Extracting, Locating and Visualizing Geospatial Rescue Information in German-language Social Media

Xi Yan
University of Hamburg
xi.yan@uni-hamburg.de

Cedric Möller*
University of Hamburg,
cedric.moeller@uni-hamburg.de

David Rath
University of Hamburg
david.rath@studium.uni-hamburg.de

Patrick Westphal
Hamburger Informatik Technologie-Center e.V.
patrick.westphal@hitec-hamburg.de

Ricardo Usbeck
Leuphana University of Lüneburg
ricardo.usbeck@leuphana.de

Martin Semmann
University of Hamburg
martin.semmann@uni-hamburg.de

ABSTRACT

Timely extraction of rescue-related data from social media is vital for emergency response, with event extraction and geolocation playing a key role. This paper presents a demo system that leverages Large Language Models (LLMs) and Knowledge Graphs (KGs) to identify rescue-related data from social media streams and integrate this information into a continuously updated KG, with a focus on the German city of Hamburg. Our approach utilizes an LLM to process unstructured social media text, accurately identifying events and relevant location references. LLMs in combination with in-context learning are applied for event extraction as well as geoparsing. The extracted and linked information is stored in a KG, which is both queryable for further analysis and supports downstream applications such as interactive map-based visualizations, providing real-time awareness for emergency services. Specifically, our geoparsing methods bridge the gap in the German setting, achieving state-of-the-art performance on the benchmark dataset MobIE.

Keywords

Large Language Models, Knowledge Graphs, Event Extraction, Geospatial Entity Linking, Interactive Map Visualization

INTRODUCTION

As natural disasters grow in frequency and severity, timely and accurate information is vital for effective response. Social media offers real-time insights, but extracting and geolocating disaster-related data remains a major challenge—particularly in low-resource settings like German (Takahashi et al. 2015; Suwaileh et al. 2023). While progress has been made for English, German geospatial extraction lags behind (Hennig et al. 2021), and inconsistent location granularity further hinders real-time map-based visualization for emergency responders.

However, with the development of LLMs that support multilingualism and demonstrate generalizability on unseen data(Zhang et al. 2024), KGs that encompass extensive and up-to-date geospatial information, and interactive geospatial visualizations, we aim to bridge these gaps and create a system that can reliably extract, link, and visualize rescue-related geospatial entities from social media data.

*corresponding author

It is important to clarify the intended operational role of our system. The system is *not* a replacement for established emergency reporting channels such as the European emergency number 112, nor does it aim to autonomously detect or triage incoming distress calls. Instead, it functions as a complementary situational-awareness layer for emergency services: a continuous stream of public social media posts is monitored, geoparsed, and fused with municipal geospatial data, so that dispatchers and incident commanders gain a broader spatial picture *around* incidents that are already known or emerging through official channels. Posts such as a user reporting that a street is impassable, that a basement is flooding, or that a shelter is overcrowded are not actionable distress calls in themselves, but, when aggregated and localized, they enrich the operational picture in ways that 112 calls alone cannot provide.

As an illustrative scenario, consider a storm-surge event in Hamburg's Altona district. While 112 dispatchers handle individual emergency calls, our system simultaneously surfaces social media posts mentioning flooded streets, blocked underpasses, or stranded residents, links each mention to OSM coordinates, and overlays them on the interactive map alongside layers for nursing homes, schools, and emergency shelters. An incident commander can then query the underlying knowledge graph for, e.g., all reported flooding within 1 km of a nursing home, supporting proactive evacuation decisions that would be difficult to derive from individual 112 calls alone. In current practice, emergency services manually monitor social media by searching for common keywords or following known accounts. Our system runs in parallel to this workflow: new geolocated events are surfaced on the interactive map, and dispatchers can click through to the originating post to triage and respond accordingly. Recent versions additionally include a flagging mechanism that lets operators mark posts as relevant, irrelevant, or requiring follow-up, providing a feedback signal for downstream filtering.

Our research enables automatic aggregation and location mapping of rescue-relevant social media signals, complementing existing emergency-response workflows. This is critical for Hamburg, which is frequently affected by water-related disasters, where timely information saves lives and optimizes resources. We collaborate with Hamburg rescue teams to build a monitoring system that provides real-time assistance. Beyond this use case, our work contributes a scalable AI framework for crisis response adaptable to various disaster types and regions (Bräker et al. 2022).

