

Towards Generative Location Awareness for Disaster Response: A Probabilistic Cross-view Geolocalization Approach[★]

Hao Li^{a,*}, Fabian Deuser^{b,c}, Wenping Yin^d, Steffen Knoblauch^{e,f}, Wufan Zhao^g, Filip Biljecki^{h,i}, Yong Xue^d and Wei Huang^j

^aDepartment of Geography, National University of Singapore, Singapore, 117570, Singapore

^bInstitute of Distributed Intelligent Systems, University of the Bundeswehr Munich, Neubiberg, 85579, Bavaria, Germany

^cProfessorship of Big Geospatial Data Management, School of Engineering and Design, Technical University of Munich, Munich, 85521, Bavaria, Germany

^dSchool of Environment and Spatial Informatics, China University of Mining and Technology, Xuzhou, 221116, Jiangsu, China

^eInterdisciplinary Center for Scientific Computing, Heidelberg University, Heidelberg, 69120, Baden-Württemberg, Germany

^fHeidelberg Institute for Geoinformation Technology, Heidelberg University, Heidelberg, 69120, Baden-Württemberg, Germany

^gUrban Governance and Design Thrust, Hong Kong University of Science and Technology (Guangzhou), Guangzhou, 511453, Guangdong, China

^hDepartment of Architecture, National University of Singapore, Singapore, 117566, Singapore

ⁱDepartment of Real Estate, National University of Singapore, Singapore, 119245, Singapore

^jCollege of Surveying and Geo-Informatics, Tongji University, Shanghai, 200092, China

ARTICLE INFO

Keywords:

Cross-View

Disaster Response

GeoAI

Climate Resilience

Generative Geolocalization

Explainability

ABSTRACT

As Earth's climate changes, it is impacting disasters and extreme weather events across the planet. Record-breaking heatwaves, drenching rainfalls, extreme wildfires, and widespread flooding during hurricanes are all becoming more frequent and more intense. Rapid and efficient response to disaster events is essential for climate resilience and sustainability. A key challenge in disaster response is to accurately and timely identify disaster locations to support decision-making and resources allocation. In this paper, we propose a Probabilistic Cross-view Geolocalization approach, called ProbGLC, exploring new pathways towards generative location awareness for rapid disaster response. Herein, we combine probabilistic and deterministic geolocalization models into a unified framework to simultaneously enhance model explainability (via uncertainty quantification) and achieve state-of-the-art geolocalization performance. Designed for rapid disaster response, the ProbGLC is able to address cross-view geolocalization across multiple disaster events as well as to offer unique features of probabilistic distribution and localizability score. To evaluate the ProbGLC, we conduct extensive experiments on two cross-view disaster datasets (i.e., MultiIAN and SAGAINDisaster), consisting diverse cross-view imagery pairs of multiple disaster types (e.g., hurricanes, wildfires, floods, to tornadoes). Preliminary results confirms the superior geolocalization accuracy (i.e., 0.86 in Acc@1km and 0.97 in Acc@25km) and model explainability (i.e., via probabilistic distributions and localizability scores) of the proposed ProbGLC approach, highlighting the great potential of leveraging generative cross-view approach to facilitate location awareness for better and faster disaster response. The data and code is publicly available at <https://github.com/bobleegogogo/ProbGLC>.

1. Introduction

Due to climate change and rapid urbanization, disasters are becoming more frequent and severe worldwide (Dietze and Ozturk, 2021; Nohrstedt et al., 2022). Compared with the average for the last 30 years (i.e., 1994-2023),

the major natural disasters in 2025 showed an 8% increase in frequency, a 74% increase in fatalities, a 42% more affected population, and a 49% increase in direct economic losses (Delforge et al., 2025). Effective and efficient disaster response depends on rapid situational awareness, by quickly detecting affected areas, estimating damage, and understanding access constraints on the ground. A preliminary and key prerequisite for rapid situational awareness is the accurate localization of impacted areas and affected communities (Li et al., 2025). Precise location information supports timely decision-making, effective allocation of emergency resources, and efficient recovery efforts (Cai et al., 2018; Huang et al., 2018; Zou et al., 2023; Feng et al., 2022; Macatulad and Biljecki, 2024). When disaster locations are correctly identified, rescue teams can better evaluate the extent of damage, prioritize urgent needs, and deliver assistance to the most vulnerable. In this context, disaster geolocalization has emerged as a critical task.

Recently, geolocalization has progressed along two main directions: text-based and image-based methods. Text-based

[★]This is the accepted manuscript (postprint) of the following article: Li H, Deuser F, Yin W, Knoblauch S, Zhao W, Biljecki F, Xue Y, Huang W (2026). Towards generative location awareness for disaster response: A probabilistic cross-view geolocalization approach. *ISPRS Journal of Photogrammetry and Remote Sensing*, 237, 130–145. The Version of Record is available at <https://doi.org/10.1016/j.isprsjprs.2026.03.050>. © 2026 Elsevier B.V. on behalf of International Society for Photogrammetry and Remote Sensing, Inc. (ISPRS). This manuscript version is made available under the CC BY-NC-ND 4.0 license, <https://creativecommons.org/licenses/by-nc-nd/4.0/>.

*Corresponding author

✉ hao.li@nus.edu.sg (H. Li); fabian.deuser@tum.de (F. Deuser); yin@cumt.edu.cn (W. Yin); steffen.knoblauch@uni-heidelberg.de (S. Knoblauch); wufanzhao@hkust-gz.edu.cn (W. Zhao); filip@nus.edu.sg (F. Biljecki); yx9@hotmail.com (Y. Xue); wei_huang@tongji.edu.cn (W. Huang)
ORCID(s): 0000-0002-6336-8772 (H. Li); 0000-0003-4511-4223 (F. Deuser); 0009-0002-2056-0091 (W. Yin); 0000-0002-6229-7749 (F. Biljecki)

approaches rely on volunteered geographic information (VGI) imagery, social media posts, and other textual data sources (Zou et al., 2018; Lam et al., 2023; Yin et al., 2025c). Natural language processing techniques extract toponyms, link them to gazetteers, and infer locations from contextual clues when coordinates are missing (Hu et al., 2023b). These methods enable rapid identification of affected regions from large volumes of online text and provide valuable real-time information during disasters. Image-based methods commonly include ground-view image retrieval, cross-view image retrieval, and two-dimensional-three-dimensional (2D-3D) matching, designed to handle different types of image data. Earlier studies emphasized feature matching (Workman et al., 2015), whereas recent work employs deep learning to learn shared representations between ground and overhead images (Lin et al., 2015). Among these, cross-view geolocation is particularly important in disaster contexts because it links crowdsourced ground images with satellite imagery to pinpoint damaged areas (Deuser et al., 2023; Zeng et al., 2025; Li et al., 2025). For example, street-view imagery (SVI) collected after Hurricane Ian, 2022, provided rapid situational awareness along Florida’s coast; however, moving from nadir satellite views to ground SVI perspectives remains challenging for automated geolocation and requires cross-view techniques (Manzini et al., 2023). Although advances in cross-view geolocation have improved disaster monitoring and response, the majority of existing methods remains black boxes as deterministic approaches with limited explainability.

Explainability has become a central concern in geospatial artificial intelligence (GeoAI), as black-box models limit transparency and hinder operational deployment (Hsu and Li, 2023; Li et al., 2024; Liu et al., 2024). To address this issue, a variety of explainable AI (XAI) methods have been proposed. Postprocessing techniques such as saliency maps, gradient-based attribution, SHAP, and Grad-CAM aim to highlight the input regions or variables most responsible for a model’s output (Höhl et al., 2024; Foroutan et al., 2025). Attention-based visualization and concept-level reasoning further provide insights into the internal decision process of AI models (Dehimi and Tolba, 2024). In addition, model-intrinsic approaches, including interpretable architectures, physics-aware designs, and rule-based systems, embed explainability directly into the model design (Hu et al., 2023a). While these advances show strong potential for geospatial applications, their feasibility and applicability in disaster-related geolocation remain limited: explanations are often coarse, unstable across scenarios, or difficult to interpret. This gap underscores the urgent need for domain-oriented explainability techniques that can deliver trustworthy and actionable explainability insights in disaster response.

In this context, this paper proposes a **Probabilistic Cross-view GeoLoCalization** approach, namely **ProbGLC**, aiming at facilitating generative location awareness for rapid disaster response. In ProbGLC, we carefully integrate the probabilistic and deterministic geolocation approaches into a unified framework, which simultaneously ensures

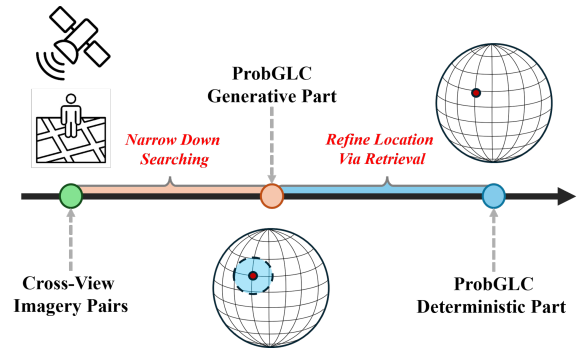


Figure 1: Illustration of a simplified diagram of ProbGLC consisting of 1) generative and 2) deterministic geolocation parts.

model explainability (via uncertainty quantification) and achieves state-of-the-art geolocation performance. Tailored for rapid disaster response, ProbGLC excels the deterministic geolocation approach by helping to narrow down the area of interest (AOI) as shown in Figure 1, which is extremely important in disaster response scenario as every second counts when people are waiting for help. The major contributions of ProbGLC are summarized as follows:

- To provide competitive geolocation performance across multiple disaster events (e.g., hurricanes, wildfires, floods, and tornadoes).
- To offer unique features of density distribution and localizability scores towards a more uncertainty-aware geolocation process.
- To improve the deterministic approach with a generative probabilistic prediction via Riemannian Flow Matching directly on the Sphere.

The remainder of this paper is organized as follows: Section 2 reviews related work on state-of-the-art GeoAI methods for disaster geolocation and uncertainty quantification. Section 3 describes the proposed ProbGLC method together with the generative and deterministic geolocation approaches. Section 4 presents experimental results on two openly available disaster geolocation datasets. Section 5 discusses key findings as well as limitations and future directions, and Section 6 concludes the paper by highlighting the main contributions.

2. Related Work

2.1. Climate Change and Disaster Resilience

Climate change is intensifying the frequency and severity of extreme weather events, including hurricanes, heavy rainfall, and coastal storm surges (Ellis et al., 2024). These phenomena are no longer isolated anomalies but increasingly recurring patterns that pose long-term risks to human settlements and infrastructure (Jones et al., 2012). Rising global temperatures have led to accelerated sea-level rise and increased atmospheric moisture, both of which contribute

directly to the likelihood and impact of urban flooding events caused by climate-induced disasters (e.g., Hurricanes, Heavy rainfalls, etc.) (Hulme et al., 1999). As a result, the intersection of climate science, disaster response, and geospatial analytics has become a crucial area for research and policy intervention.

Disaster risk and resilience received insufficient emphasis in sustainable development planning, despite the obvious relationship between disasters and development (Hallegatte et al., 2016; Cai et al., 2018). Disaster resilience building (e.g., prevention, preparedness, and disaster response systems) for predictable events like the major climate extreme weather conditions such as cyclones, hurricanes, wildfires, heat and cold waves (Li et al., 2020; Zhou et al., 2025; Shives et al., 2025; Russo and Domeisen, 2023), helps to protect both human and economic assets. With the growing availability of remote sensing data and advanced GeoAI methods, disaster resilience can be enhanced through near real-time monitoring, fine-grained damage mapping, and predictive modeling (Li et al., 2025). In recent years, events like Hurricane Ian (Yin et al., 2025b) and California wildfires (Shives et al., 2025) have highlighted the need for integrated systems that combine environmental sensing, citizen-contributed data, and AI-powered analytics to improve situational awareness and guide emergency response.

2.2. GeoAI for Disaster Mapping and Localization

GeoAI, the integration of geospatial data and artificial intelligence, has become a transformative force in disaster mapping and localization (Mai et al., 2025; Zou et al., 2023). By combining remote sensing, computer vision, and machine learning, GeoAI enables rapid extraction of critical information from multimodal data sources. In particular, satellite-based Earth Observation (EO) and SVI provide complementary views: the former ensures wide-area coverage for macro-scale assessment, while the latter offers detailed, ground-level context (Biljecki and Ito, 2021; Zhang et al., 2024). GeoAI models can fuse these views to support timely, scalable, and accurate disaster mapping across spatial resolutions and sensor modalities.

GeoAI plays a critical role in identifying the location of ground-level visual content, especially when metadata is missing or unreliable. Cross-view geolocalization, powered by deep feature matching between ground-level and overhead imagery, allows GeoAI models to infer precise locations of images or videos. This is particularly valuable in disaster scenarios where citizen-generated content or drone imagery lacks Global Positioning System (GPS) information. Yin et al. (2025b) develops two novel GeoAI methods, StaGeo and TriGeo, to improve the cross-view geolocalization accuracy of disaster-related VGI imagery based on RSI by using SVI as a bridge. Vivanco Cepeda et al. (2023) proposes GeoCLIP, a CLIP-inspired image-to-GPS retrieval framework that aligns images with their geographic locations using continuous spatial encoding and multi-scale feature representations for accurate worldwide geolocalization.

Moreover, recent trend in developing multi-modal and vision-language models (VLMs) sheds stimulating light on integrating large-scale remote sensing with semantic or multimodal reasoning for disaster understanding. Gupta et al. (2019) established a large-scale disaster mapping benchmark focusing on building damages and change detection, which largely facilitates relevant disaster model development. Wang et al. (2025) further curated a remote sensing vision-language dataset, namely DisasterM3, for global-scale disaster response, feature multi-modal visual perception and reasoning in various disaster scenes. More recent, the potential of VLMs in coordinating multi-view imagery (i.e., SVI and satellite imagery) is confirm in tackling sustainable urban development challenges (Wang et al., 2026). However, we notice a remaining research gap on examining the uncertainty and explainability of GeoAI-based disaster response models, which largely motivates the ProbGLC approach in this paper.

