

Generating Synthetic Oracle Datasets to Analyze Noise Impact: A Study on Building Function Classification Using Tweets

Shanshan Bai^{1,3,*}, Anna Kruspe² and Xiaoxiang Zhu^{1,3}

¹Technical University of Munich, Munich, Germany

²Munich University of Applied Sciences, Munich, Germany

³Munich Center for Machine Learning, Munich, Germany

Abstract

Tweets provides valuable semantic context for earth observation tasks and serves as a complementary modality to remote sensing imagery. In building function classification (BFC), tweets are often collected using geographic heuristics and labeled via external databases—an inherently weakly supervised process that introduces both label noise and sentence-level feature noise (e.g., irrelevant or uninformative tweets). While label noise has been widely studied, the impact of sentence-level feature noise remains underexplored, largely due to the lack of clean benchmark datasets for controlled analysis. In this work, we propose a method for generating a synthetic oracle dataset using LLM, designed to contain only tweets that are both correctly labeled and semantically relevant to their associated buildings. This oracle dataset enables systematic investigation of noise impacts that are otherwise difficult to isolate in real-world data. To assess its utility, we compare model performance using Naïve Bayes and mBERT classifiers under three configurations: real vs. synthetic training data, and cross-domain generalization. Results show that noise in real tweets significantly degrades the contextual learning capacity of mBERT, reducing its performance to that of a simple keyword-based model. In contrast, the clean synthetic dataset allows mBERT to learn effectively, outperforming Naïve Bayes by a large margin. These findings highlight that addressing feature noise is more critical than model complexity in this task. Our synthetic dataset offers a novel experimental environment for future noise injection studies and is publicly available on GitHub¹.

Keywords

Noise, Tweets, Synthetic Data Generation, Building Function Classification, LLM,

1. Introduction

Building Function Classification (BFC) aims to determine the functional purpose of buildings, such as commercial or residential use, using various data sources. Traditionally, remote sensing imagery has been the dominant modality for this task. However, such imagery often lacks the semantic granularity needed to distinguish nuanced urban functions at the building level, for example, differentiating between dormitories and hotels.

Textual data from social media, such as geo-tagged tweets, offers a complementary perspective. Tweets sent from the same location as a building may reveal human activities and behaviors that provide contextual clues about building use. For instance, a tweet like "*Great coffee at this café!*" implies a commercial function, while "*Silent night at the dorm.*" suggests residential use.

Despite this promise, social media datasets collected for BFC often rely on weak supervision—tweets are heuristically matched to buildings (e.g. tweets sent within 50 meters radius of a building are assigned to this building) and labeled using voluntarily-provided building tags on the external databases such as OpenStreetMap (OSM). This process introduces multiple types of noise: (1) label noise, from incorrect or outdated building tags; and (2) sentence-level feature noise, in the form of irrelevant, ambiguous,

¹https://github.com/zhu-xlab/synthetic_tweets

GeoExt 2025: Third International Workshop on Geographic Information Extraction from Texts at ECIR 2025, April 10, 2025, Lucca, Italy

*Corresponding author.

✉ shanshan.bai@tum.de (S. Bai); anna.kruspe@hm.edu (A. Kruspe); xiaoxiang.zhu@tum.de (X. Zhu)

🆔 0009-0000-4786-8809 (S. Bai); 0000-0002-2041-9453 (A. Kruspe); 0000-0001-5530-3613 (X. Zhu)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

or uninformative tweets. While label noise has been extensively studied—often through controlled label-flipping experiments—sentence-level feature noise remains harder to investigate. This is because it requires access to a dataset that is known to contain only relevant and correctly labeled examples, something that human annotators often cannot guarantee due to the subjective and implicit nature of semantic relevance.

To address this, we explore the feasibility of using LLM to generate a synthetic oracle dataset for BFC: a noise-free benchmark in which all tweets are both correctly labeled and semantically relevant to their associated buildings. We define an oracle dataset as an idealized collection designed not for deployment, but to serve as a clean experimental environment for systematically analyzing the effects of noise.