Our system pipeline is illustrated in Figure 1. When a social media text containing rescue information is received, a geospatial entity extractor processes the text to identify location entities. These entities are then linked to their corresponding entries in OpenStreetMap (OSM) (Bennett 2010), which provides the geographic coordinates. The extracted entities and their relations, enriched with metadata (timestamps, provenance, entity types), are populated into a knowledge graph (KG). This KG forms the central integration point and can be queried for complex information needs (e.g., "all reported floods within 1 km of nursing homes in Altona district"), in addition to powering the interactive map, which is just one application visualizing KG content based on query results. Refer to Appendix for our code base .

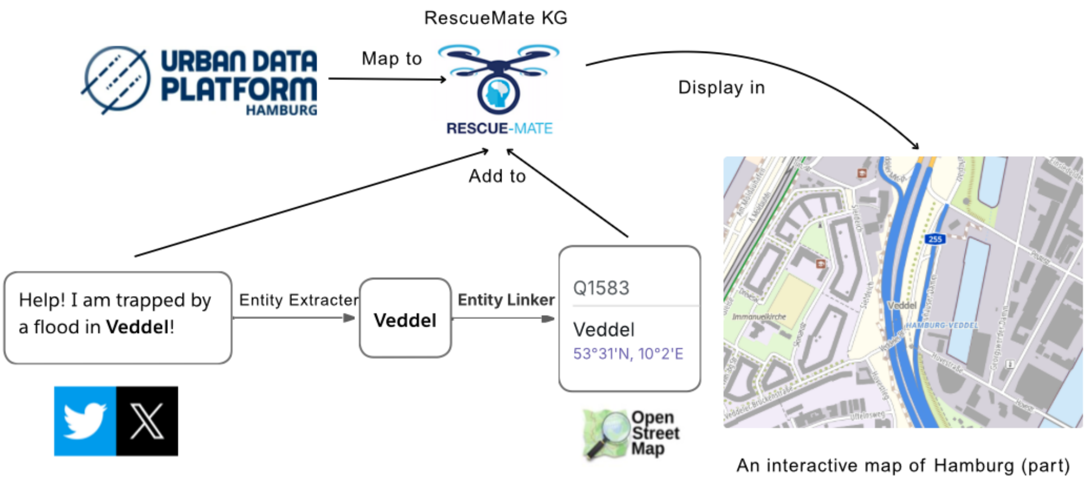

Figure 1. An illustrative example demonstrates our pipeline, which processes a social media post by extracting a geospatial entity, linking it to OpenStreetMap, adding it to the RESCUE-MATE KG, linking to Urban data platform and visualizing it on a map. English instance is used to enhance readability. The depicted post is illustrative of the kind of unstructured text the pipeline processes; the system is not designed to act on individual distress messages, which should be reported through official channels (e.g., 112).

In this work, we present the following key contributions:

1. **A German Geospatial Entity Extractor and Linker** We develop a few-shot learning framework utilizing LLM for extracting geospatial entities from social media text in the German context. This is the first and state-of-the-art system on German geo-entity extraction and linking based on evaluation from the MobIE dataset (Hennig et al. 2021).

2. **Interactive Geo-Visualization** We populate a knowledge graph with extracted and linked geospatial entities, storing not just locations but also event relations, provenance, and entity types. Downstream applications—including an interactive map of Hamburg—are powered by real-time queries over this KG, demonstrating its flexibility for situational awareness and emergency response.

This work is one component of the broader **RESCUE-MATE** project, which develops AI-supported situational awareness tools for emergency services in complex crisis situations. Within RESCUE-MATE, a continuous stream of social media posts is processed: an event extraction and clustering module (developed separately) is employed to detect and group posts corresponding to specific incidents. These grouped posts are then processed by the geospatial linking system presented in this paper. The results (events, locations, organizations, and their relations) are instantiated as nodes and edges in the underlying KG. Each KG entity is annotated with provenance information (e.g., source post, extraction timestamp, confidence score), facilitating traceability and reliability. The KG supports flexible, multi-hop queries (e.g., event chains, proximity searches, organization-event involvement) and serves as the canonical data store for applications, including, but not limited to, the presented interactive map.