2.3. Uncertainty and Explainability of GeoAI

Uncertainty and explainability have emerged as two central research themes in GeoAI, reflecting growing demands for reliable, transparent, and accountable decision-making in geospatial science (Hu et al., 2024; Li, 2025). GeoAI models frequently operate on heterogeneous, multi-source spatial data, introducing uncertainty from data noise, temporal dynamics, model assumptions, and output variability. At the same time, the complex nature of deep models often leads to limited explainability. Tackling both challenges is essential for enhancing model robustness, fostering trust, and enabling informed decisions in critical applications such as disaster response and climate risk management.

Uncertainty has long been recognized as a key issue in geographic information science, as both data and models are generally considered to carry uncertainty (Zhang and Goodchild, 2002). Many researchers highlight the importance of making uncertainty interpretable and actionable, where Bayesian approximation and ensemble learning are two commonly applied methods for uncertainty quantification. Bayesian techniques for uncertainty quantification include Monte Carlo dropout (Srivastava et al., 2014), Markov chain Monte Carlo (Kupinski et al., 2003), and Variational inference (Blei et al., 2017). Common ensemble learning methods include deep ensemble, deep ensemble Bayesian/Bayesian deep ensemble, and uncertainty in Dirichlet deep networks (Abdar et al., 2021). Moving forward, integrating explainability and uncertainty handling will be crucial for advancing the scientific and practical impact of GeoAI.

However, while uncertainty quantification is crucial, the explainability of GeoAI models remains a challenge, especially in spatial applications like environmental monitoring and disaster management (Li, 2025). Traditional “black-box” models often fail to provide intuitive explanations of their decision-making processes, making it difficult to trust their outputs. To address this, research has increasingly focused on post-hoc explainability methods, such as local

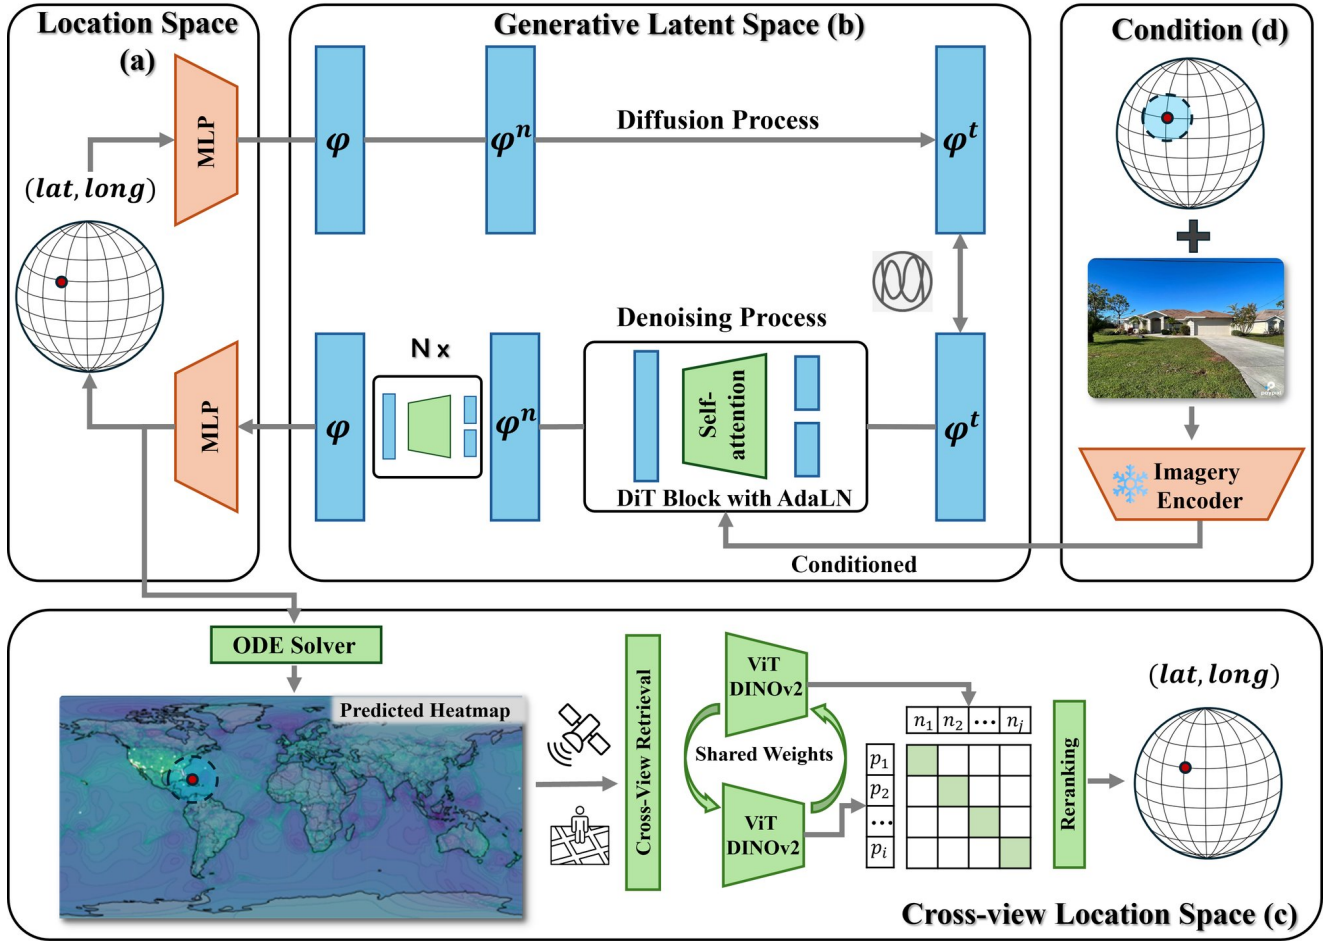


Figure 2: Overview of the **Probabilistic Cross-view GeoLocalization (ProbGLC)** approach. The ProbGLC consists of mainly four part: (1) Location Space with geographical coordinates (lat and long) of disaster-related imagery; (2) Generative Latent Space where the generative model is trained to learn a latent space for probabilistic geolocation tasks; (3) Cross-view Location Space with the deterministic retrieval approach is trained to refine the location predictions; (4) Condition of disaster-related imagery embeddings and model hyperparameter. Herein, φ refers to diffusion model parameters, p_i and n_i are cross-view imagery embeddings.

interpretable model-agnostic explanations (Ribeiro et al., 2016), Shapley additive explanations (SHAP) (Lundberg and Lee, 2017), and class activation mapping (CAM) (Zhou et al., 2016), which provide insights into model behavior by attributing output decisions to specific input features. However, these methods primarily focus on interpreting outputs rather than understanding the internal mechanisms of geospatial models. Therefore, there is growing emphasis on developing intrinsically interpretable models that integrate spatial knowledge, physical principles, and domain expertise, while also improving uncertainty and explainability methods in GeoAI models.

3. Methodology

3.1. Notation and Mission Statement

Given a set of disaster-related VGI imagery $\{L_c\}$ and satellite imagery $\{L_s\}$, we aim to predict the most likely location x_0 (longitude and latitude) where the imagery pair

was taken. Unlike the pure deterministic approach, we model the conditional probabilistic distribution $P(x_n | \mathbf{c})$ as a generative process, where x_n can be any point on Earth, modeled as the unit sphere \mathbb{S}^2 in \mathbb{R}^3 , and \mathbf{c} is the imagery embedding of the VGI imagery used as a query input. Herein, the random noise is denoted as ϵ and the noisy coordinates as x_t for a timestep t . Moreover, φ refers to a set of trainable parameters in the neural network (Figure 2).

Based on the probabilistic distribution $P(x_n | \mathbf{c})$, we then run the cross-view imagery retrieval as a deterministic process, but only within a refined area of interest radiused with \mathbf{r} . This deterministic geolocation process learns a cross-view embedding space \mathbb{R}_{CV} so that the exact geographical location of VGI imagery $\{L_c\}$ can be retrieved based on its similarities to geo-tagged satellite imagery $\{L_s\}$ in the learned embedding space. Figure 2 shows how we achieve the generative location awareness with a probabilistic geolocation approach, leveraging cross-view imagery and state-of-the-art GeoAI models to facilitate the critical

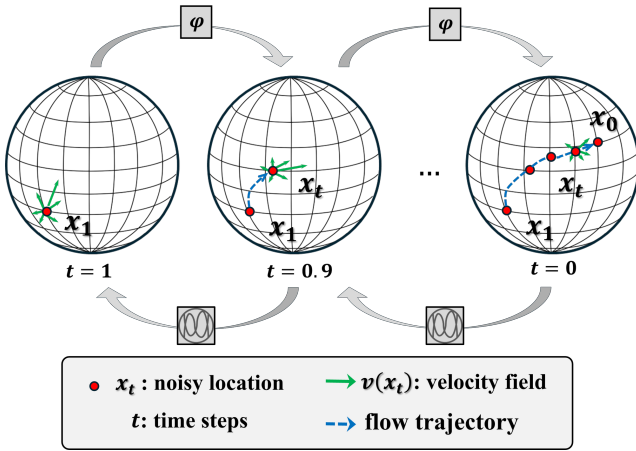


Figure 3: RFM on the sphere. Visualization of the generative flow process where noisy locations evolve along a learned velocity field toward their denoised positions on the spherical manifold.

disaster response scenario. In the rest of the section, we will elaborate on the detailed methodology and design specifics.

3.2. Generative Probabilistic Geolocalization

In the disaster response context, modeling geolocalization as a generative process has mainly two benefits: 1) to explicate and quantify uncertainty of visual geolocalization which mostly remains a black-box and deterministic process, 2) to offer a generalizable approach of addressing spatial ambiguity across multiple disaster types. In the rest of this section, we will elaborate on how **ProbGLC** models the probabilistic geolocalization process using a Riemannian Flow Matching diffusion approach, which is inspired by Dufour et al. (2024).

3.2.1. Riemannian Flow Matching on the Sphere

The previous steps model the classic imagery geolocalization task as a generative process that can be trained with state-of-the-art diffusion training pipelines. Figure 2 shows the overview workflow in the ProbGLC with part (b) refers to a generative diffusion model for imagery geolocalization. Specifically, we employed a Riemannian Flow Matching (RFM) approach from Dufour et al. (2024) to extend the classic geographic diffusion process (see Appendix A) allowing training and inferencing directly on the Sphere as shown in Figure 3.

Flow Matching in the Euclidean space: Classic diffusion process is built on Denoising Diffusion Probabilistic Models (DDPM). DDPM is useful but can be expensive and difficult to train as the convergence is heavily dependent on the noise schedule function $\beta(t)$. Flow matching offers an elegant solution to this problem (Lipman et al., 2022).

As a context, in diffusion models, the forward process adds noise, and the backward process removes noise. Both the forward and backward processes are regarded as Stochastic Differential Equation (SDEs), though the forward process is integrable in closed-form, so it can be done at no computational cost. However, the backward denoising process

is not integrable in closed-form, so it must be integrated step-by-step using standard SDE solvers, which makes it computationally expensive. In flow-based diffusion models, the forward process will be a deterministic flow along a time-dependent vector field $v(x_t)$, and the backward process is redefined also a deterministic flow along the very same vector field, but backwards. To this end, both processes become Ordinary Differential Equation (ODE). This means if the vector field is well-behaved, the ODE will also be well-behaved, making the training easier and faster.

Given the time-dependent vector field $v(x_t)$, it can be calculated as the partial derivative of the noisy location (namely Equation 13) over the timestep t :

$$v(x_t) = \frac{d\mathbf{x}_t}{dt} = \frac{d\beta(t)}{dt} (\epsilon - \mathbf{x}_0), \quad (1)$$

The training target of flow matching is then to minimize the prediction error of $\varphi(\mathbf{x}_t | c)$ to the vector field similarly as in Equation 18:

$$\mathcal{L}_{FM} = \mathbb{E}_{\mathbf{x}_0, c, \epsilon, t} [|\varphi(\mathbf{x}_t | c) - v(\mathbf{x}_t)|^2], \quad (2)$$

Inference Process: For inference, one can then use ODE solver initialized at a random location $\mathbf{x}_1 = \epsilon$, integrating backward from $t = 1$ to $t = 0$ using the trained parameters $\varphi(\mathbf{x}_t | c)$:

$$\mathbf{x}_{t-\Delta t} = \mathbf{x}_t - \varphi(\mathbf{x}_t | c) \Delta t \quad (3)$$

But noticed this is still in the Euclidean space, and one will have to project the final location \mathbf{x}_0 back to the Spherical coordinates of \mathbb{S}^2 .

Extend to Riemannian Flow Matching on the Sphere: To constrain the noisy location \mathbf{x}_t always on the Sphere \mathbb{S}^2 without of projections, the trick is to extend the aforementioned flow matching to Riemannian manifolds (Chen and Lipman, 2024) which ensure both the location variable \mathbf{x}_t and the noise ϵ remain on \mathbb{S}^2 . The former one is easy as the natural format of a location is in spherical coordinates, and the latter one can be achieved by manipulating Equation 13 as follows:

$$\mathbf{x}_t = \exp_{\mathbf{x}_{t-1}} (\beta(t) \log_{\mathbf{x}_{t-1}} (\epsilon)), \quad (4)$$

Where $\log_{\mathbf{x}_{t-1}}$ is the logarithmic map, mapping a point on \mathbb{S}^2 to the tangent space at a previous timestep \mathbf{x}_{t-1} , and $\exp_{\mathbf{x}_{t-1}}$ refers to the exponential map back to the Riemannian manifold. For more details about the projection, see Appendix B.