It is also important to acknowledge that geo-tagged tweets are no longer widely accessible due to platform changes (e.g., Twitter removing precise geolocation tagging [1]). However, our focus is not on Twitter per se, but on the broader challenge of noise in weakly supervised user-generated datasets—an issue that persists across many social media applications.

Our key contributions are as follows:

A Synthetic Oracle Tweets Dataset for Building Function Classification: We construct a clean, LLM-generated dataset that enables controlled experiments on feature noise, facilitating analysis that is otherwise infeasible using real-world data.

A Data Generation Pipeline: We introduce a reproducible method for generating synthetic datasets, guided by real-world building and tweet distributions to ensure statistical realism.

Evaluation of Data Quality: We compare the correctness and diversity of the synthetic data against real-world datasets using both classifier performance and linguistic metrics (Self-BLEU, perplexity).

Insights into Feature Noise Impact: Our results show that handling feature noise is more critical than increasing model complexity for BFC tasks.

The remainder of this paper is structured as follows: Section 2 reviews related work on building function classification and synthetic data generation. Section 3 describes our data generation pipeline. Section 4 presents the evaluation methodology and results. Section 5 analyzes the quality and characteristics of the dataset. Section 6 summarizes our findings, and Section 7 outlines directions for future research.

2. Related Work

Building Function Classification with Tweets

Recent studies have investigated the feasibility of using geo-tagged tweets for Building Function Classification (BFC), confirming their potential while emphasizing key challenges related to noise and data quality. Häberle et al. [2] were among the first to treat BFC as a text classification task, assigning tweets located near buildings as inputs. However, their use of FastText limited robustness in multilingual urban settings. In contrast, Häberle et al. [3] shifted focus from sentence-level classification to feature engineering, using tweet embeddings to extract function-indicative features. While some keywords were found to correlate strongly with building types, overlapping lexical features introduced noise, reducing overall reliability.

To enhance classification accuracy, Häberle et al. [4] proposed a decision-level fusion strategy, integrating textual features from tweets with remote sensing imagery. Their results showed that while tweets offered useful contextual cues, the discriminative power largely came from imagery. More recently, Häberle [5] scaled BFC globally, reframing it as a multilingual classification task using mBERT. Despite achieving moderate success (55% accuracy across three functional categories), noise in the tweet data—both in labeling and content—remained a significant bottleneck.

A consistent finding across these studies is that tweets, while informative, are inherently noisy. Label noise often originates from incorrect or incomplete tags in OpenStreetMap (OSM) [6, 7], and geo-tagging mismatch occurs when tweets are attributed to the wrong building due to spatial heuristics [8]. Additionally, sentence-level feature noise—such as irrelevant tweets—remains a major challenge [9]. While label noise has been the subject of many experimental studies, there has been limited systematic analysis of sentence-level feature noise and its effects on model performance.

Synthetic Data Generation with LLMs

Large Language Models (LLMs) have shown impressive ability to generate text that captures patterns and structures from real-world corpora [10]. Recent work has explored using LLMs to create synthetic datasets via instruction-tuned prompting conditioned on class labels [11, 12], showing that LLMs can serve as scalable alternatives to human annotators. However, category-conditioned generation poses a critical challenge: the synthetic text distribution often diverges from real data in linguistic diversity, topical focus, or context relevance [13, 14].

To address these issues, researchers have proposed human-in-the-loop refinement strategies [15] and model-guided feedback loops [16], which adjust prompts based on classifier performance. Others have explored persona-based synthesis [17, 18, 19, 20] to introduce style and perspective variation. While these approaches enhance fluency and diversity, they are not explicitly designed for tasks requiring spatial and contextual grounding—such as aligning tweets with building-specific metadata in urban environments.

Despite these advances, there remains a gap in applying LLM-based data generation to fine-grained, geospatially anchored classification tasks like BFC. Most prior work assumes structured label spaces and ignores the semantic constraints imposed by geographic or functional context. Our work addresses this gap by conditioning LLM prompts on real-world building attributes (e.g., function, name, location) and matching tweet language distributions. This enables the creation of a high-quality oracle dataset that supports controlled analysis of noise in text-based BFC.