GERMAN GEOLOCATION EXTRACTION AND LINKING

Related Work

To our knowledge, few works have addressed **German geolocation extraction and linking systems**. While some studies (Halike et al. 2023; De Cao et al. 2022) have focused on German entity extraction and linking in the general domain utilizing prompt engineering and LLM, they primarily focus on broad entity types such as persons, locations, organizations, human-made products, and temporal expressions (Hamdi et al. 2021). Also, they primarily focus on news or articles such as Wikipedia and are trained on news articles data.

The only dataset focusing on German geolocation extraction and linking from social media text is MobIE (Hennig et al. 2021). This corpus comprises 3,232 social media texts and traffic reports, containing 20.5K manually annotated entities linked to OpenStreetMap[1] (OpenStreetMap contributors 2017). The annotation schema of MobIE encompasses 20 distinct entity types related to the mobility domain. For our study, we concentrate on seven entity types that can be linked to geographic entities in our map: location, location-city, location-route, location-stop, location-street, organization, and organization-company.

Geolocation Extraction

To address the nested entity problem of **named entity recognition (NER)**, where multiple entities may exist within a single mention (Katiyar and Cardie 2018), we employ a text generation approach rather than traditional sequence tagging. On top, we implement a few-shot learning(Wang et al. 2020) approach, which sends the LLM a prompt that states the instruction, *n* examples, and the unlabeled social media text. We discovered that dynamic few-shot learning, which selects semantically more relevant examples as part of the input at inference time, is achieving outstanding performance. This is done by calculating the semantic similarity between the input text and the training data. This example selection framework has demonstrated effectiveness across various tasks (Nashid et al. 2023; Parnami and Lee 2022). The output will include extracted geo entities, along with their predicted entity types by the model. The prompts are available on our GitHub.

Evaluation on the MobIE dataset demonstrates the effectiveness of our approach, achieving strong performance with precision, recall, and F1 scores of 0.8113, 0.7624, and 0.7861, respectively. We utilize LLaMA-3[2] (Touvron et al. 2023) as our base model, which effectively processes the German language while maintaining computational efficiency within our available resources. We experimented with several methods, including LoRA tune (Hu et al. 2022) and prompt tuning, with dynamic few-shot learning yielding the best results. Comparative performance metrics are presented in Table 1. As can be seen from the table, a dynamic few-shot on top of LLaMA greatly improves the model performance over the baseline and reaches an ideal score. For deployment, we transitioned to the Qwen model with a larger parameter size.

[1] https://www.openstreetmap.org/
[2] https://huggingface.co/meta-LLaMA/Meta-LLaMA-3-8B-Instruct

Table 1. Entity Extraction Model Performance Comparison on MobIE Dataset. The best performance is marked in bold. Note that LLaMA is short for Meta-LLaMA-3-8B-Instruct. BERT is short for dbmdz/bert-base-german-cased, and Qwen for Qwen/Qwen2.5-14B-Instruct.

Model	Precision	Recall	F1
Qwen (Dynamic Few-Shot)	**0.8049**	**0.7815**	**0.7930**
LLaMA (Dynamic Few-Shot)	0.7148	0.7961	0.7532
LLaMA (Few-Shot)	0.7127	0.7198	0.7162
LLaMA + LoRA	0.8343	0.4796	0.6091
BERT (Baseline)	0.5687	0.7006	0.6278

Geolocation Linking

Following entity extraction, we perform a two-step entity linking process to map extracted mentions to their corresponding entries in OSM. This involves **candidate retrieval** followed by **disambiguation** based on contextual and geographic relevance.

Candidate retrieval is done by employing locally running nominatim [3] and photon [4] services to retrieve candidates for each identified entity. Instead of just querying for each mentioned entity as identified, we enrich the query with additional information, such as the mentioned district, street, or similar, if found in the text.

For **disambiguation**, we represent each location candidate by its full address as well as its OpenStreetMap type and class [5]. The set of all candidates and the initial text are fed into an LLM to decide on the final candidate. The LLM is also prompted to identify whether no location matches. This is a necessary step as candidate retrieval might be incomplete, leading to false positives.