This representation will allow us to keep the noisy coordinates always along the geodesic, and let us rewrite the vector field on the tangent space of \mathbf{x}_t as follows:

$$v(\mathbf{x}_t) = \frac{d\mathbf{x}_t}{dt} = \frac{d\beta(t)}{dt} \log_{\mathbf{x}_t} (\mathbf{x}_0), \quad (5)$$

Where $\log_{x_t}(x_0)$ refers to the tangent vector at \mathbf{x}_t pointing along the geodesic from \mathbf{x}_0 to ϵ (as shown in Figure 3). Similarly to Equation 20, we can predict the velocity field by minimizing the prediction error constrained by the Riemannian manifold:

$$\mathcal{L}_{RFM} = \mathbb{E}_{\mathbf{x}_0, \mathbf{c}, \epsilon, t} \left[\|\varphi(\mathbf{x}_t | \mathbf{c}) - v(\mathbf{x}_t)\|_{\mathbf{x}_t}^2 \right], \quad (6)$$

Where the only difference compared to Equation 20 is the $\|\cdot\|_{\mathbf{x}_t}$ denoting the norm induced by the Riemannian metric on the tangent space at \mathbf{x}_t .

Probabilistic Inference on the Sphere: To predict with RFM, once the $\varphi(\mathbf{x}_t | \mathbf{c})$ is trained, we can use it to compute the geolocalization probability of any location on the Sphere as $P(x_n | \mathbf{c})$ conditional on a query imagery embedding c . This enables us to explicate and quantify the uncertainty of visual geolocalization, contrasting with classic approaches, which mostly remain deterministic. The probability distribution is calculated in the following way:

$$\log P(x_n | \mathbf{c}) = \log P_\epsilon(x(1) | \mathbf{c}) - f(1), \quad (7)$$

Where P_ϵ is the known distribution of the pure noise ϵ . And $f(t)$ is a function that accumulates the negative divergence of the velocity field along the trajectory x_t using the following ODEs:

$$\frac{d}{dt} \begin{bmatrix} x_t \\ f(t) \end{bmatrix} = \begin{bmatrix} \varphi(\mathbf{x}(t) | \mathbf{c}) \\ -\text{div} \varphi(\mathbf{x}_t | \mathbf{c}) \end{bmatrix}, \quad \text{with} \quad \begin{bmatrix} \mathbf{x}_0 \\ f(0) \end{bmatrix} = \begin{bmatrix} \mathbf{x}_n \\ 0 \end{bmatrix}. \quad (8)$$

For a detailed explanation, please refer to the original implementation in TorchDiffEq¹. This probabilistic inference directly on the sphere finally allows us to quantify the uncertainty and calculate the "localizability" score of the generative geolocalization process. Following the idea from Dufour et al. (2024), we define the "localizability" score as a measure of spatial ambiguity in geolocalization using the following equation:

$$\text{Localizability}(\mathbf{c}) = \int_{\mathbb{S}^2} p(\mathbf{x}_n | \mathbf{c}) \log_2 P(\mathbf{x}_n | \mathbf{c}) d\mathbf{x}_n, \quad (9)$$

Where the localizability of an image c is the negative entropy of the predicted distribution. In practice, we estimate this integral with Monte-Carlo sampling (Metropolis and Ulam, 1949) using 10,000 samples.

3.3. Cross-view Deterministic Geolocalization

In this work, **ProbGLC** takes the generative probabilistic geolocalization as a first step, then combines with a deterministic cross-view geolocalization approach only with a certain radius thresholding by \mathbf{r} . Herein, the consideration is

mainly twofold: 1) though the RFM approach offers a unique feature of uncertainty quantification and spatial ambiguity measurement, the deterministic approach still provide a more fine-grained geolocalization accuracy comparing to the generative approach with the cost of additional computational effort; 2) taking advantage of the probabilistic distribution generated from the RFM approach, we can improve the deterministic approach by narrowing down the area for imagery retrieval, which is extremely important in disaster response scenario as every second counts when people are waiting for help. The rest of this section will elaborate on the deterministic cross-view geolocalization approach.

3.3.1. Siamese Cross-view Image Encoder

Herein, we use a Siamese Vision Transformer with DINOv2 (i.e., ViT DINOv2) encoder to learn a cross-view embedding space \mathbb{R}_{CV} so that the exact geographical location of VGI imagery $\{L_c\}$ can be retrieved based on its similarities to geo-tagged satellite imagery $\{L_s\}$. Specifically, ViT DINOv2 treats each image as a sequence of fixed-size patches and captures global dependencies through a multi-layer self-attention (Dosovitskiy et al., 2021). As a state-of-the-art vision model, ViT DINOv2 provides a strong and generalizable backbone for the deterministic cross-view geolocalization approach, where a range of previous studies have confirmed its semantic representation capability, robustness to large viewpoint changes, and cross-domain generalizability (Oquab et al., 2024; Yin et al., 2025b).

Moreover, we adopt a Siamese design for learning the cross-view embedding space \mathbb{R}_{CV} . As shown in the Figure. 2, the same backbone encoder (i.e., ViT DINOv2) is used, but with independent weights for each data modality. This design addresses the issue of different image sizes across modalities, where RSI and VGI images can share a similar encoder. Next, the weight-sharing strategy across other modalities further guarantees that both inputs are embedded into a common feature space under the same semantic criteria (Bromley et al., 1993), which is crucial for accurate cross-view imagery matching and retrieving. Finally, this Siamese imagery encoder, denoted as $f()$, is trained on the imagery pair of VGI and RSI with a contrastive learning objective as follows:

Given a context vector c , the positive sample is drawn from a conditional distribution $p(x|c)$, where $(N - 1)$ negative samples are drawn from the same distribution $p(x)$ but without condition. In this context, the probability of correctly selecting the positive samples can be formulated as follows:

$$P(C = \text{pos} | X, \mathbf{c}) = \frac{f(\mathbf{x}_{\text{pos}}, \mathbf{c})}{f(\mathbf{x}_{\text{pos}}, \mathbf{c}) + \sum_{j=1}^{N-1} f(\mathbf{x}_j, \mathbf{c})} \quad (10)$$

Here, N is the total number of samples in a batch, and $f(\mathbf{x}, \mathbf{c}) \propto \frac{P(\mathbf{x}|\mathbf{c})}{P(\mathbf{x})}$ is the similarity or scoring function between two samples.

Then, the contrastive InfoNCE loss can be calculated as follows:

¹<https://github.com/rmqichen/torchdiffeq>

$$\mathcal{L}_{\text{InfoNCE}} = -\mathbb{E} [\log p(C = \text{pos} | X, \mathbf{c})] \quad (11)$$

Where the objective of $\mathcal{L}_{\text{InfoNCE}}$ is to optimize the negative log probability of selecting the positive samples.

In this work, we use the InforNCE as our contrastive learning loss in both the pre-training and fine-tuning stages for deterministic cross-view geolocalization. During the fine-tuning, we take the model weights pre-trained on CVUSA data given its relatively large size and geographical closeness (Workman et al., 2015), then fine-tune the model on the cross-view imagery pairs collected from two disaster mapping datasets.

3.3.2. Self-Supervised Feature Learning with DINOv2

The cross-view image encoder used in this work is pretrained with DINOv2, a state-of-the-art self-supervised learning method based on the teacher-student architecture (Dosovitskiy et al., 2021). Herein, DINOv2 enables the model to learn semantically rich and domain-invariant features by enforcing consistency between representations of diverse augmented views of the same image, without additional supervision signals. This is achieved by minimizing a cross-view consistency loss between the teacher and student outputs (Oquab et al., 2024). The self-supervised loss function for DINO is defined as follows:

$$\mathcal{L}_{DINO} = - \sum_{i=1}^N P^{(t)}(x_i) \cdot \log P^{(s)}(x_i) \quad (12)$$

Where $P^{(t)}(x_i)$ and $P^{(s)}(x_i)$ are the teacher and student output distributions for the i -th feature, respectively, and N is the number of feature dimensions. This training objective optimizes the student network to align its outputs with those of the teacher across varying augmentations, resulting in stable and semantically-consistent representations of input data from diverse view angles.

In this context, DINOv2 enhances the ViT backbone's generalization ability across heterogeneous modalities from ground-view to head-view, which often differ in their resolution, granularity and scene complexity. Compared to CNN-based models, ViT pretrained with DINOv2 tends to reserve long-range dependencies and semantic coherence, making it particularly effective for aligning multi-view features (Oquab et al., 2024). In this context, we use the same encoder architecture to encode both VGI and RSI imagery to ensure a consistent feature extraction and facilitate efficient joint training, ultimately improving geolocalization accuracy and reliability in complex disaster response scenarios.

Moreover, to further improve the cross-view geolocalization accuracy, we add an anchor-based reranking strategy (Deuser et al., 2024) as a post-processing step after the initial retrieval (See Appendix C). As a result, the reranking process enforces local neighborhood consistency and mitigates potential mismatches from the initial retrieval. For more details ablation, one can refer to Deuser et al. (2024).

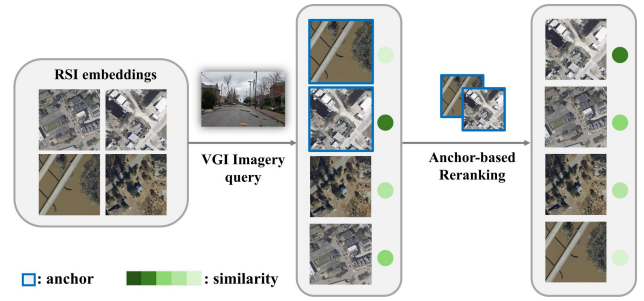


Figure 4: Anchor-based reranking for cross-view geolocalization. Retrieved RSI candidates are refined using anchor samples to improve similarity consistency with the ground-view query.

4. Experiment

4.1. Dataset Overview

In this work, we evaluate the proposed approach against two openly available disaster geolocalization datasets, namely MultiIAN and SAGINDisaster datasets, with a different geographical focus and disaster types as shown in Figure 5 and 6.

For the MultiIAN dataset (Yin et al., 2025b), it focuses on southeastern United States and western Cuba following the catastrophic Hurricane Ian, a Category 5 Atlantic hurricane occurring from late September to early October 2022. It reached peak intensity on September 28 near southwestern Florida, and became the third costliest weather disaster globally (Wikipedia, 2022). As a result, MultiIAN includes 5,706 imagery pairs containing both VGI and RSI imagery.

As for SAGINDisaster dataset (Yin et al., 2025a), multiple representative disaster types are selected across the continental United States including the Hawaiian Islands to diversify the potential cross-view disaster geolocalization scenarios. As a context, we take an advantage of the emergency response imagery (over 50 disaster events) released by the National Oceanic and Atmospheric Administration (NOAA) Remote Sensing Division², ranging from hurricanes, wildfires, floods, to tornadoes. Specifically, these disasters include the California Fire (2025) and Maui Fire (2023), the aforementioned Hurricane IAN (2025), the Nashville Tornadoes (2020) and Illinois Tornadoes (2015), and the Louisiana Flooding (2016) and Midwest U.S. Flooding (2015). In this context, SAGINDisaster results in a total of 2,080 cross-view imagery pairs (as shown in Figure 5), featuring the diversity and challenge of cross-view geolocalization in real-world disaster response scenarios.

VGI: VGI imagery was collected from multiple online platforms where people voluntarily shared images and posts during and after the disasters. For Hurricane Ian, the Paypixl platform provided a large number of geo-tagged photos. In addition, we included images from Twitter and Flickr. Though Paypixl and Flickr already contain location metadata, the Twitter samples required manual annotation to determine their geographic position. These crowdsourced data

²<https://storms.ngs.noaa.gov/>

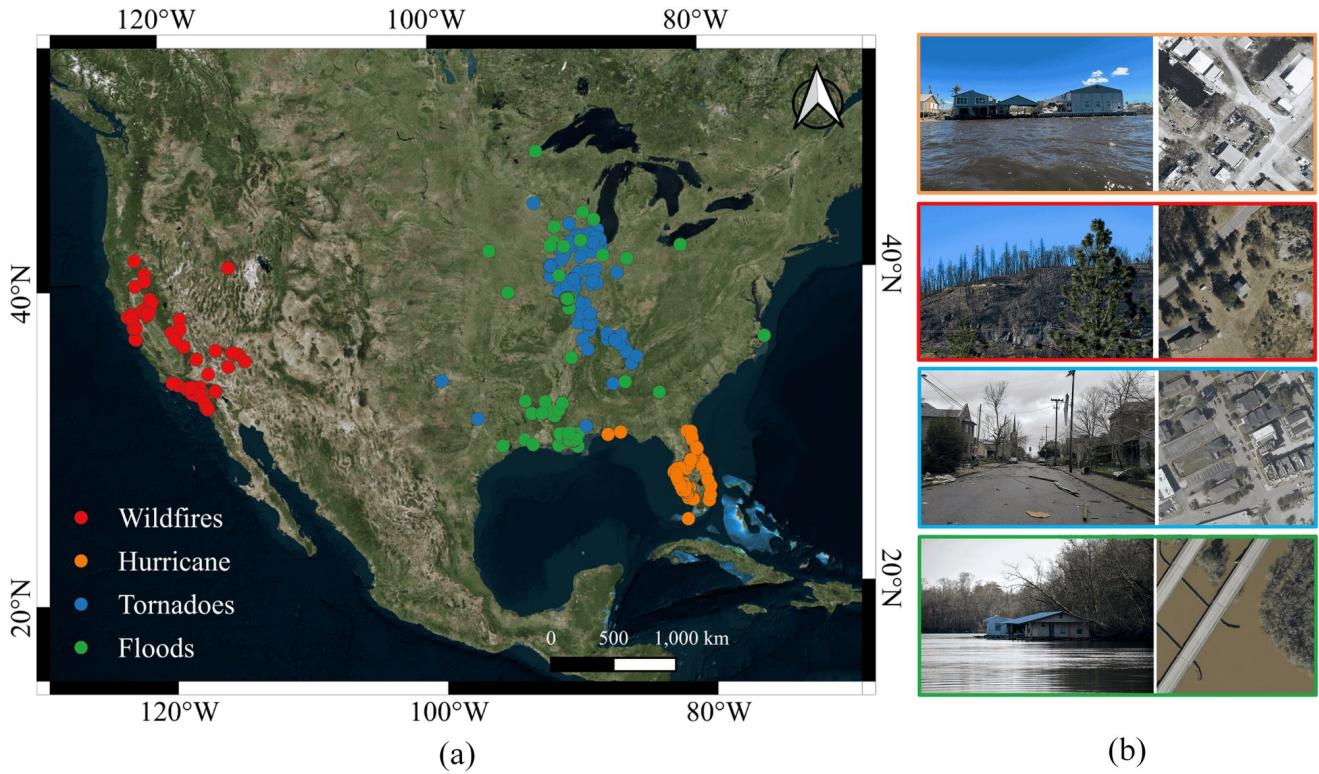


Figure 5: Spatial distribution and examples of the SAGINDisaster dataset. (a) Distribution of samples for different disaster types. (b) Example images of each disaster type, with border colors matching the categories in (a).