3. Synthetic Tweets Generation Pipeline

This section presents our three-step pipeline for generating synthetic multilingual tweets using a large language model (LLM), conditioned on building-level metadata. The pipeline includes: (1) retrieving metadata, (2) cleaning it, and (3) prompting the LLM to generate contextually grounded tweets.

3.1. Step 1: Retrieving Metadata

We define metadata as structured descriptive information associated with each building—such as its function, location, and the intended number and language of tweets to be generated. This metadata is either inherited from a prior labeled dataset [5] or retrieved from OSM, ensuring that the resulting synthetic dataset reflects the statistical characteristics of real-world data.

- **Building Attributes:**

- **Building_tag:** A fine-grained functional label from OSM (e.g., "restaurant", "apartment"), distinct from the binary ground-truth classification ("commercial" or "residential").
- **Building_name:** A descriptive identifier used in tweet content (e.g., "Merlex Auto Group").
- **Building_city:** The city where the building is located.

- **Tweets Language Distribution:** A list specifying the language of each synthetic tweet to be generated. For instance, ["English", "English"] indicates two English tweets should be generated for that building. These values are inherited from the dataset in [5], ensuring alignment with real-world tweet frequency and language usage.

3.2. Step 2: Preprocessing Metadata

Since building attributes derived from OSM can be noisy, and we also want to control the maximum number of tweets each building, we apply a cleaning preprocessing pipeline to previous collected metadata. This step ensures LLM generating synthetic tweets without introducing unwanted noise and also a balanced dataset.

- Removing formatting artifacts: We sanitize building names and tags by stripping special characters (e.g., underscores `_`, slashes `/`) to prevent malformed prompts.
- Filtering generic or erroneous tags: Entries with non-informative tags such as `"yes"` or `"roof"` are excluded, as they do not convey meaningful function.
- Ensuring unique tag: Many buildings are associated with multiple use-specific tags that suggest different functional categories (e.g., both `"church"` and `"restaurant"`).
- Ensuring label-tag consistency: For buildings pre-labeled as `"commercial"` or `"residential"` (from prior datasets), we remove those whose OSM tags clearly contradict the assigned label. For example, if a building labeled `"commercial"` has the tag `"mosque"`, it is removed from the dataset.
- Restricting to at most five tweets per building: This step applies to the number of synthetic tweets to be generated and also to real-world tweets. We cap the number at five to balance the dataset across buildings and to control for prompt length and generation consistency. This helps avoid overrepresentation of a few buildings during training or evaluation phases.

The preprocessed metadata is stored in structured JSON format. An example is shown below:

```
{ "Building city": "WashingtonDC",  
  "Building tag": "Retail",  
  "Building name": "Merlex Auto Group",  
  "Tweets distribution": [ "English", "English" ] }
```

While these preprocessing steps help reduce ambiguity and maintain consistency, we acknowledge that it may exclude buildings with complex, multi-functional roles (e.g., shopping malls with food courts, theaters, and offices). Although this limits the coverage of our analysis, our focus in this study is to establish a controlled, noise-free experimental setting—not to comprehensively model all real-world function types.

3.3. Step 3: Generating Tweets using an LLM

We use the `Llama-3.3-70B-Instruct-bnb-4bit` model from Hugging Face¹ to generate multilingual tweets. Each prompt consists of two key components:

- System Prompt: Defines the overall task, outlining style and formatting guidelines for tweet generation. The system prompt also includes a one-shot example demonstrating the desired tweet format. This remains constant for all buildings.
- User Prompt: Contains building-specific metadata, to ensure diverse and contextually relevant outputs.

Here is the system prompt:

¹<https://huggingface.co/unsloth/Llama-3.3-70B-Instruct-bnb-4bit>

Generate tweets as if they were posted by real Twitter users in a specific building. Tweets should be sent from the type of building described in the 'building tag'. Ensure that each tweet reflects a unique perspective or experience. Imagine switching personas for each tweet, simulating thoughts from different types of users, such as tourists, professionals, or families. Consider varying the tone (e.g., humorous, cynical, formal, casual), the length (short and concise, or longer and detailed), and the use of mentions or hashtags. Highlight varied aspects of the building, such as its architecture, services, location, history, or events. You must generate only one tweet in each language specified under 'tweet language distribution,' written directly in that language.