We implement two key constraints to optimize entity linking precision[6]: 1. Location type constraints that specify entities must belong to predefined location categories (e.g., city, museum, library, nightclub, etc.) relevant to our use case. 2. Geographic range constraints that limit matches to locations within Hamburg, including all districts, such as Altona (we remove this constraint in the evaluation phase).

We link entities to OSM for benchmarking purposes, as OSM entries consistently include geographic coordinates. This enables precise spatial mapping of identified entities.

We evaluate our method on the MobIE test set, which includes 623 social media posts with 1,824 linked entities—184 with OSM IDs and 1,086 with coordinates. Due to the limited number and poor quality of OSM annotations, we opt for coordinate-based evaluation. Many locations of interest, such as *Hamburg*, are either unannotated or incorrectly linked. In some cases, mismatches between the annotated label and the entity type returned by OSM prevent successful linking, resulting in NIL assignments. Additionally, the NER annotations include errors such as missing entities and malformed mentions (e.g., extra "#" characters). Following recent work in geoparsing (Liu et al. 2022), we adopt **Accuracy@161**, which considers a location correctly resolved if it lies within 161 km of the gold reference. The result is shown in Table 2. While Accuracy@161 provides a coarse measure of linking correctness, it does not capture per-type performance differences. Importantly, our geolinker can predict NIL when none of the retrieved candidates are deemed correct by the LLM, effectively trading recall for precision. This distinguishes it from standalone geocoders like Nominatim and Photon, which always return a top-ranked result. We therefore additionally report precision, recall, and F1 per entity type in Table 3, which reveals that the system performs strongly on well-defined types such as *location-city* (F1 0.967) and *location-stop* (F1 0.972), while *location-route* (F1 0.119) and *location-street* (F1 0.453) remain challenging. The per-type breakdown should be read with this recall–precision trade-off in mind.

Table 2. Evaluation results for the geo parsing system on the test set.

Metrics	LLaMA	Qwen
Accuracy@161	0.4700	0.6576

[3] https://nominatim.org
[4] https://photon.komoot.io
[5] The type and class information is necessary as often multiple candidates have the same address, such as a street and a bus stop in the street.
[6] Noted that both type and range constraints are parameterized to maintain system generalizability across different geographic regions and use cases.

Error analysis. A manual inspection of a sample of incorrectly linked or NIL-assigned mentions reveals three recurring failure modes. (1) *Candidate retrieval gaps:* informal abbreviations, colloquial street names, and partial mentions (e.g., "Reeperbahn" without district context) sometimes yield no relevant Nominatim/Photon candidate, making correct disambiguation impossible downstream. (2) *Annotation noise in MobIE:* a non-trivial fraction of gold mentions are unlinked, mis-typed, or contain artifacts such as residual hashtag characters, which depress the upper bound of any system evaluated on this corpus. (3) *Ambiguous toponyms:* generic names shared across multiple German cities (e.g., "Hauptbahnhof", "Marktplatz") remain difficult even with our Hamburg range constraint relaxed during evaluation. Together, (1) and (3) account for the majority of the gap between Qwen and a perfect linker, and motivate the fine-grained-geoparsing direction noted in Section 4.

To isolate the contribution of LLM-based disambiguation from the entity extraction step, we additionally evaluate the linking component in isolation using gold entity mentions from the MobIE test set. Table 3 compares our full geolinker (Nominatim + Photon + LLM disambiguation) against using Nominatim and Photon directly as standalone geocoders.

Table 3. Linking performance comparison using gold entity mentions. Nominatim and Photon are used as standalone geocoders; geolinker adds LLM-based disambiguation on top of both.