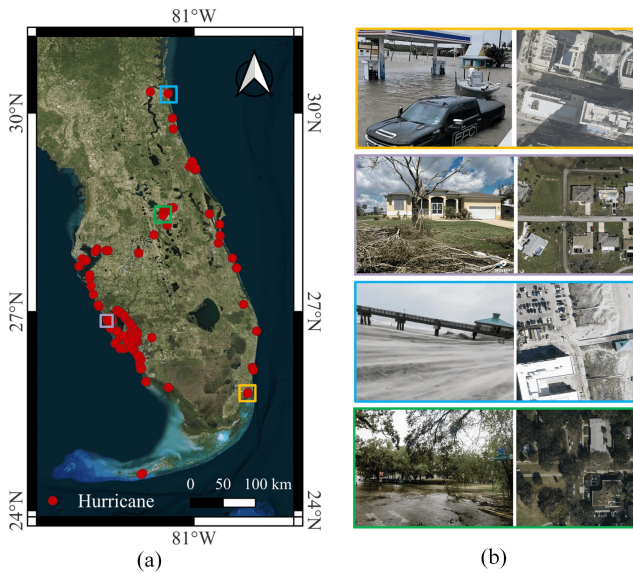


Figure 6: Spatial distribution and examples of the MultilAN dataset. (a) Distribution of samples for different disaster types. (b) Example scenes of VGI and Satellite image pairs, where colors refer to locations in (a).

reflect real on-site observations from affected communities, which are key for effective and rapid disaster responses.

RSI: High-resolution satellite images were sourced from the U.S. National Oceanic and Atmospheric Administration

(NOAA). These images were captured during the disaster periods, especially Hurricane Ian, and processed into accessible mosaics and tiles. With a ground sampling distance of 15–30 cm per pixel, the satellite data offer a detailed overview of damaged areas and provide a critical large-scale perspective to complement ground-level information. Serving as the most important reference imagery in cross-view geolocation, RSI plays a crucial role in aligning and anchoring information from a head-view perspective.

4.2. Evaluation Metrics and Implementations

The main idea of ProbGLC is to combine the probabilistic and deterministic geolocation approaches to simultaneously achieve uncertainty quantification and performance boosting. In this context, the experiment aims to validate a two-level improvement against a range of baseline models and variations. First, the improvement over pre-trained generative approaches, for example in Dufour et al. (2024) without fine-tuning for disaster-related VGI imagery; second, the improvement over either solely fine-tuned generative approaches or classic imagery retrieval approaches.

For generative baselines, we consider three types of pre-trained visual geolocation models, namely 1) Diffusion: the geographic diffusion approach in the Euclidean space (Appendix A), 2) Flow Matching in Euclidean space with time-dependent vector field, and 3) Riemannian Flow Marching (RFM) directly on the Sphere S^2 as suggested in Dufour et al. (2024). Herein, their pre-training weights

are obtained from three state-of-the-art global visual geolocalization datasets (i.e., YFCC (Thomee et al., 2016), iNaturalist (Van Horn et al., 2021), OSV-5M (Astruc et al., 2024)), respectively.

Moreover, to investigate the impact of combining probabilistic and deterministic geolocalization approaches, we carefully design a set of comparisons between the ProbGLC approach with either solely probabilistic or deterministic geolocalization approaches, especially with fine-tuning on the disaster-related cross-view datasets. We examine the pattern of our performance gain under different thresholds r , which is considered as a key hyperparameter of the ProbGLC approach to balance accuracy and explainability from either deterministic or probabilistic approaches. For deterministic cross-view geolocalization approaches, we include the state-of-the-art models as baselines as well, namely SAIG-D (Zhu et al., 2023) and TransGeo (ViT-based) (Zhu et al., 2022).

To evaluate the model performance, since the ProbGLC approach is one the very first to tackle the geolocalization as a generative process rather than classic imagery retrieval, we have to use a novel metric called Acc@K rather than the classic Recall@K as the main performance metric. Herein, Acc@K measures the proportion of correct matches (between predicted locations and exact locations) within a certain geographical radius. In this study, Acc@1km, Acc@25km, Acc@50km, and Acc@200km are used to measure the performance of both probabilistic and deterministic approaches. In addition, we also report the median and mean distances over the test set of our disaster cross-view datasets. All evaluations are performed using both 8:2 and 7:3 train-test splits for two datasets, respectively.

A unique metric of the ProbGLC approach is the localizability score featuring a quantitative measure of how difficult it is to geolocalize the disaster-related VGI imagery with the generative approach using the Riemannian Flow Marching. This localizability score in combination with the density maps (see in Figure 2) offers novel insights towards generative location awareness for disaster response.

Regarding the training and implementation, for ProbGLC variations, we use a learning rate of 8×10^{-4} and a step size (Δt) of 16 with the LAMB optimizer (You et al., 2019). This choice is motivated by the large batch size of 512. The learning rate is scheduled using cosine decay down to zero. All experiments are implemented in PyTorch (Paszke et al., 2019) and conducted on a DGX-2 system equipped with 16 GPUs (i.e., Nvidia V100), each with 32 GB of memory. For the ProbGLC architecture, we adopt the same network design as Dufour et al. (2024), consisting of an MLP with 8 layers and a hidden dimension of 256. The trained retrieval model follows the Sample4Geo architecture (Deuser et al., 2023) and employs a ConvNeXt-Base backbone (Liu et al., 2022). It is trained with a learning rate of 1×10^{-3} , a cosine decay learning rate scheduler, and a batch size of 128.

4.3. Baseline Geolocalization Results

Deterministic baselines: Table 1 presents the baseline performances of three state-of-the-art cross-view geolocalization approaches (i.e., TransGeo, SAIG-D, and Sample4Geo) across the two disaster datasets (i.e., SAGINDisaster, MultiIAN) using the generative evaluation metrics (Acc@K). Specifically, we take the Recall@1 match from the classic imagery-retrieval models as predicted locations then measure the distance to actual locations of VGI imagery.

As a baseline for deterministic approaches, ProbGLC demonstrates the strongest zero-shot performance across both datasets, showing notably higher retrieval accuracy and lower localization errors. For SAGINDisaster dataset (20% test), it achieves Acc@50km of 0.632 and Acc@200km of 0.659, while reducing the mean distance to 903.68 km and the median distance to only 3.40 km. Similarly for MultiIAN (20% test), Sample4Geo also excels with Acc@25km reaching 0.963 and Acc@200km reaching 0.990, alongside very low mean and median distances of 7.26 km and 2.40 km. In contrast, TransGeo and SAIG-D yield lower accuracies, such as Acc@50km of 0.339 and 0.387 on SAGINDisaster, with much higher mean distances exceeding 1500 km. However, a key finding herein is that all baseline approaches perform poorly for Acc@1km with an average accuracy level below 0.150 despite of datasets, which highlight the pressing need for closing this "last kilometer" gap in disaster geolocalization.

One can easily notice the fact that two datasets (i.e., SAGINDisaster, MultiIAN) are not equally difficult. As shown in Figure 5 and 6, SAGINDisaster includes multiple disaster types and events across the continental United States, therefore it is considered as more challenging than MultiIAN when it comes to geolocalization accuracy. This fact is also confirmed in Table 1, resulting in much lower Mean and Median Dist for MultiIAN that all falls in 50 km range. However, as a nature of deterministic approach, the lack of explainability and uncertainty quantification during cross-view geolocalization often makes performance boosting an exhaustive and time-demanding task.

Generative baselines: Table 2 and 3 provide baseline performance of three different variants of ProbGLC (e.g., Diffusion, Flow Matching, and Riemannian Flow Matching) across the SAGINDisaster and MultiIAN dataset, respectively.

Regarding the ProbGLC pre-training options, those baseline models pre-trained on YFCC and OSV-5M consistently lead to higher accuracy and lower localization errors compared to iNaturalist. In the SAGINDisaster dataset (Table 2), RFM-80 achieves Acc@50km of 0.101 with Mean/Median Dist of 941.71 km and 165.56 km on YFCC, and Acc@50km of 0.113 with Mean/Median Dist of 513.13 km and 105.02 km on OSV-5M. In contrast, baseline models pre-trained on iNaturalist perform poorly, with Acc@50km below 0.04 and Mean Dist exceeding 1760 km, even Diffusion-70 only reaches Acc@200km of 0.269, and FlowR3-70 has Mean Dist of 2493.61 km. As for the MultiIAN dataset (Table 3),

Table 1

Zero-shot retrieval results on the Multidisaster and IANdisaster datasets for 20% and 30% splits.

Dataset	Model	Split	Acc@1km	Acc@25km	Acc@50km	Acc@200km	Mean Dist (km)	Median Dist (km)
SAGINDisaster	TransGeo	20%	0.062	0.219	0.339	0.440	1517.75	789.76
	TransGeo	30%	0.045	0.191	0.300	0.426	1404.18	753.69
	SAIG-D	20%	0.053	0.226	0.387	0.423	1746.45	810.78
	SAIG-D	30%	0.040	0.186	0.285	0.413	1709.95	886.49
	Sample4Geo	20%	0.137	0.623	0.632	0.659	903.68	3.40
	Sample4Geo	30%	0.125	0.556	0.566	0.599	1302.82	3.79
MultiIAN	TransGeo	20%	0.132	0.583	0.735	0.946	39.91	3.34
	TransGeo	30%	0.130	0.606	0.762	0.947	38.07	3.39
	SAIG-D	20%	0.096	0.475	0.715	0.960	40.26	27.54
	SAIG-D	30%	0.095	0.494	0.710	0.961	39.76	26.87
	Sample4Geo	20%	0.130	0.963	0.967	0.990	7.26	2.40
	Sample4Geo	30%	0.117	0.954	0.961	0.988	8.38	2.41

Table 2

Zero-shot performance of RFM-70 and RFM-80 models on the SAGINDisaster dataset across different pre-training datasets.

Pre-training	Method	Acc@1km	Acc@25km	Acc@50km	Acc@200km	MeanDist (km)	MedianDist (km)
iNaturalist	Diffusion-70	0.000	0.008	0.035	0.269	1763.69	534.15
	Diffusion-80	0.000	0.005	0.024	0.231	1777.19	758.33
	FlowR3-70	0.000	0.014	0.040	0.212	2493.61	1221.62
	FlowR3-80	0.000	0.007	0.029	0.209	2305.87	1317.18
	RFM-70	0.000	0.005	0.016	0.205	1800.86	770.04
	RFM-80	0.000	0.007	0.019	0.202	2030.33	818.69
YFCC	Diffusion-70	0.000	0.038	0.080	0.516	1136.12	195.25
	Diffusion-80	0.000	0.034	0.084	0.524	926.89	195.12
	FlowR3-70	0.002	0.058	0.119	0.429	1065.15	221.77
	FlowR3-80	0.002	0.053	0.094	0.394	1060.01	226.22
	RFM-70	0.006	0.054	0.096	0.537	1078.54	187.37
	RFM-80	0.014	0.055	0.101	0.558	941.71	165.56
OSV-5M	Diffusion-70	0.000	0.021	0.099	0.752	515.18	106.70
	Diffusion-80	0.000	0.026	0.099	0.762	609.78	105.02
	FlowR3-70	0.000	0.021	0.090	0.607	566.05	126.63
	FlowR3-80	0.000	0.026	0.065	0.618	513.13	122.45
	RFM-70	0.000	0.042	0.111	0.582	643.31	158.68
	RFM-80	0.000	0.043	0.113	0.594	513.13	105.02

a similar trend is observed. RFM-80 reaches Acc@50km of 0.477 on OSV-5M with Mean/Median Dist of 99.93 km and 174.55 km, while on YFCC it achieves Acc@50km of 0.291 with Mean/Median Dist of 117.98 km and 623.79 km. iNaturalist-pre-trained models have lower accuracies, with Acc@50km below 0.23 and median distances exceeding 1199 km. Notice, for Acc@1km, all baseline models despite

of pre-training and test split, failed almost completely in the SAGINDisaster dataset, and achieved very poor accuracy even in the local-scale MultiIAN dataset. This observation highlights the substantial gap when introducing generative geolocalization approaches to the critical disaster response scenario.

Table 3

Zero-shot performance of RFM-70 and RFM-80 models on the MultiIAN dataset across different pre-training datasets.

Pre-training	Method	Acc@1km	Acc@25km	Acc@50km	Acc@200km	MeanDist (km)	MedianDist (km)
iNaturalist	RFM-70	0.000	0.004	0.027	0.291	424.40	1213.17
	RFM-80	0.000	0.010	0.038	0.291	420.69	1199.87
YFCC	RFM-70	0.001	0.019	0.077	0.666	125.76	654.83
	RFM-80	0.000	0.026	0.084	0.709	117.98	623.79
OSV-5M	RFM-70	0.000	0.034	0.109	0.746	103.91	169.60
	RFM-80	0.000	0.032	0.094	0.741	99.93	174.55

Table 4

Fine-tuning results on the SAGINDisaster dataset with a 50 km threshold, showing the impact of refinement and retrieval. Herein, RFM refer to standalone generative model of Riemannian Flow Matching, Retrieval is the deterministic cross-view imagery retrieval baseline, and ProbGLC is the proposed approach. Best results for each pre-trained dataset are in bold.