Example:

```
{"Building_city": "WashingtonDC",  
  "Building_tag": "Retail",  
  "Building_name": "Merlex Auto Group",  
  "Tweets_language_distribution": ["English", "English"]}
```

["Bought new rims here at Merlex Auto yesterday, totally transformed my ride! @merlexautogroup #AutoCare", "Merlex Auto Group really knows how to treat car lovers right. The staff? Super knowledgeable. The selection? If you're in DC and thinking about upgrading your ride, this is the place! #CarShopping #DCLuxury"]

Using this three-steps pipeline, we generate a synthetic oracle dataset covering 6,000 real-world buildings across 41 global cities. The final dataset includes 15,222 tweets in 45 languages. We ensure that the distribution of building types, tweet language frequencies closely mirror those of the real-world dataset used in [5]. A quantitative comparison between real and synthetic tweets is provided in Table 1.

4. Dataset Evaluations and Results

In this section, we evaluate the quality of the generated synthetic dataset along two key dimensions: diversity and correctness. Diversity ensures that the dataset captures a broad range of sentence structures and vocabulary variations, rather than overly repetitive content that could oversimplify the classification task. Correctness assesses whether the synthetic dataset fulfills its intended purpose as an oracle dataset, containing only perfectly aligned tweets that semantically correspond to their respective target buildings.

Since each building in the synthetic dataset mirrors a real-world building exactly, and the tweet distributions in terms of volume and language match their real-world counterparts, we evaluate correctness and diversity by comparing the synthetic dataset against the real-world dataset.

4.1. Diversity Evaluation

To assess diversity, we use 4-gram Self-BLEU [21] as the primary metric, following [22]. Self-BLEU measures how similar each sentence is to the rest of the dataset by calculating BLEU [21] scores for every sentence against all others. A lower Self-BLEU score indicates higher diversity, suggesting that the dataset contains a richer variety of expressions. The results are reported in Table 2.

As a secondary measure, we compute perplexity [23]. While traditionally used to assess how well a language model predicts text, perplexity can also serve as an indirect proxy for vocabulary alignment between the synthetic and real-world datasets. Specifically, we compute unigram perplexity using a unigram language model pre-trained on a held-out set of 100,000 real-world tweets from buildings not included in the experimental dataset. By evaluating perplexity for both datasets, a similar score between synthetic and real-world data suggests that the synthetic dataset shares comparable linguistic distribution and vocabulary complexity. Conversely, a significant deviation would indicate substantial differences in word usage. The perplexity results are also included in Table 2.

Category	Commercial	Residential
Tag	Retail	Dormitory
Name	Superfresh	Baxter Hall
City	New York	Cape Town
Tweet Distribution	[English, English, English]	[English, English, English, English, English]
Real Tweets	<p>"Just posted a photo @ SuperFresh Supermarket of Bay Ridge."</p> <p>"Turkey Dinner with all the Fixin's, that your kids will actually gobble up?! Yep! Cake #nailedit Spotted at superfresh_bayridge_dyker (where they do NOT make vegan pumpkin pie btw) SuperFresh Bay Ridge."</p> <p>"New York - The food in the deli, had to go to the hospital where they pumped my stomach out and determined the deli salads poisoned me. I was so... Food Poisoning."</p>	<p>"Just posted a photo @ Baxter Hall."</p> <p>"Happiness is a decision @ Baxter Hall."</p> <p>"Just posted a video @ Baxter Hall."</p> <p>"Shooting visserlieb clarinet improv #artistic #research #clarinet #performancebasedresearch #phdlife #phdstudent #archive #freeimprovisation #kirbycollection #music..."</p> <p>"Shooting visserlieb clarinet improv. #artistic #research #clarinet #artisticresearch #performancebasedresearch #phdlife #phdstudent #archive #freeimprovisation #kirbycollection #music #experimentation #sound..."</p>
Synthetic Tweets	<p>"Fresh produce at unbeatable prices, Superfresh has become my go-to grocery spot in New York!"</p> <p>"Why did I wait so long to discover the best deals on electronics are actually inside this store?"</p> <p>"Does anyone else notice how vibrant the flower arrangements are near the entrance every morning? It sets the mood right for shopping."</p>	<p>"Cramming for exams is way more bearable when you have an amazing view like ours at Baxter Hall!"</p> <p>"Dorm life can be crazy but our community here makes all the difference."</p> <p>"Battling to get work done before deadline hits - anyone else feeling the pressure? Hope caffeine kicks in soon."</p> <p>"Sharing laughs over breakfast on the rooftop has become daily highlight always something new."</p> <p>"Wonderful place full great people making every day count."</p>