	Nominatim			**Photon**			**Geolinker**		
Type	P	R	F1	P	R	F1	P	R	F1
location	.857	.851	.854	.753	.751	.752	**.910**	**.844**	**.876**
location-city	.960	.955	.957	.921	.919	.920	**.982**	**.953**	**.967**
location-route	**.380**	**.372**	**.376**	.000	.000	.000	.435	.069	.119
location-stop	.968	.941	.954	.941	.941	.941	**.990**	**.955**	**.972**
location-street	.266	.251	.258	.099	.099	.099	**.598**	**.364**	**.453**
overall	.788	.770	.779	.694	.692	.693	**.918**	**.770**	**.838**

The geolinker achieves the highest overall precision (0.918) and F1 (0.838), demonstrating that LLM-based disambiguation substantially improves linking quality over standalone geocoders. The gains are most pronounced for *location-street* (F1 from 0.258/0.099 to 0.453), where ambiguous street names benefit most from contextual disambiguation. For well-defined entity types such as *location-city* and *location-stop*, all three systems perform strongly, though the geolinker still adds a consistent margin. The exception is *location-route*, where the geolinker's strict filtering yields high precision (0.435) but very low recall (0.069), as many route mentions lack sufficiently specific candidates in OSM. Nominatim achieves a better balance here (F1 0.376) because its default ranking by population often selects the most prominent candidate, which happens to be correct, whereas the LLM attempts to use contextual cues that may not match any available candidate and thus returns NIL more frequently. This suggests that combining population-based priors with contextual disambiguation for route-type entities is a promising direction for improvement.

INTERACTIVE MAP OF HAMBURG CITY

The interactive map is an interface that visualizes rescue events and locations by querying the underlying knowledge graph in real time. All map data are retrieved from the KG, maintaining consistency with both the raw extractions and their ongoing updates. A representative output is shown in Figure 2.

Features

One key feature of the interactive map is its support for customizable **layers**, which allow users to filter and highlight designated geospatial entity types directly on the map interface. These include *schools, nursing homes, hospitals, emergency shelters, fire departments, the main dike, water rescue points, flooding areas, train stations, bus stops*, and *ferry stops*. The layers are derived from the Urban Data Platform Hamburg, which we describe in Section , and enable emergency responders to quickly assess the proximity of critical facilities to reported rescue events.

Each map marker is tailored to the event or entity type, providing an intuitive visual cue—e.g., a hospital icon for medical facilities. Clicking on a marker reveals detailed metadata. For example, a state school marker displays its name, type, number of students and contact info. This feature allows users to locate resources and access relevant operational details with a single click.

Another central component of the system is the **Geocoder** module, which processes incoming social media posts. This module combines geospatial entity recognition and linking to identify and link places from unstructured text to

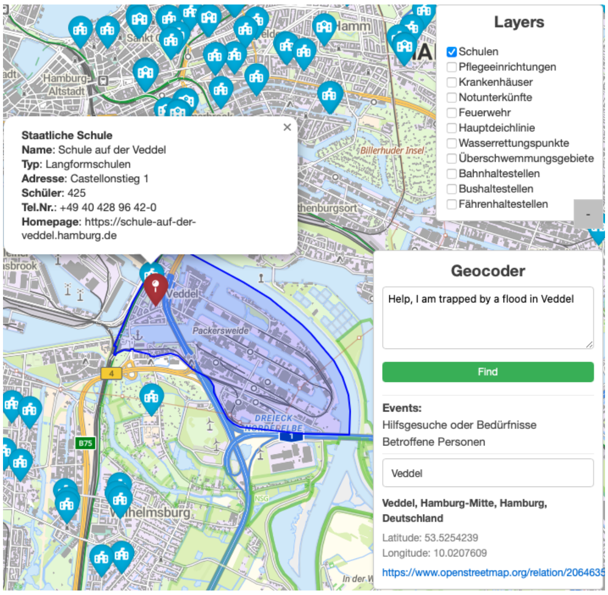

Figure 2. An example of our map, which supports two features: 1. Layers to include the interested location types. 2. Geocoder to take a social media text, return the geo entity(s) and mark it on the map.

OSM. Once linked, the corresponding coordinates are retrieved, and the location polygons are marked on the map along with the event type of the text. This feature enables dynamic, location-aware mapping of rescue needs, even when the text input is informal, ambiguous, or written in low-resource language varieties.

Map Data and Framework

We utilize the Urban Data Platform Hamburg[7] as the backbone of our interactive map. This platform is accessible through an API that implements the OGC API - Features standard (developed by the Open Geospatial Consortium[8]) and includes a wide range of entity types relevant to emergency response, particularly public facilities such as administrative boundaries, road networks and public infrastructure. Its standardized interfaces enable efficient integration and interconnection of urban data.