Pre-trained	Method	Fine-tuned	Acc@1km	Acc@25km	Acc@50km	Acc@200km	Mean Dist (km)	Median Dist (km)
YFCC	RFM-70	✗	0.006	0.054	0.096	0.537	1078.54	187.37
	RFM-80	✗	0.014	0.055	0.101	0.558	941.71	165.56
	RFM-70	✓	0.000	0.646	0.681	0.784	273.97	7.68
	RFM-80	✓	0.000	0.625	0.666	0.796	267.58	7.71
	Retrieval@1-70	✓	0.446	0.678	0.692	0.753	222.59	0.94
	Retrieval@1-80	✓	0.459	0.668	0.685	0.750	154.05	0.86
	ProbGLC-70	✓	0.446	0.678	0.692	0.787	268.85	1.31
	ProbGLC-80	✓	0.459	0.668	0.685	0.796	261.62	1.32
OSM-5M	RFM-70	✗	0.000	0.042	0.111	0.582	643.31	158.68
	RFM-80	✗	0.000	0.043	0.113	0.594	513.13	105.02
	RFM-70	✓	0.000	0.638	0.710	0.859	142.32	12.24
	RFM-80	✓	0.000	0.570	0.680	0.844	165.31	14.40
	Retrieval@1-70	✓	0.468	0.716	0.736	0.801	72.91	0.86
	Retrieval@1-80	✓	0.469	0.683	0.707	0.762	70.74	0.77
	ProbGLC-70	✓	0.468	0.716	0.736	0.861	133.52	1.18
	ProbGLC-80	✓	0.469	0.683	0.707	0.849	156.25	1.27

Furthermore, two important findings herein deserve attention, thus affecting the upcoming evaluation design: 1) since the RFM consistently outperforms the other two variants, we decided to use it as the default training strategy in the ProbGLC. Similarly, we selected the YFCC and OSM-5M as pre-training baselines for further comparison given their performance superiority. 2) The key advantage of the generative approach lies in the fact that it directly predicts real locations of VGI imagery on the Earth rather than matching it with RSI to decide the location.

As the generative approach generates a probabilistic prediction of the imagery location, this brings a unique advantage for disaster geolocation in terms of model explainability and uncertainty quantification, besides the potential of performance boosting. In the next section, we will elaborate on the performance gain as well as the uncertainty quantification of the ProbGLC as a combined approach of both deterministic and generative geolocation methods.

4.4. ProbGLC Geolocation Results

Two-level of Performance Gain: The proposed ProbGLC approach aims at two levels of improvements over 1) baseline pre-training models and 2) generative or retrieval only fine-tuning models. Table 4 and 5 summarize our experiment results of evaluating ProbGLC approaches (applying a threshold r equal to 50 km), across two disaster datasets, respectively. As a protocol, the comparison between generative-only (i.e., RFM), retrieval-only, and the ProbGLC approaches highlights 1) the performance gain over baselines, 2) the advantage of uncertainty quantification and localizability measure via probabilistic predictions.

As one can see in Table 4, our ProbGLC led to substantial performance gains on the SAGINDisaster dataset across both YFCC and OSV-5M pre-trained baseline models. Compared to the pre-training baseline, where Acc@50km remained below 0.11 (e.g., 0.096 for YFCC with RFM-70 and

0.111 for OSM-5M with RFM-70), the ProbGLC boosted it to around 0.7 in Acc@50km. This first-level improvement is accompanied by a major reduction in Mean Dist, from over 1000 km to about 270 km for YFCC and 140 km for OSM-5M, and a sharp decrease in Median Dist from more than 150 km to around 7–12 km, which is favorable toward closing the "last kilometer" challenge in a disaster response scenario. Moreover, the ProbGLC variants as well as the deterministic Retrieval@1 notably enhanced short-range accuracy, especially for reaching up to 0.469 for Acc@1km and around 0.736 for Acc@25km, reducing the Median Dist to sub-kilometer levels (0.86 km for OSV-5M and 0.86–1.32 km for YFCC). However, we want to emphasize again that the key advantage of ProbGLC over deterministic Retrieval@1 is not about superior performance as they both yield very competitive results, but more about the explainability feature that comes with the generative probabilistic geolocation approach.

Among all settings, OSM-5M consistently achieved lower Mean Dist (down to 70.74 km) and higher Acc@25km to Acc@200km, suggesting stronger spatial generalization. In contrast, YFCC exhibited slightly larger residual Mean Dist despite comparable Acc@50km, indicating a slightly weaker localization accuracy. We attribute this phenomena to the fact that OSM-5M focuses on street-level features, which present higher similarity to the disaster-related VGI imagery in the latent space. Overall, we can see that ProbGLC is able to simultaneously benefit from the fine-tuning of generative approach and the deterministic refinement of retrieval-based approach, producing the best balance between accuracy and distance metrics across all scales.

Compared with the SAGINDisaster results (Table 4), the fine-tuning on the MultiIAN dataset led to dramatically higher localization accuracy and lower distance errors as shown in Table 5. For ProbGLC variants, Acc@50km rose

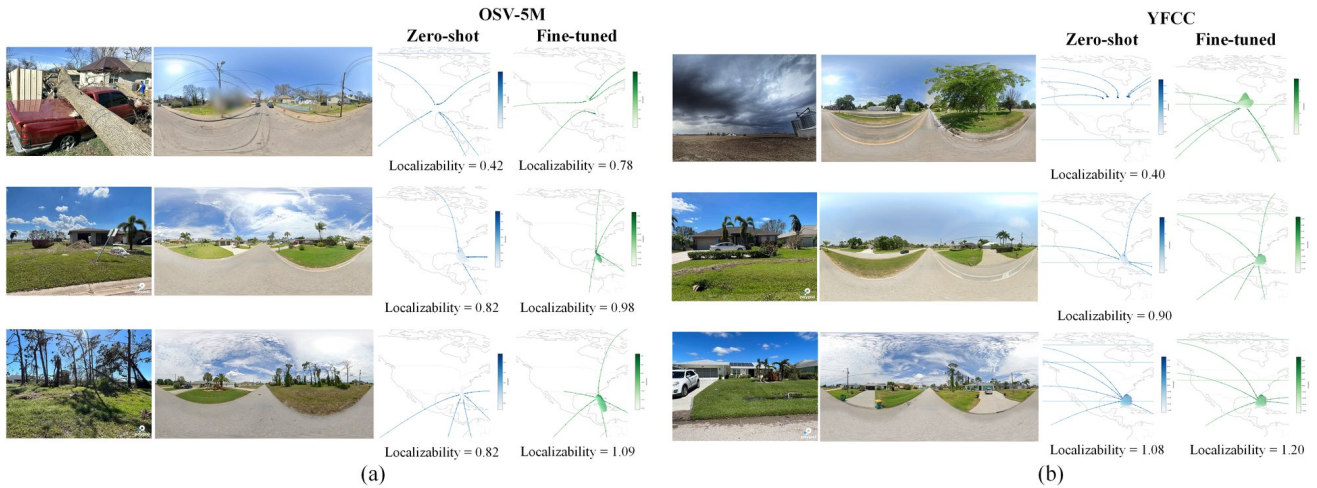


Figure 7: Examples of VGI imagery and corresponding SVI, and the estimation of localizability under different training strategies for the SAGINDisaster dataset. (a) Results based on OSV-5M pre-training; (b) Results based on YFCC pre-training.

Table 5

MultiiAN dataset: Fine-tuning with 50km threshold and impact of refinement and retrieval. Herein, RFM refer to standalone generative model of Riemannian Flow Matching, Retrieval is the deterministic cross-view imagery retrieval baseline, and ProbGLC is the proposed approach. Best results per pre-trained dataset are highlighted.

Pre-trained	Method	Fine-tuned	Acc@1km	Acc@25km	Acc@50km	Acc@200km	Mean Dist (km)	Median Dist (km)
YFCC	RFM-70	✗	0.001	0.019	0.077	0.666	125.76	654.83
	RFM-80	✗	0.000	0.026	0.084	0.709	117.98	623.79
	RFM-70	✓	0.114	0.949	0.963	0.987	9.88	2.50
	RFM-80	✓	0.090	0.952	0.968	0.994	8.01	2.92
	Retrieval@1-70	✓	0.838	0.962	0.970	0.981	4.25	0.00
	Retrieval@1-80	✓	0.868	0.970	0.977	0.994	3.85	0.00
	ProbGLC-70	✓	0.838	0.962	0.970	0.986	6.77	0.00
	ProbGLC-80	✓	0.868	0.970	0.977	0.996	4.00	0.00
OSV-5M	RFM-70	✗	0.000	0.034	0.109	0.746	103.91	169.60
	RFM-80	✗	0.000	0.032	0.094	0.741	99.93	174.55
	RFM-70	✓	0.085	0.957	0.973	0.992	9.09	3.12
	RFM-80	✓	0.100	0.963	0.975	0.994	7.48	2.86
	Retrieval@1-70	✓	0.848	0.970	0.974	0.987	3.55	0.00
	Retrieval@1-80	✓	0.870	0.972	0.978	0.991	3.09	0.00
	ProbGLC-70	✓	0.848	0.970	0.974	0.992	5.49	0.00
	ProbGLC-80	✓	0.870	0.972	0.978	0.994	3.85	0.00

from roughly 0.68–0.71 to above 0.96 for both YFCC and OSV-5M models, while the Mean Dist dropped sharply from over 270 km to under 10 km. The combined ProbGLC approach further boosted short-range accuracy, achieving up to 0.870 at 1 km and reducing the Median Dist to 0.00 km, indicating near-perfect geolocation. This is intuitive as we already mentioned that SAGINDisaster is considered in general more challenging than MultiAN. However, the implication on a local-scale disaster dataset is indeed exciting as it means by incorporating local samples, the ProbGLC approach can easily boost up a very competitive approach which is crucial when deploying to new disaster on a timely basis.

Probabilistic Distribution and Localizability: As aforementioned, the unique advantage of generative approach (with RFM) is about the predicted location comes always

with a probabilistic distribution as well as their fixed velocity field as illustrated in Figure 3, which enables us to visualize the conferencing trajectory of different VGI imagery directly in the Earth as well as calculating their localizability score as a uncertainty proxy.

Figure 7 and 8 offer a visual-intuitive comparison of generative RFM trajectories for baseline models and the fine-tuned ProbGLC approach alongside the corresponding localizability scores as a major uncertainty metric. First, we see a clear improvement of the ProGLC over generative baselines as the probability distributions are more spatially clustered as well as the increase localizability scores. In this context, by interpreting the trajectory patterns, one can now better explain where exactly the predicted locations and how confident the models are, which provides key insights into designing more focused and uncertainty-aware disaster

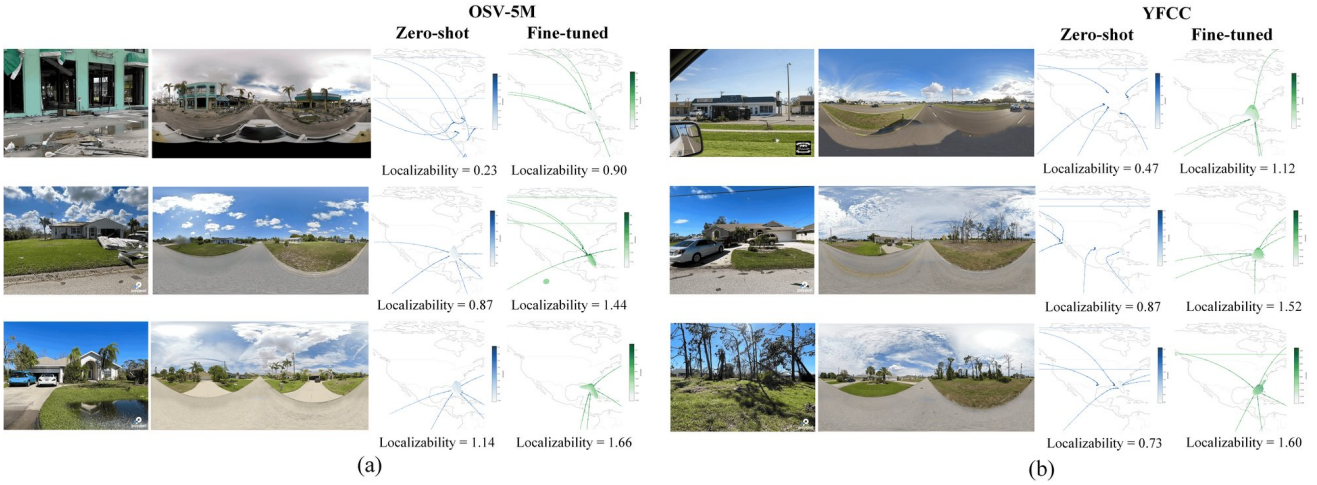


Figure 8: Examples of VGI imagery and corresponding SVI, and the estimation of localizability under different training strategies for the Hurricane IAN dataset. (a) Results based on OSV-5M pre-training; (b) Results based on YFCC pre-training.