Table 1
Comparison of Real-World and Synthetic Tweets of Two Buildings

Dataset	Self-BLEU↓	Perplexity(log ₁₀)↓
Synthetic	48.37%	4.49
Real-world	40.78%	4.36

Table 2
Diversity Comparison Results by Self-BLEU and Perplexity on Synthetic and Real-world Datasets

4.2. Correctness Evaluation

A direct way to evaluate correctness is by testing the dataset's utility in a downstream classification task. We compare two distinct classifiers: Naïve Bayes (NB) and multilingual BERT (mBERT). These models were chosen to contrast a traditional statistical classifier (NB), which relies on word co-occurrence patterns, with a transformer-based language model (mBERT), which leverages contextual embeddings for nuanced language understanding. We conduct evaluations under three configurations:

- **Real-world:** Train and test models using real-world data to establish a baseline.
- **Synthetic:** Train and test models using synthetic data to determine if the dataset achieves correctness.
- **Cross-domain:** Train models on synthetic data and test on real-world data to evaluate generalization performance.

Configurations	Model	Commercial			Residential			Overall Accuracy
		Precision	Recall	F1	Precision	Recall	F1	
Real-world	NB	0.67 ± 0.03	0.57 ± 0.02	0.61 ± 0.02	0.61 ± 0.02	0.71 ± 0.02	0.66 ± 0.02	0.64 ± 0.02
	mBERT	0.66 ± 0.01	0.68 ± 0.03	0.67 ± 0.02	0.66 ± 0.02	0.62 ± 0.01	0.64 ± 0.01	0.65 ± 0.01
Synthetic	NB	0.86 ± 0.02	0.82 ± 0.03	0.84 ± 0.02	0.81 ± 0.03	0.85 ± 0.01	0.83 ± 0.02	0.84 ± 0.02
	mBERT	0.92 ± 0.03	0.90 ± 0.01	0.91 ± 0.01	0.90 ± 0.01	0.92 ± 0.02	0.91 ± 0.01	0.91 ± 0.01
Cross-Domain	NB	0.60 ± 0.03	0.67 ± 0.01	0.63 ± 0.02	0.60 ± 0.02	0.52 ± 0.02	0.56 ± 0.02	0.60 ± 0.01
	mBERT	0.59 ± 0.06	0.77 ± 0.06	0.66 ± 0.01	0.66 ± 0.02	0.44 ± 0.15	0.52 ± 0.10	0.61 ± 0.05

Table 3

Performance comparison of Naïve Bayes (NB) and mBERT models across three configurations: Real-world (train and test on real-world data), Synthetic (train and test on synthetic data), and Cross-Domain (train on synthetic data and test on real-world data).

First two configurations allow us to assess whether the synthetic dataset can function as a valid oracle dataset. The cross-domain evaluation is particularly insightful in determining how well models trained on noise-free synthetic data adapt to noisy real-world conditions.

The real-world dataset used in our experiments was derived from [5], filtered to include the same 6,000 buildings as the synthetic set. For each building, we retain the same number of tweets in the same languages, resulting in a dataset of 15,222 real tweets. We split both datasets at the building level to prevent data leakage, allocating 10,729 tweets from 80% of buildings for training and validation, and 2279 tweets of the rest 20% buildings for testing. We fine-tune bert-base-multilingual-cased for five epochs using the Adam optimizer with a learning rate of 5e-6, a batch size of 16, a warm-up ratio of 0.01, and a weight decay of 0.01 to prevent overfitting. All training hyperparameters remain consistent across experiments. The results are presented in Table 3.