Our interactive map is built using dash-leaflet[9], a Python wrapper to the Leaflet library[10]. Leaflet is an open-source JavaScript library for creating web-based mapping applications. We chose a Python framework for easier integration with other processing pipelines, e.g., the entity extractor and linker.

SCALABILITY, ROBUSTNESS, AND LIMITATIONS

Throughput and latency. Per-post processing is dominated by the two LLM calls (extraction and disambiguation). With locally hosted geocoders and LLM instances, the deployed pipeline sustains a throughput of roughly 2 posts per second on our hardware (2× NVIDIA RTX A6000) — sufficient for the moderate volume of Hamburg-region rescue-relevant posts observed in practice, though high-volume disaster spikes would still require batching, request-level caching of repeat toponyms, or a smaller distilled model in the disambiguation stage.

Robustness. The system is currently sensitive to noisy input in two ways: (i) very informal or code-mixed posts can degrade extraction quality, and (ii) the geocoder's recall is bounded by Nominatim/Photon coverage of fine-grained or colloquial toponyms, as discussed in the error analysis.

Evaluation scope. Our quantitative results rely on a single benchmark (MobIE), which is the only publicly available German geoparsing corpus covering social-media-style text. We therefore cannot fully disentangle dataset-specific artifacts from genuine model behavior, and a comparison against multilingual or LLM-based geoparsers (e.g., Edinburgh Geoparser, recent prompt-based approaches) on a shared corpus remains future work.

[7] https://api.hamburg.de/datasets/v1/
[8] https://docs.ogc.org/is/17-069r4/17-069r4.html
[9] https://www.dash-leaflet.com
[10] https://leafletjs.com

CONCLUSION AND FUTURE WORK

This paper introduces a system for extracting, linking, and integrating rescue-related geospatial data from German social media into a continuously updated knowledge graph. Using dynamic few-shot learning with LLMs and OSM-based entity linking, it achieves state-of-the-art results. Extracted entities are geolocated and displayed on an interactive Hamburg map to support real-time emergency response. Our KG-centric architecture enables structured queries across events, locations, and organizations, powering both real-time visualization (e.g., Hamburg map) and analytical tasks for emergency response. Future work will explore integrating multimodal data—such as images and videos—to enhance contextual understanding, alongside adding the event clustering and classification module. Furthermore, we are currently exploring ways to improve geoparsing for fine-grained locations. In particular, we are evaluating the system on the fine-grained location-focused DLRGeoTweet dataset[11] using stricter distance-based metrics such as Accuracy@1 km and Accuracy@100 m to better characterize performance on detailed toponyms, and we are experimenting with reinforcement learning to improve query generation for the geocoder, with promising initial results.

ONLINE RESOURCES

The models described in Section 2 and 3 are separately maintained on GitHub[12][13].

ETHICAL STATEMENT

No personal or sensitive data is collected, and all data sources (e.g., social media, news reports) are publicly available and processed in compliance with applicable data protection regulations (e.g., GDPR). The system does not make autonomous decisions; it supports humans by providing structured information. This research respects privacy, avoids harm, and intends to contribute positively to public safety and crisis management.

ACKNOWLEDGEMENT

Generative AI tools were used in the preparation of this manuscript to assist with language refinement, phrasing, and structural suggestions. All scientific content, methodology, analysis, and conclusions were developed and verified by the authors. The authors take full responsibility for the integrity, accuracy, and originality of the work. This work is supported by the Ministry of Research and Education, which supported this project within the SifoLIFE project "RESCUE-MATE: Dynamische Lageerstellung und Unterstützung für Rettungskräfte in komplexen Krisensituationen mittels Datenfusion und intelligenten Drohnenschwärmen" (FKZ 13N16836). We utilized 2 x NVIDIA RTX A6000, kindly provided by the NVIDIA Academic Hardware Grant Program. We use Claude[14] for grammar check and rephrasing.

REFERENCES

Bennett, J. (2010). *OpenStreetMap*. Packt Publishing Ltd.