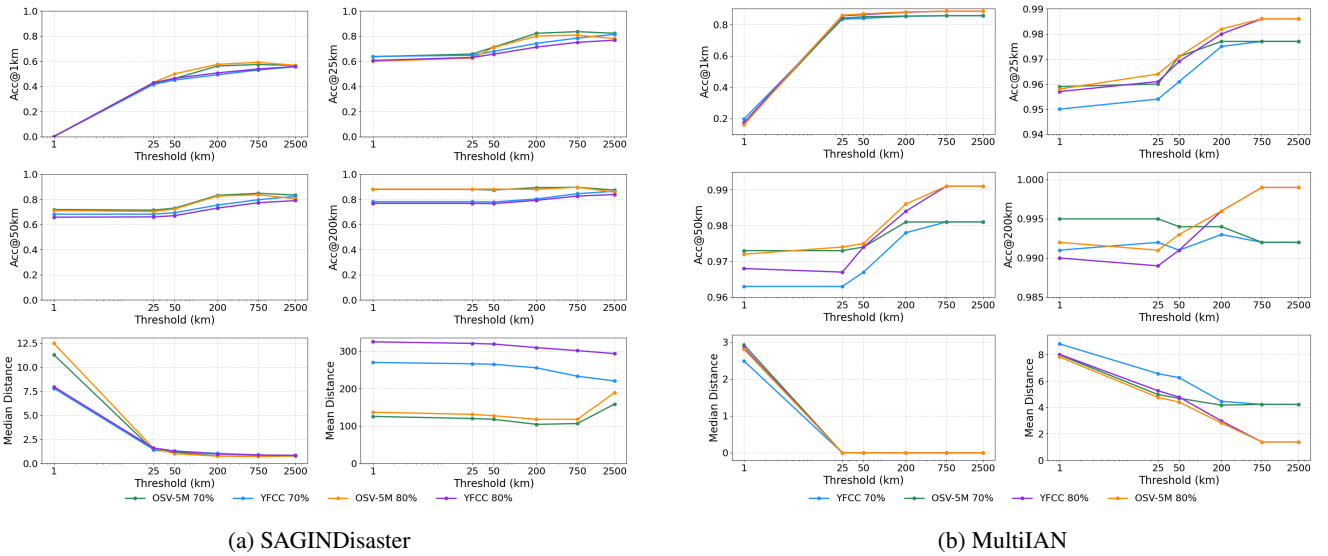


Figure 9: Effect of different decision thresholds on OSV-5M and YFCC across two datasets.

response actions. Second, we notice a common character of high localizability scores still refers to street layout (e.g., pavement and cross-road) as well as landscape features (e.g., palm tree and building styles), which is helpful to guide the further development in global-scale imagery geolocation pre-training and how to best deploy the ProbGLC approach in future disaster events. The capability of understanding how and why a model predicts a location is one of the most desired features and a persistent limitation for deterministic cross-view geolocation approaches.

Ablation Study: To provide a systematic justification, we examine the impact of different threshold r , which is considered as a trade-off between narrowing the probabilistic searching zone and maximizing deterministic retrieval accuracy. To remind its function, the threshold r determines the extent to which the Retrieval@1 models consider the predicted locations from the RFM generative models. Figure

9a and 9b elaborate on the impact of different distance thresholds on the SAGINDisaster and MultiIAN datasets, respectively. Two common patterns deserve notice here: 1) Median and Mean distances drop dramatically already when relaxing the threshold from 1km to 25km, aligning with the steep increase in fine-grained accuracy (Acc@1km and Acc@25km). This observation confirms the effectiveness of the combined ProbGLC approach to simultaneously benefit from the generative and deterministic geolocation approaches, ensuring both model explainability and accuracy performances. 2) One can notice that the change of threshold doesn't influence the performance much when the distance is set bigger than 50km, which actually means that the largest performance gain comes with the fact that the generative approach is able to narrow down the location to a relatively small area where the deterministic approach is able to hit the exact location via imagery retrieval. This is also the reason

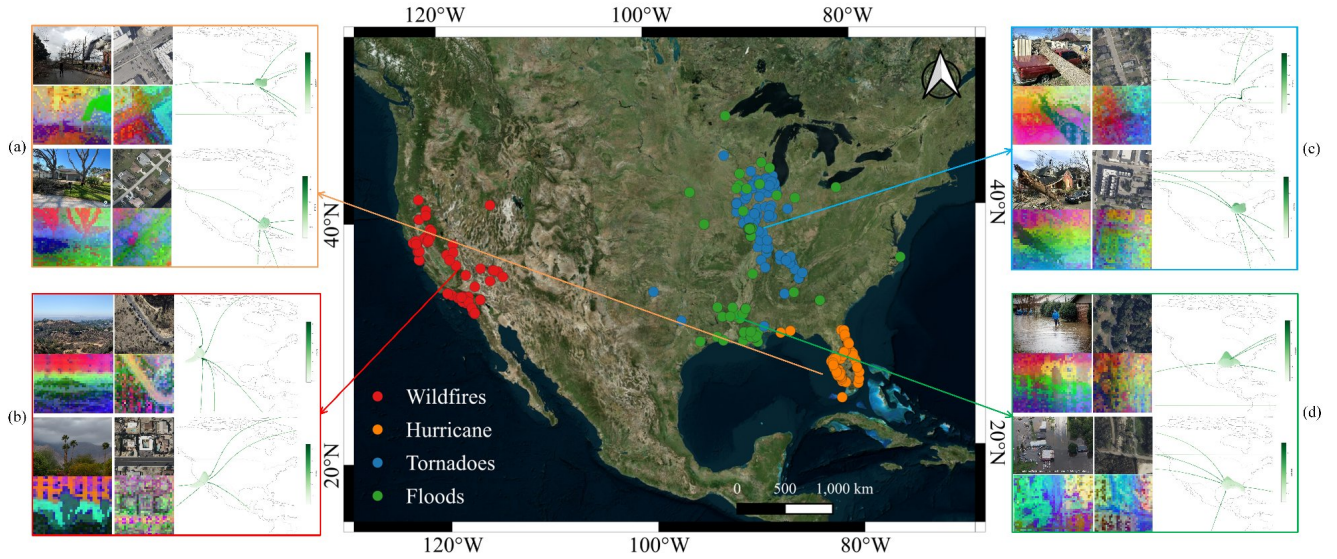


Figure 10: Visualization results of several representative examples from the SAGINDisaster dataset, corresponding to four major disaster types: (a) Hurricane, (b) Wildfire, (c) Tornado, and (d) Flood.

why we set the threshold r to 50km in Table 4 and 5. Of course, this ablation is not supposed to be comprehensive as one can still try out different settings in both RFM or Retrieval@1 approaches. However, as aforementioned, the threshold r is considered a key hyperparameter of the ProbGLC approach to balance accuracy and explainability from either deterministic or probabilistic approaches. If we take ProbGLC as a general framework for cross-view geolocation, one could easily extend this ablation to other important parameters.

4.5. Implications in Disaster Response

In this section, we would like to elaborate on the implications in real-world disaster response scenarios that we foresee for the proposed ProbGLC.

First, as shown in Figure 7 and 8, the localizability score together with the probabilistic density distribution is key to prioritize resource allocation. For example, decision-makers can use the localizability score to filter real-time VGI disaster imagery. Images with high localizability (low uncertainty) can be geolocated and immediately included into a crisis map for rapid damage assessment, while images with low localizability (high uncertainty) often indicate difficult terrains or severe destruction where visual cues are lost. These specific cases can be routed to human analysts for expert review, or directing UAVs to the estimated general area to capture clearer imagery where ground-level ambiguity is too high, preventing the dispatch of rescue teams to the wrong location.

Second, the ProbGLC allows for optimizing searching zones in real disaster response. For instance, in case of a lost person or a trapped group identified via real-time VGI imagery, the density map defines the geographical extent for the search operation. A tight cluster implies a focused search radius (e.g., Acc@1km), while a dispersed cloud suggests a

wider search grid is necessary. By combining with the localizability score, the whole system allows decision-makers to calibrate their confidence level in real-time and optimize their response actions with more informed decisions.

Moreover, thanks to the deterministic image encoder (i.e., ViT DINOv2), we are now able to visualize and compare the learned cross-view image patterns regarding different disaster types. Herein, Figure 10 presents a qualitative visual comparison using DINOv2 PCA techniques (AI, 2023) with selective examples from four major types of disaster in the SAGINDisaster dataset, respectively. Herein, DINOv2 PCA Visualization does the following trick, we have an image, process it through the Transformer and then have 32×32 patches. All have their features in the dimension size of the Transformer (i.e., 1024 in the case of ViT-L). Then we fit a PCA to 3 components meaning instead of $32 \times 32 \times 1024$ we get $32 \times 32 \times 3$. This is very intriguing because we can now interpret the 3 dimensions as colors, namely RGB. As we can see in Figure 10 when we compare PCA of VGI imagery the features tend to have the same color through the same PCA. For instance, the water is blue in (d), the tree and vegetation are purple (c), and the sky is orange (b). In the RSI image, there are more colors as more is visible compared to the VGI imagery. But this can be very helpful because it means the learned cross-view embeddings for the ProbGLC approach encode direct semantic relations, which is what is in the image and what is very similar. When combined with the probabilistic density map, this kind of visualization offers good insights into common and distinct features in different disaster-related geolocation tasks.

5. Discussions

In this work, we present a **Probabilistic Cross-view GeoLoCalization** approach, in short **ProbGLC**, as a novel

pathway towards generative location awareness for rapid disaster response. Despite the motivating preliminary results we got, there are a few points that deserve additional discussion and awareness.

Generative cross-view approaches: A major contribution of the proposed ProbGLC approach is to introduce the generative training approach (via diffusion and flow matching) into imagery geolocalization tasks. However, the pretraining of these generative backbones (e.g., Diffusion, FlowR3, RFM) is limited largely by two practical factors: 1) the geographical distribution, density, and focus of three global visual geolocalization datasets, namely YFCC, OSV-5M, and iNaturalist; 2) the sole training target of locations (i.e., lat and long). The first factor implies the underlying biases in the crowdsourced imagery datasets, where the density and spatial coverage of images, no matter street view or social media, vary significantly across countries and regions. Therefore, the geographical generalizability (Li et al., 2023), similarly the replicability across space and time (Goodchild and Li, 2021), deserves more dedicated and comprehensive evaluations and treatments. The second factor can be attributed to the fact that there is still a research gap in a global-scale geospatial pretraining scheme that can provide fine-grained and general-purpose backbones. Recent works, including but not limited to SatCLIP (Klemmer et al., 2025), Clay Foundation Model (Clay, 2024), and the most recent Google AlphaEarth Embedding (Brown et al., 2025), shed light on this direction, while their performance in time-critical disaster response scenarios still needs careful investigations and the potential of generative cross-view training remains largely unexplored.

GeoAI for disaster resilience: In this work, we focus on ensuring fast and accurate cross-view geolocalization in a post-disaster scenario. Preliminary results from diverse disaster types (i.e., hurricane, wildfire, tornado, and flood) highlighted the increasing role of GeoAI to predict, analyze, and enhance urban disaster resilience against climate extreme events. A comprehensive assessment of disaster resilience involves not only physical infrastructure such as drainage systems and flood barriers, but also socio-economic factors like population density, vulnerability, and impact assessment.

Towards multi-temporal monitoring and fine-grained disaster mapping, we wish the lessons learned from this work can be integrated into dedicated local GeoAI initiatives, especially for those global south areas as argued in (Huang, 2025), to better prepare our cities for future disaster and climate extreme events. Recent works by Xue and Zhu (2025); Wang et al. (2026) highlight the potential of incorporating Vision-Language Models (VLMs) in understanding fine-grained and regression-like geospatial predictions. It would be interesting to see how one can leverage VLMs to understand the multi-faceted disaster resilience aspects (before, during, and after) in combination with socio-economic context information.

Representation learning for multimodal geospatial data: A fundamental task for GeoAI models is to align the

multimodal geospatial data in the latent space so that their latent representations can be better used for downstream tasks (Mai et al., 2024; Liu et al., 2025), no matter whether it is regression, classification, or geolocalization tasks. However, so far the major advances in the field focus on raster-based imagery data (such as satellite and street view imagery) due to the uniform format of pixels and rich content from the visual perception. However, in the geospatial domain, there is still a large amount of vector data (e.g., points, polylines, polygons) that remains largely underexplored for representation learning, for instance, in OpenStreetMap, which also contains rich and helpful location and semantic information.

To extract meaningful features from vector data for predictive tasks, existing researchers either perform feature engineering based on domain knowledge (Yan et al., 2022), or convert spatial data from their original formats (e.g., points, polylines, and polygons) into formats that are easier for neural networks to handle, or map vector files into raster image tiles (Kang et al., 2019; Li and Zipf, 2022). Herein, the first stream heavily relies on domain knowledge and can not be easily generalized to new tasks, while the latter stream suffers from reduced data precision and loss of linked semantic context. Either way, the representation capability and performance are limited due to the lack of end-to-end learning. Therefore, a promising direction for future work is to incorporate massive vector datasets (e.g., from OpenStreetMap) into the current learning framework with a dedicated vector data encoder, for instance, Poly2Vec (Siampou et al., 2024) and Geo2Vec (Chu and Shahabi, 2025), facilitating a synergy of multimodal geospatial data into more efficient and effective representation learning.

6. Conclusions

In this paper, we present a novel cross-view disaster geolocalization approach, called ProbGLC, standing for **Probabilistic Cross-view GeoLoCalization**, to facilitate generative location awareness for rapid disaster response purpose. The ProbGLC is the first of this kind of model to combine probabilistic and deterministic geolocalization models into a unified framework, which brings unique benefits in terms of model explainability (via uncertainty quantification) and superior geolocalization accuracy compared to state-of-the-art baseline approaches. More importantly, by leveraging the generative training strategy, for instance, Riemannian Flow Matching on the Sphere, the ProbGLC predicts imagery locations directly on the Sphere with a probabilistic distribution as a proxy of uncertainty and localizability, which enables us to better understand how the model makes decisions during the geolocalization process. Extensive experiments on two cross-view disaster geolocalization datasets (SAGINDisaster and MultiIAN) show competitive performance and explainability features of the proposed ProbGLC model in handling multiple disaster types with distinct disaster characteristics. In addition, we carefully elaborate on the implications of ProbGLC in various disaster response scenarios and the benefits it brings

towards leveraging GeoAI in disaster response. Building on these insights, our future work will focus on developing responsible GeoAI approaches that enhance disaster risk management and societal resilience in the changing climate.