5. Discussions

Our diversity evaluation reveals a trade-off between structural variation and lexical realism in the synthetic dataset. The higher self-BLEU score of the synthetic tweets (48.37%) compared to the real-world tweets (40.78%) suggests that the synthetic content is more repetitive at the sentence structure level. However, the perplexity scores—4.49 for synthetic vs. 4.36 for real—indicate that the vocabulary distribution is comparable across both datasets. This suggests that while the LLM may favor certain patterns or templates in sentence generation, it still captures a realistic range of word usage and topic-relevant terms.

In the real-world configuration, mBERT (65%) performs only marginally better than Naïve Bayes (64%). This finding suggests that the contextual capabilities of transformer models are suppressed by the high level of noise in real tweets. Naïve Bayes, which relies on surface-level word co-occurrences, performs nearly as well—indicating that complex models may be overkill in noisy, weakly supervised settings where semantic signals are diluted.

By contrast, in the synthetic configuration, mBERT (91%) significantly outperforms Naïve Bayes (84%). This performance gap highlights how the noise-free, semantically aligned tweets in the oracle dataset allow transformer model to leverage its full potential. The performance gain achieved by mBERT also reflects the presence of fine-grained semantic cues that word-frequency based models cannot exploit. This validates the intended function of our synthetic dataset as an oracle environment—models capable of contextual understanding should perform well when noise is removed.

These findings also show a key insight: noise handling is more critical than model complexity for improving classification accuracy in weakly supervised text settings. Even the best models fail to learn effectively when irrelevant or ambiguous inputs dominate the data.

We also noted the cross-domain evaluation shows that models trained on the synthetic oracle dataset generalize poorly to noisy real-world tweets. This result underscores the significant impact of domain

shift introduced by the removal of noise by LLM-generated dataset. It is important to recognize that our synthetic dataset is not designed to replace real-world data. Rather, it serves a distinct purpose: to create a controlled, noise-free environment for systematic experimentation—particularly for studying the isolated effects of label or feature noise via noise injection strategies. We also acknowledge the potential risk of semantic drift in LLM-generated data—that is, the possibility that synthetic tweets may reflect biases, stereotypes, or linguistic abstractions learned by the model rather than replicating authentic user behavior. However, our goal is not to reproduce social media behavior with full fidelity. Instead, we intentionally trade off some realism for semantic clarity and control, enabling cleaner experimental analysis.

6. Conclusion

This study investigates the feasibility of using LLM to generate a synthetic oracle dataset for text classification tasks. We focus on the domain of building function classification (BFC) using geo-tagged tweets, a modality that offers semantic richness but suffers from label noise and, more critically, sentence-level feature noise—irrelevant, ambiguous, or uninformative tweets that obscure learning signals.

While label noise has been extensively studied through noise injection experiments, sentence-level feature noise remains underexplored due to the lack of a truly clean dataset. Human annotators often disagree on what constitutes relevance in social media text, making it difficult to filter out such noise consistently. To address this, we construct an oracle dataset using an LLM, where all tweets are guaranteed to be correctly labeled and semantically aligned with their associated building types. This provides a controlled, noise-free environment suitable for systematic experimentation.

Our evaluations show that:

- The synthetic dataset, while slightly less diverse in sentence structure (higher Self-BLEU), exhibits comparable lexical richness (similar perplexity) to real-world tweets.
- On real-world data, sophisticated models like mBERT perform only marginally better than simple classifiers like Naïve Bayes, due to noise degrading contextual learning.
- On synthetic data, mBERT significantly outperforms Naïve Bayes, validating that the synthetic dataset retains meaningful semantic distinctions and fulfills its oracle role.

These findings suggest that addressing feature noise may be more critical than increasing model complexity for tasks involving weakly supervised text. While our synthetic dataset is not intended to replicate real user behavior, it offers a valuable testbed for controlled experiments that would be infeasible with noisy real-world data.