Bräker, J., Fischer, M., and Semmann, M. (2022). "RescueMate–Rethinking Innovation in Natural Disasters: Building an AI & Data Platform to Better Deal with Storm Surges". In: *ICIS 2022 TREOs. 66*. Copenhagen, Denmark.

De Cao, N., Wu, L., Popat, K., Artetxe, M., Goyal, N., Plekhanov, M., Zettlemoyer, L., Cancedda, N., Riedel, S., and Petroni, F. (2022). "Multilingual autoregressive entity linking". In: *Transactions of the Association for Computational Linguistics* 10, pp. 274–290.

Halike, A., Wumaier, A., and Yibulayin, T. (2023). "Zero-Shot Relation Triple Extraction with Prompts for Low-Resource Languages". In: *Applied Sciences* 13.7.

Hamdi, A., Linhares Pontes, E., Boros, E., Nguyen, T. T. H., Hackl, G., Moreno, J. G., and Doucet, A. (2021). "A multilingual dataset for named entity recognition, entity linking and stance detection in historical newspapers". In: *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 2328–2334.

[11] https://github.com/xukehu/DLRGeoTweet
[12] https://github.com/semantic-systems/sems-digital-twin-map
[13] https://github.com/rescue-mate/social-media-sensor
[14] https://claude.ai

Hennig, L., Truong, P. T., and Gabryszak, A. (June 2021). "MobIE: A German Dataset for Named Entity Recognition, Entity Linking and Relation Extraction in the Mobility Domain". In: *Proceedings of the 17th Conference on Natural Language Processing (KONVENS 2021)*. Ed. by K. Evang, L. Kallmeyer, R. Osswald, J. Waszczuk, and T. Zesch. Düsseldorf, Germany: KONVENS 2021 Organizers, pp. 223–227.

Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., and Chen, W. (2022). "LoRA: Low-Rank Adaptation of Large Language Models". In.

Katiyar, A. and Cardie, C. (June 2018). "Nested Named Entity Recognition Revisited". In: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. Ed. by M. Walker, H. Ji, and A. Stent. New Orleans, Louisiana: Association for Computational Linguistics, pp. 861–871.

Liu, Z., Janowicz, K., Cai, L., Zhu, R., Mai, G., and Shi, M. (2022). "Geoparsing: Solved or Biased? An Evaluation of Geographic Biases in Geoparsing". In: *AGILE: GIScience Series* 3, p. 9.

Nashid, N., Sintaha, M., and Mesbah, A. (2023). "Retrieval-Based Prompt Selection for Code-Related Few-Shot Learning". In: *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*, pp. 2450–2462.

OpenStreetMap contributors (2017). *Planet dump retrieved from https://planet.osm.org*. https://www.openstreetmap.org.

Parnami, A. and Lee, M. (2022). "Learning from Few Examples: A Summary of Approaches to Few-Shot Learning". In: *CoRR* abs/2203.04291. arXiv: 2203.04291.

Suwaileh, R., Elsayed, T., and Imran, M. (2023). "IDRISI-RE: A generalizable dataset with benchmarks for location mention recognition on disaster tweets". In: *Information Processing & Management* 60.3, p. 103340.

Takahashi, B., Tandoc Jr, E. C., and Carmichael, C. (2015). "Communicating on Twitter during a disaster: An analysis of tweets during Typhoon Haiyan in the Philippines". In: *Computers in human behavior* 50, pp. 392–398.

Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al. (2023). "LLaMA: Open and Efficient Foundation Language Models". In: *CoRR* abs/2302.13971. arXiv: 2302.13971.

Wang, Y., Yao, Q., Kwok, J. T., and Ni, L. M. (2020). "Generalizing from a few examples: A survey on few-shot learning". In: *ACM computing surveys (csur)* 53.3, pp. 1–34.

Zhang, S., Gao, C., Zhu, W., Chen, J., Huang, X., Han, X., Feng, J., Deng, C., and Huang, S. (Nov. 2024). "Getting More from Less: Large Language Models are Good Spontaneous Multilingual Learners". In: *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. Ed. by Y. Al-Onaizan, M. Bansal, and Y.-N. Chen. Miami, Florida, USA: Association for Computational Linguistics, pp. 8037–8051.