Acknowledgments

We would like to thank the editor and reviewers for their valuable comments, which have helped improve the quality of our manuscript. This work was supported by the Start-Up Grant (SUG) project ‘‘Geospatial Artificial Intelligence for Climate Resilient Urban Environment’’ from the National University of Singapore (E-109-00-0036-01).

Disclosure Statement

No potential conflict of interest was reported by the authors.

A. Geographic Diffusion in the Euclidean space

In this appendix, we elaborate on the basic of geographic diffusion approach in the Euclidean space. Since this work is among the early attempts in the geospatial community to address imagery geolocalization from a generative perspective, we think it would be beneficial to provide an introduction to the basics of diffusion models. The concept of diffusion models was first introduced by Sohl-Dickstein et al. (2015), which only gained popularity after Ho et al. (2020) formally outlined their use as generative models.

Forward Diffusion Process: Given an input variable, in our case a random location x_n (longitude and latitude), the core idea of diffusion is to iteratively add Gaussian noise to this location variable. We denote the location at the 0th timestamp (i.e., the original location where no noise has been added yet) as \mathbf{x}_0 . With each iteration, we increment the timestamp such that $\mathbf{x}_t = \mathbf{x}_{t-1} + \epsilon$, where $\epsilon \in \mathbb{R}^D$ refer to the Gaussian noise. After a sufficient number of iterations, we posit that $\mathbf{x}_T \sim \mathcal{N}(0, \mathbf{I})$. In other words, if we add enough noise, the resulting location variable will be indistinguishable from a random one drawn from a unit normal distribution. This process is called the forward diffusion process, which can be technically formulated as follows:

$$\mathbf{x}_t = \sqrt{1 - \beta(t)}\mathbf{x}_{t-1} + \sqrt{\beta(t)}\epsilon, \quad (13)$$

Where $\mathbf{x}_1, \dots, \mathbf{x}_T$ are the noisy latent variables at different time stamps $t \in [0, 1]$ and $\beta(t) : [0, 1] \rightarrow [0, 1]$ as a scheduling function with $\beta(0) = 0$ and $\beta(1) = 1$ to control the noise level added to the coordinates.

Herein, these latent variables are not simply fixed versions of the noisy x_0 , but a condition variable. We assume that they follow a distribution conditioned on the previous step, denoted as $\mathcal{N}(0, \beta(t)\mathbf{I})$ that represents a multivariate Gaussian distribution with covariance of $\beta(t)\mathbf{I}$. The only exception will be the original location, denoted as $\mathbf{x}_0 \sim$

$q(x_0)$, which is not conditioned on anything else. Therefore, the sequence of diffused locations \mathbf{x}_1 to \mathbf{x}_T forms a Markov chain and can be formulated as:

$$q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \mathcal{N}(\sqrt{1 - \beta(t)}\mathbf{x}_{t-1}, \beta(t)\mathbf{I}), \quad (14)$$

Which further leads to the following equation to present the whole forward diffusion process:

$$q(\mathbf{x}_{1:T} | \mathbf{x}_0) = \prod_{i=1}^T q(\mathbf{x}_i | \mathbf{x}_{i-1}), \quad (15)$$

Backward Denoising Process: Based on the forward diffusion process, the key idea of Denoising Diffusion Probabilistic Models (DDPM) introduced by Ho et al. (2020) is to learn a neural network parametrized by φ , which takes the diffused location \mathbf{x}_t and the noisy level $\beta(t)$ as input and most importantly is conditioned on the imagery embeddings \mathbf{c} .

Considering the final timestep T of added noise will result in an entirely random location. We can represent this as $P(\mathbf{x}_T) = \mathcal{N}(0, \mathbf{I})$, where $P_\varphi(\mathbf{x}_{t-1} | \mathbf{x}_t)$ denotes the learned distribution of the de-noised \mathbf{x}_t . This represents the previous step in the noise addition process, \mathbf{x}_{t-1} . Just opposite to the forward diffusion process, the backward denoising will yield:

$$P_\varphi(\mathbf{x}_{t-1} | \mathbf{x}_t) = \mathcal{N}(\mu_\varphi(\mathbf{x}_t, t), \Sigma_\varphi(\mathbf{x}_t, t)), \quad (16)$$

Where μ_φ and Σ_φ refer to the mean and covariance noisy vector ϵ for the previous state \mathbf{x}_{t-1} , both are functions of \mathbf{x}_t and the timestamp t . Herein, φ refers to a set of trainable parameters in the neural network (Figure 2), with which we can learn a function to de-noise the latent variables.

The overall denoising process can then be formulated similarly as a Markov chain as follows:

$$P_\varphi(\mathbf{x}_{0:T}) = P(\mathbf{x}_T) \prod_{i=1}^T P_\varphi(\mathbf{x}_{i-1} | \mathbf{x}_i), \quad (17)$$

The key is to find a solution to optimize φ to achieve a one-to-one denoising of the locations. As introduced by Ho et al. (2020), it is suggested to fix the variance to 1, so the only remaining learnable parameter is $\mu_\varphi(\mathbf{x}_t, t)$. Further, the μ_φ predictor is replaced with the ϵ , where the predictor is trained to predict the reverse noise, in order to denoise the latent variable. As a result, the whole generative model can be trained to minimize the diffusion loss function as follows:

$$\mathcal{L}_{DF} = \mathbb{E}_{\mathbf{x}_0, \mathbf{c}, \epsilon, t} [\|\varphi(\mathbf{x}_t | \mathbf{c}) - \epsilon\|^2], \quad (18)$$

Where the expectation is over four variables, saying the origin location \mathbf{x}_0 , the imagery embedding \mathbf{c} , the Gaussian noise ϵ , and the timestamp t . Herein, $\varphi(\mathbf{x}_t | \mathbf{c})$ refers to the generative DDPM network based on a noisy location \mathbf{x}_t and conditional on the imagery embedding \mathbf{c} .

Inference Process: For inference, given a new imagery c , the geographic diffusion start with a random location with $\mathbf{x}_1 = \epsilon$, then iteratively denoise the location latent variable \mathbf{x}_t over a specific timesteps from $t = 1$ to $t = 0$, using the location variable update equation as follows:

$$\mathbf{x}_{t-\Delta t} = \sqrt{1 - \beta(t)} \hat{\mathbf{x}}_t + \sqrt{\beta(t)} \varphi(\mathbf{x}_t | c) \quad (19)$$

$$\hat{\mathbf{x}}_t = \frac{1}{\sqrt{1 - \beta(t)}} \mathbf{x}_t - \sqrt{\beta(t)} \varphi(\mathbf{x}_t | c) \quad (20)$$

Where Δt and $\hat{\mathbf{x}}_t$ refer to the time step size and the denoised location variables, respectively. Finally, the $\hat{\mathbf{x}}_0$ in Euclidean space \mathbb{R}^3 will be projected to the Spherical coordinates of \mathbb{S}^2 (i.e., longitude and latitude).

B. The Projection between Spherical and Euclidean Spaces

Herein, the logarithmic map \log_x maps a point $y \in \mathcal{S}_2$ onto T_x , the tangent space at point x (Sommer et al., 2020):

$$\log_x(y) = \frac{\theta}{\sin \theta} (y - \cos \theta x), \quad (21)$$

where $\theta = \arccos(\langle x, y \rangle)$ is the angle between x and y .

Next, the exponential map \exp_x of a point $x \in \mathcal{S}_2$ maps a tangent vector $v \in T_x$ back onto the spherical space:

$$\exp_x(v) = \cos(\|v\|)x + \frac{\sin(\|v\|)}{\|v\|} v, \quad (22)$$

where $\|v\|$ is the Euclidean norm of v .

Moreover, the projection here is fully aligned with classic map projection theory, therefore it can be potentially extend to different ellipsoid coordinates as well.

C. Anchor-based Reranking as Postprocessing

In a nutshell, the anchor-based reranking (as shown in Figure 4) helps to refine the initial retrieval ranking by selecting representative anchors (i.e., often the query and top results) and re-scoring candidates based on both their similarity to the query and their consistency with these anchors.

Given the learned cross-view embedding space \mathbb{R}_{CV} and a query embedding of VGI imagery $L_c \in \mathbb{R}^d$ and candidate RSI embeddings $\{L_s^i\}_{i=1}^N$, the initial retrieval computes a base cosine similarity as follows:

$$S(L_c, L_i) = \frac{L_c \cdot L_s^i}{\|L_c\| \|L_s^i\|}. \quad (23)$$

By selecting the top- k retrieved items as anchors $\{a_j\}_{j=1}^k$, where $a_j \in \{L_s^i\}$, we can rescore each RSI candidate by

combining query-candidate similarity and anchor-candidate consistency using the following equation:

$$R(L_s^i) = \alpha S(q, x_i) + \frac{1 - \alpha}{k} \sum_{j=1}^k S(a_j, x_i), \quad (24)$$

Where $\alpha \in [0, 1]$ is a ratio to balance the query and anchor influence, x_i and q refers to the query VGI imagery and the candidate RSI embeddings, respectively.

References

- Abdar, M., Pourpanah, F., Hussain, S., Rezazadegan, D., Liu, L., Ghavamzadeh, M., Fieguth, P., Cao, X., Khosravi, A., Acharya, U.R., et al., 2021. A review of uncertainty quantification in deep learning: Techniques, applications and challenges. *Information fusion* 76, 243–297.
- AI, M., 2023. Dinov2: Learning robust visual features without supervision. <https://github.com/facebookresearch/dinov2>. Accessed: 2025-10-27.
- Astruc, G., Dufour, N., Siglidis, I., Aronsson, C., Bouia, N., Fu, S., Loiseau, R., Nguyen, V.N., Raude, C., Vincent, E., et al., 2024. Openstreetview-5m: The many roads to global visual geolocation, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 21967–21977.
- Biljecki, F., Ito, K., 2021. Street view imagery in urban analytics and GIS: A review. *Landscape and Urban Planning* 215, 104217.
- Blei, D.M., Kucukelbir, A., McAuliffe, J.D., 2017. Variational inference: A review for statisticians. *Journal of the American statistical Association* 112, 859–877.
- Bromley, J., Guyon, I., LeCun, Y., Säckinger, E., Shah, R., 1993. Signature verification using a "siamese" time delay neural network. *Advances in neural information processing systems* 6, 737–744.
- Brown, C.F., Kazmierski, M.R., Pasquarella, V.J., Rucklidge, W.J., Samsikova, M., Zhang, C., Shelhamer, E., Lahera, E., Wiles, O., Ilyushchenko, S., et al., 2025. Alphaearth foundations: An embedding field model for accurate and efficient global mapping from sparse label data. *arXiv preprint arXiv:2507.22291*.
- Cai, H., Lam, N.S., Qiang, Y., Zou, L., Correll, R.M., Mihinov, V., 2018. A synthesis of disaster resilience measurement methods and indices. *International journal of disaster risk reduction* 31, 844–855.
- Chen, R.T., Lipman, Y., 2024. Flow matching on general geometries, in: *The Twelfth International Conference on Learning Representations*.
- Chu, C., Shahabi, C., 2025. Geo2vec: Shape-and distance-aware neural representation of geospatial entities. *arXiv preprint arXiv:2508.19305*.
- Clay, 2024. Clay foundation model. <https://huggingface.co/made-with-clay>. Accessed: 2025-10-28.
- Dehimi, N.E.H., Tolba, Z., 2024. Attention mechanisms in deep learning: Towards explainable artificial intelligence, in: *2024 6th International Conference on Pattern Analysis and Intelligent Systems (PAIS)*, IEEE, pp. 1–7.
- Delforge, D., Wathelet, V., Below, R., Sofia, C.L., Tonnelier, M., van Loenhout, J.A., Speybroeck, N., 2025. Em-dat: the emergency events database. *International Journal of Disaster Risk Reduction*, 105509.
- Deuser, F., Habel, K., Oswald, N., 2023. Sample4geo: Hard negative sampling for cross-view geo-localisation, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 16847–16856.
- Deuser, F., Werner, M., Habel, K., Oswald, N., 2024. Optimizing geo-localization with k-means re-ranking in challenging weather conditions, in: *Proceedings of the 2nd Workshop on UAVs in Multimedia: Capturing the World from a New Perspective*, pp. 9–13.
- Dietze, M., Ozturk, U., 2021. A flood of disaster response challenges. *Science* 373, 1317–1318.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al., 2021. An image is worth 16x16 words: Transformers for image

