ali-Bey, Brahim Chaib-Draa, and Philippe Giguere. 2023. Mixvpr: Feature mixing for visual place recognition. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 2998–3007.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, and 1 others. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.
- Baidu Maps. 2025. *Geocoding API*. <https://lbsyun.baidu.com/faq/api?title=webapi/guide/webservice-geocoding-base>.
- Gabriele Berton, Carlo Masone, and Barbara Caputo. 2022. Rethinking visual geo-localization for large-scale applications. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4878–4888.
- Xiao Bi, Deli Chen, Guanting Chen, Shanhuang Chen, Damai Dai, Chengqi Deng, Honghui Ding, Kai Dong, Qiu Shi Du, Zhe Fu, and 1 others. 2024. Deepseek llm: Scaling open-source language models with longtermism. *arXiv preprint arXiv:2401.02954*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and 1 others. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Yuhan Cheng, Zhengcong Yin, Diya Li, and Zhuoying Li. 2024. Assessing urban safety: A digital twin approach using streetview and large language models. In *2024 IEEE 100th Vehicular Technology Conference (VTC2024-Fall)*, pages 1–5. IEEE.
- Konstantin Clemens. 2020. *Geocoding user queries*. Ph.D. thesis, Technische Universität Berlin.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. 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)*, pages 4171–4186.
- Abhilash Durgam, Sidike Paheding, Vikas Dhiman, and Vijay Devabhaktuni. 2024. Cross-view geo-localization: a survey. *IEEE Access*.
- Daniel W Goldberg, John P Wilson, and Craig A Knoblock. 2007. From text to geographic coordinates: the current state of geocoding. *URISA journal*, 19(1):33–46.
- Xuke Hu, Zhiyong Zhou, Hao Li, Yingjie Hu, Fuqiang Gu, Jens Kersten, Hongchao Fan, and Friederike Klan. 2023. Location reference recognition from texts: A survey and comparison. *ACM Computing Surveys*, 56(5):1–37.
- Yingjie Hu. 2018. Geo-text data and data-driven geospatial semantics. *Geography Compass*, 12(11):e12404.
- Pengyue Jia, Yiding Liu, Xiaopeng Li, Xiangyu Zhao, Yuhao Wang, Yantong Du, Xiao Han, Xuetao Wei, Shuaiqiang Wang, and Dawei Yin. 2024. G3: an effective and adaptive framework for worldwide geolocalization using large multi-modality models. *Advances in Neural Information Processing Systems*, 37:53198–53221.
- Diya Li, Harshita Chaudhary, and Zhe Zhang. 2020. Modeling spatiotemporal pattern of depressive symptoms caused by covid-19 using social media data mining. *International Journal of Environmental Research and Public Health*, 17(14):4988.
- Min Li, Zeyu Liu, Gang Li, Mingle Zhou, and Delong Han. 2023. Enex-fp: A bert-based address recognition model. *Electronics*, 12(1):209.
- Fernando Melo and Bruno Martins. 2017. Automated geocoding of textual documents: A survey of current approaches. *Transactions in GIS*, 21(1):3–38.
- Shekoofeh Mokhtari, Ahmad Mahmoodi, Dragomir Yankov, and Ning Xie. 2019. Tagging address queries in maps search. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 9547–9551.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, and 1 others. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, and 1 others. 2018. Improving language understanding by generative pre-training.
- Benjamin J Radford. 2021. Regressing location on text for probabilistic geocoding. *arXiv preprint arXiv:2107.00080*.
- Zahra Rezaei Ghahroodi, Hassan Ranji, and Alireza Rezaee. 2024. Address matching using machine learning methods: An application to register-based census. *Statistical Journal of the IAOS*, 40(1):25–40.
- Kathryn M Rose, Joy L Wood, Sarah Knowles, Ricardo A Pollitt, Eric A Whitsel, Ana V Diez Roux, DongKeun Yoon, and Gerardo Heiss. 2004. Historical measures of social context in life course studies:

- retrospective linkage of addresses to decennial censuses. *International journal of health geographics*, 3(1):27.
- Gerard Rushton, Marc P Armstrong, Josephine Gittler, Barry R Greene, Claire E Pavlik, Michele M West, and Dale L Zimmerman. 2007. *Geocoding health data: the use of geographic codes in cancer prevention and control, research and practice*. CRC Press.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, and 1 others. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Qwen Team and 1 others. 2024. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2(3).
- Tencent Maps. 2025. *Geocoding API*. <https://lbs.qq.com/service/webService/webServiceGuide/address/Geocoder>.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, and 1 others. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- 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.
- Daniel Wilson, Xiaohan Zhang, Waqas Sultani, and Safwan Wshah. 2024. Image and object geolocalization. *International Journal of Computer Vision*, 132(4):1350–1392.
- Shitao Xiao, Zheng Liu, Peitian Zhang, Niklas Muenighoff, Defu Lian, and Jian-Yun Nie. 2024. C-pack: Packed resources for general chinese embeddings. In *Proceedings of the 47th international ACM SIGIR conference on research and development in information retrieval*, pages 641–649.
- Shixiong Xu, Chenghao Zhang, Lubin Fan, Gaofeng Meng, Shiming Xiang, and Jieping Ye. 2024. Addressclip: Empowering vision-language models for city-wide image address localization. In *European Conference on Computer Vision*, pages 76–92. Springer.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, and 1 others. 2024. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*.
- Zhengcong Yin, Daniel W Goldberg, Binbin Lin, Bing Zhou, Diya Li, Andong Ma, Ziqian Ming, Heng Cai, Zhe Zhang, Shaohua Wang, and 1 others. 2025. Toward building next-generation geocoding systems: a systematic review. *arXiv preprint arXiv:2503.18888*.
- Zhengcong Yin, Diya Li, and Daniel W Goldberg. 2023. Is chatgpt a game changer for geocoding—a benchmark for geocoding address parsing techniques. In *Proceedings of the 2nd ACM SIGSPATIAL International Workshop on Searching and Mining Large Collections of Geospatial Data*, pages 1–8.
- Zhengcong Yin, Andong Ma, and Daniel W Goldberg. 2019a. A deep learning approach for rooftop geocoding. *Transactions in GIS*, 23(3):495–514.
- Zhengcong Yin, Chong Zhang, Daniel W Goldberg, and Sathya Prasad. 2019b. An nlp-based question answering framework for spatio-temporal analysis and visualization. In *Proceedings of the 2019 2nd international conference on geoinformatics and data analysis*, pages 61–65.
- Wei Zhang and Judith Gelernter. 2014. Geocoding location expressions in twitter messages: A preference learning method. *Journal of Spatial Information Science*, (9):37–70.
- Lei Zou, Nina SN Lam, Shayan Shams, Heng Cai, Michelle A Meyer, Seungwon Yang, Kisung Lee, Seung-Jong Park, and Margaret A Reams. 2020. Social and geographical disparities in twitter use during hurricane harvey. In *Social Sensing and Big Data Computing for Disaster Management*, pages 103–121. Routledge.

Table 6: Quantitative Results of Each Method on the Singapore Dataset

Data	Methods	ADD (m)	Acc@100 (%)	Acc@200 (%)	Acc@500 (%)
Base Data	NER+Levenshtein	4347.6	24.1	33.1	46.1
	Vector-Retrieval@Top1	2964.6	54.0	60.8	68.1
	Vector-Retrieval@Top5+Reranker	2790.5	59.2	64.4	70.7
	ReaGeo (Ours)	<b>2152.3</b>	<b>65.9</b>	<b>69.6</b>	<b>74.2</b>
Anchor Offset Data	NER+Levenshtein	5826.0	3.7	10.6	33.5
	Vector-Retrieval@Top1	3455.6	7.6	19.7	59.4
	Vector-Retrieval@Top5+Reranker	3662.6	7.8	19.7	58.8
	ReaGeo (Ours)	<b>2395.9</b>	<b>10.6</b>	<b>25.3</b>	<b>70.8</b>

Table 7: Details of Data Partitioning for the Singapore Dataset

Data	Training	Test	Total
Base	60,000	6,669	66,669
Anchor Offset	60,000	6,669	66,669

## A Quantitative Evaluation on a Second Dataset

To further demonstrate the advantages of the proposed method, an additional dataset covering the city of Singapore is constructed. The dataset will be publicly released on <https://github.com/ziyue246/ReaGeo> to facilitate future research. As the second dataset in the experiments, it follows the same partitioning strategy as the main dataset, consisting of the base data and the anchor-offset data. Each subset contains 60,000 training samples and 6,669 test samples, as presented in Table 7.

The quantitative evaluation results of each method on the Singapore dataset are shown in Table 6. Notably, some methods are not included in the tables because Qwen3-Max largely relied on guesswork, resulting in substantial errors, and the APIs of Baidu Maps and Tencent Maps do not support English addresses. The experimental results indicate that the proposed end-to-end method still outperforms other retrieval-based approaches in terms of performance metrics.