- recognition at scale, in: International Conference on Learning Representations, pp. 1–21.
- Dufour, N., Picard, D., Kalogeiton, V., Landrieu, L., 2024. Around the world in 80 timesteps: A generative approach to global visual geolocation. *arXiv preprint arXiv:2412.06781*.
- Ellis, P.W., Page, A.M., Wood, S., Fargione, J., Masuda, Y.J., Carrasco Denney, V., Moore, C., Kroeger, T., Griscom, B., Sanderman, J., et al., 2024. The principles of natural climate solutions. *Nature Communications* 15, 547.
- Feng, Y., Huang, X., Sester, M., 2022. Extraction and analysis of natural disaster-related vgi from social media: review, opportunities and challenges. *International Journal of Geographical Information Science* 36, 1275–1316. URL: <http://dx.doi.org/10.1080/13658816.2022.2048835>, doi:10.1080/13658816.2022.2048835.
- Foroutan, E., Hu, T., Li, Z., 2025. Revealing key factors of heat-related illnesses using geospatial explainable ai model: A case study in texas, usa. *Sustainable Cities and Society* 122, 106243.
- Goodchild, M.F., Li, W., 2021. Replication across space and time must be weak in the social and environmental sciences. *Proceedings of the National Academy of Sciences* 118, e2015759118.
- Gupta, R., Goodman, B., Patel, N., Hosfelt, R., Sajeev, S., Heim, E., Doshi, J., Lucas, K., Choset, H., Gaston, M., 2019. Creating xbd: A dataset for assessing building damage from satellite imagery, in: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops, pp. 10–17.
- Hallegatte, S., Vogt-Schilb, A., Bangalore, M., Rozenberg, J., 2016. *Unbreakable: building the resilience of the poor in the face of natural disasters*. World Bank Publications.
- Ho, J., Jain, A., Abbeel, P., 2020. Denoising diffusion probabilistic models. *CoRR abs/2006.11239*. URL: <https://arxiv.org/abs/2006.11239>, arXiv:2006.11239.
- Höhl, A., Obadic, I., Fernández-Torres, M.Á., Oliveira, D., Zhu, X.X., 2024. Recent trends challenges and limitations of explainable ai in remote sensing, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 8199–8205.
- Hsu, C.Y., Li, W., 2023. Explainable geoai: can saliency maps help interpret artificial intelligence’s learning process? an empirical study on natural feature detection. *International Journal of Geographical Information Science* 37, 963–987.
- Hu, C., Jia, S., Zhang, F., Xiao, C., Ruan, M., Thrasher, J., Li, X., 2023a. Updexplainer: An interpretable transformer-based framework for urban physical disorder detection using street view imagery. *ISPRS Journal of Photogrammetry and Remote Sensing* 204, 209–222.
- Hu, X., Zhou, Z., Li, H., Hu, Y., Gu, F., Kersten, J., Fan, H., Klan, F., 2023b. Location reference recognition from texts: A survey and comparison. *ACM Comput. Surv.* 56. URL: <https://doi.org/10.1145/3625819>, doi:10.1145/3625819.
- Hu, Y., Goodchild, M., Zhu, A.X., Yuan, M., Aydin, O., Bhaduri, B., Gao, S., Li, W., Lunga, D., Newsam, S., 2024. A five-year milestone: reflections on advances and limitations in geoai research. *Annals of GIS* 30, 1–14.
- Huang, X., 2025. Geospatial artificial intelligence (geoai) is widening the digital divide. *Annals of the American Association of Geographers* 115, 2017–2029.
- Huang, X., Wang, C., Li, Z., 2018. A near real-time flood-mapping approach by integrating social media and post-event satellite imagery. *Annals of GIS* 24, 113–123.
- Hulme, M., Barrow, E.M., Arnell, N.W., Harrison, P.A., Johns, T.C., Downing, T.E., 1999. Relative impacts of human-induced climate change and natural climate variability. *Nature* 397, 688–691.
- Jones, H.P., Hole, D.G., Zavaleta, E.S., 2012. Harnessing nature to help people adapt to climate change. *Nature climate change* 2, 504–509.
- Kang, Y., Gao, S., Roth, R.E., 2019. Transferring multiscale map styles using generative adversarial networks. *International Journal of Cartography* 5, 115–141.
- Klemmer, K., Rolf, E., Robinson, C., Mackey, L., Rußwurm, M., 2025. Satclip: Global, general-purpose location embeddings with satellite imagery, in: Proceedings of the AAAI Conference on Artificial Intelligence, pp. 4347–4355.
- Kupinski, M.A., Hoppin, J.W., Clarkson, E., Barrett, H.H., 2003. Ideal-observer computation in medical imaging with use of markov-chain monte carlo techniques. *Journal of the Optical Society of America A* 20, 430–438.
- Lam, N.S., Meyer, M., Reams, M., Yang, S., Lee, K., Zou, L., Mihunov, V., Wang, K., Kirby, R., Cai, H., 2023. Improving social media use for disaster resilience: challenges and strategies. *International Journal of Digital Earth* 16, 3023–3044.
- Li, H., Deuser, F., Yin, W., Luo, X., Walther, P., Mai, G., Huang, W., Werner, M., 2025. Cross-view geolocalization and disaster mapping with street-view and vhr satellite imagery: A case study of hurricane ian. *ISPRS Journal of Photogrammetry and Remote Sensing* 220, 841–854.
- Li, H., Herfort, B., Huang, W., Zia, M., Zipf, A., 2020. Exploration of openstreetmap missing built-up areas using twitter hierarchical clustering and deep learning in mozambique. *ISPRS Journal of Photogrammetry and Remote Sensing* 166, 41–51. URL: <https://doi.org/10.1016/j.isprsjprs.2020.05.007>.
- Li, H., Wang, J., Zollner, J.M., Mai, G., Lao, N., Werner, M., 2023. Rethink geographical generalizability with unsupervised self-attention model ensemble: A case study of openstreetmap missing building detection in africa, in: Proceedings of the 31st ACM International Conference on Advances in Geographic Information Systems, pp. 1–9.
- Li, H., Zipf, A., 2022. A conceptual model for converting openstreetmap contribution to geospatial machine learning training data. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* 43, 253–259.
- Li, W., Arundel, S., Gao, S., Goodchild, M., Hu, Y., Wang, S., Zipf, A., 2024. Geoai for science and the science of geoai. *Journal of spatial information science*, 1–17.
- Li, Z., 2025. Explainable ai in spatial analysis. *arXiv preprint arXiv:2505.00591*.
- Lin, T.Y., Cui, Y., Belongie, S., Hays, J., 2015. Learning deep representations for ground-to-aerial geolocalization, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 5007–5015.
- Lipman, Y., Chen, R.T., Ben-Hamu, H., Nickel, M., Le, M., 2022. Flow matching for generative modeling, in: The Eleventh International Conference on Learning Representations.
- Liu, P., Yan, Z., Biljecki, F., 2024. Explainable spatially explicit geospatial artificial intelligence in urban analytics. *Environment and Planning B: Urban Analytics and City Science* 51, 1104–1123. doi:10.1177/23998083231204689.
- Liu, Y., Wang, X., Wang, Y., Huang, F., Huang, Y., Li, Y., Zhang, W., Gong, S., Mai, G., Yao, Y., et al., 2025. Representation learning for geospatial data. *Annals of GIS*, 1–27.
- Liu, Z., Mao, H., Wu, C.Y., Feichtenhofer, C., Darrell, T., Xie, S., 2022. A convnet for the 2020s, in: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 11976–11986.
- Lundberg, S.M., Lee, S.I., 2017. A unified approach to interpreting model predictions. *Advances in neural information processing systems* 30.
- Macatulad, E., Biljecki, F., 2024. Continuing from the sendai framework midterm: Opportunities for urban digital twins in disaster risk management. *International Journal of Disaster Risk Reduction* 102, 104310. doi:10.1016/j.ijdrr.2024.104310.
- Mai, G., Xie, Y., Jia, X., Lao, N., Rao, J., Zhu, Q., Liu, Z., Chiang, Y.Y., Jiao, J., 2025. Towards the next generation of geospatial artificial intelligence. *International Journal of Applied Earth Observation and Geoinformation* 136, 104368.
- Mai, G., Yao, X., Xie, Y., Rao, J., Li, H., Zhu, Q., Li, Z., Lao, N., 2024. Srl: Towards a general-purpose framework for spatial representation learning, in: Proceedings of the 32nd ACM International Conference on Advances in Geographic Information Systems, pp. 465–468.
- Manzini, T., Murphy, R.R., Heim, E., Robinson, C., Zarrella, G., Gupta, R., 2023. Harnessing ai and robotics in humanitarian assistance and disaster response. *Science robotics* 8, eadj2767.
- Metropolis, N., Ulam, S., 1949. The monte carlo method. *Journal of the American statistical association* 44, 335–341.

- Nohrstedt, D., Hileman, J., Mazzoleni, M., Di Baldassarre, G., Parker, C.F., 2022. Exploring disaster impacts on adaptation actions in 549 cities worldwide. *Nature communications* 13, 3360.
- Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al., 2024. Dinov2: Learning robust visual features without supervision. *Transactions on Machine Learning Research* 1, 1–32.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al., 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* 32.
- Ribeiro, M.T., Singh, S., Guestrin, C., 2016. "why should i trust you?" explaining the predictions of any classifier, in: *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 1135–1144.
- Russo, E., Domeisen, D.I., 2023. Increasing intensity of extreme heatwaves: The crucial role of metrics. *Geophysical Research Letters* 50, e2023GL103540.
- Shives, E., Chen, T.H.K., Seto, K.C., 2025. Multiple hazards and exposure in california: A space-time analysis of temperature, drought, and wildfire. *International Journal of Disaster Risk Reduction* 120, 105391.
- Siampou, M.D., Li, J., Krumm, J., Shahabi, C., Lu, H., 2024. Poly2vec: Polymorphic fourier-based encoding of geospatial objects for geoai applications. *arXiv preprint arXiv:2408.14806*.
- Sohl-Dickstein, J., Weiss, E.A., Maheswaranathan, N., Ganguli, S., 2015. Deep unsupervised learning using nonequilibrium thermodynamics. *CoRR abs/1503.03585*. URL: <http://arxiv.org/abs/1503.03585>, *arXiv:1503.03585*.
- Sommer, S., Fletcher, T., Pennec, X., 2020. Introduction to differential and riemannian geometry, in: *Riemannian geometric statistics in medical image analysis*. Elsevier, pp. 3–37.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R., 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* 15, 1929–1958.
- Thomee, B., Shamma, D.A., Friedland, G., Elizalde, B., Ni, K., Poland, D., Borth, D., Li, L.J., 2016. Yfcc100m: The new data in multimedia research. *Communications of the ACM* 59, 64–73.
- Van Horn, G., Cole, E., Beery, S., Wilber, K., Belongie, S., Mac Aodha, O., 2021. Benchmarking representation learning for natural world image collections, in: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 12884–12893.
- Vivanco Cepeda, V., Nayak, G.K., Shah, M., 2023. Geoclip: Clip-inspired alignment between locations and images for effective worldwide geolocalization. *Advances in Neural Information Processing Systems* 36, 8690–8701.
- Wang, J., Xuan, W., Qi, H., Chen, Z., Chen, H., Zheng, Z., Xia, J., Zhong, Y., Yokoya, N., 2026. Cityvlm: Towards sustainable urban development via multi-view coordinated vision–language model. *ISPRS Journal of Photogrammetry and Remote Sensing* 232, 62–74.
- Wang, J., Xuan, W., Qi, H., Liu, Z., Liu, K., Wu, Y., Chen, H., Song, J., Xia, J., Zheng, Z., et al., 2025. Disasterm3: A remote sensing vision-language dataset for disaster damage assessment and response. *arXiv preprint arXiv:2505.21089*.
- Wikipedia, 2022. Hurricane ian. https://en.wikipedia.org/wiki/Hurricane_Ian.
- Workman, S., Souvenir, R., Jacobs, N., 2015. Wide-area image geolocalization with aerial reference imagery, in: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 3961–3969.
- Xue, X., Zhu, X.X., 2025. Regression in eo: Are vlms up to the challenge? *arXiv preprint arXiv:2502.14088*.
- Yan, X., et al., 2022. A graph deep learning approach for urban building grouping. *Geocarto International* 37, 2944–2966.
- Yin, W., Deuser, F., Jiang, J., Liu, Z., Luo, X., Zheng, Z., Werner, M., Li, H., Xue, Y., 2025a. Sagingeo: A space-aerial-ground integrated framework for multi-disaster cross-view geolocalization. Under review in *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*.
- Yin, W., Deuser, F., Liu, Z., Wei, J., Luo, X., Werner, M., Li, H., Xue, Y., 2025b. Triple-objective cross-view geolocalization of disaster-related vgi: the case of hurricane ian. *International Journal of Geographical Information Science*, 1–23.
- Yin, W., Xue, Y., Liu, Z., Li, H., Werner, M., 2025c. Llm-enhanced disaster geolocalization using implicit geoinformation from multimodal data: A case study of hurricane harvey. *International Journal of Applied Earth Observation and Geoinformation* 137, 104423.
- You, Y., Li, J., Reddi, S., Hseu, J., Kumar, S., Bhojanapalli, S., Song, X., Demmel, J., Keutzer, K., Hsieh, C.J., 2019. Large batch optimization for deep learning: Training bert in 76 minutes. *arXiv preprint arXiv:1904.00962*.
- Zeng, Q., Wu, J., Ren, Y., Feng, G., 2025. Cross-view geolocation via segmentation and common region feature matching. *ISPRS Journal of Photogrammetry and Remote Sensing*.
- Zhang, F., Salazar-Miranda, A., Duarte, F., Vale, L., Hack, G., Chen, M., Liu, Y., Batty, M., Ratti, C., 2024. Urban visual intelligence: Studying cities with artificial intelligence and street-level imagery. *Annals of the American Association of Geographers* 114, 876–897.
- Zhang, J., Goodchild, M.F., 2002. Uncertainty in geographical information. CRC press.
- Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A., 2016. Learning deep features for discriminative localization, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2921–2929.
- Zhou, B., Zou, L., Yang, M., Lin, B., Mandal, D., Abedin, J., Cai, H., Ji, S., Klein, A., Tian, H., 2025. Rapid disaster response and damage estimation with social media and pretrained large language models: Insights from multiple hurricanes. *Annals of the American Association of Geographers*, 1–23.
- Zhu, S., Shah, M., Chen, C., 2022. Transgeo: Transformer is all you need for cross-view image geo-localization, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1162–1171.
- Zhu, Y., Yang, H., Lu, Y., Huang, Q., 2023. Simple, effective and general: A new backbone for cross-view image geo-localization. *arXiv preprint arXiv:2302.01572*.
- Zou, L., Lam, N.S., Cai, H., Qiang, Y., 2018. Mining twitter data for improved understanding of disaster resilience. *Annals of the American Association of Geographers* 108, 1422–1441.
- Zou, L., Mostafavi, A., Zhou, B., Lin, B., Mandal, D., Yang, M., Abedin, J., Cai, H., 2023. Geoai for disaster response, in: *Handbook of geospatial artificial intelligence*. CRC Press, pp. 287–304.