7. Future Work

One key challenge in synthetic data generation is balancing diversity and semantic fidelity. Our findings show that while our dataset maintains realistic vocabulary usage, its sentence structure is less varied than real-world tweets. Future work could explore techniques such as controlled paraphrasing, persona variation, or prompt augmentation to increase structural diversity without compromising label alignment. Reducing this domain gap would make synthetic datasets more useful not only for experimentation but also for training and transfer learning.

Our synthetic oracle dataset provides a clean, well-defined starting point for studying the impact of noise in text classification. Future research can build on this by systematically injecting different types and levels of noise—such as label flipping, irrelevant content, or off-topic language—to isolate their individual and combined effects on model performance.

Given that our synthetic tweets are grounded in real cities and building names, future work could use this dataset to examine how geospatial context influences text classification. For example, studies could compare performance with and without location-specific references, or investigate how city-level

variation affects model generalization. This has implications not only for BFC, but also for broader geo-aware NLP tasks such as place-based sentiment analysis or urban event detection.

Declaration on Generative AI

During the preparation of this work, the author(s) used ChatGPT to identify and correct grammatical errors, typos, and other writing mistakes, in order to improve the clarity and professionalism of writing. After using this tool, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

References

- [1] A. Kruspe, M. Häberle, E. J. Hoffmann, S. Rode-Hasinger, K. Abdulahhad, X. X. Zhu, Changes in Twitter geolocations: Insights and suggestions for future usage, in: W. Xu, A. Ritter, T. Baldwin, A. Rahimi (Eds.), *Proceedings of the Seventh Workshop on Noisy User-generated Text (W-NUT 2021)*, Association for Computational Linguistics, Online, 2021, pp. 212–221. URL: <https://aclanthology.org/2021.wnut-1.24/>. doi:10.18653/v1/2021.wnut-1.24.
- [2] M. Häberle, M. Werner, X. X. Zhu, Building type classification from social media texts via geo-spatial textmining (2019) 10047–10050. URL: <https://ieeexplore.ieee.org/document/8898836>.
- [3] M. Häberle, M. Werner, X. X. Zhu, Geo-spatial text-mining from twitter – a feature space analysis with a view toward building classification in urban regions, *European Journal of Remote Sensing* 52 (2019) 2–11. URL: <https://doi.org/10.1080/22797254.2019.1586451>. doi:10.1080/22797254.2019.1586451.
- [4] M. Häberle, E. J. Hoffmann, X. X. Zhu, Can linguistic features extracted from geo-referenced tweets help building function classification in remote sensing?, *ISPRS Journal of Photogrammetry and Remote Sensing* 188 (2022) 255–268. URL: <https://www.sciencedirect.com/science/article/pii/S0924271622001058>. doi:<https://doi.org/10.1016/j.isprsjprs.2022.04.006>.
- [5] M. Häberle, Fusion of remote sensing images and social media text messages for building function classification (2022) 134. URL: <https://mediatum.ub.tum.de/1637448>.
- [6] B. Frenay, M. Verleysen, Classification in the presence of label noise: A survey, *IEEE Transactions on Neural Networks and Learning Systems* 25 (2014) 845–869. doi:10.1109/TNNLS.2013.2292894.
- [7] Z. Gu, Z. Chen, Y. Shi, X. Zhu, Building attributes recognition with noisy and incomplete labels, 2024, pp. 230–233. doi:10.1109/IGARSS53475.2024.10641987, publisher Copyright: © 2024 IEEE.; 2024 IEEE International Geoscience and Remote Sensing Symposium, IGARSS 2024 ; Conference date: 07-07-2024 Through 12-07-2024.
- [8] R. Lamsal, A. Harwood, M. R. Read, Where did you tweet from? inferring the origin locations of tweets based on contextual information, in: *2022 IEEE International Conference on Big Data (Big Data)*, IEEE, 2022, p. 3935–3944. URL: <http://dx.doi.org/10.1109/BigData55660.2022.10020460>. doi:10.1109/bigdata55660.2022.10020460.
- [9] A. Kruspe, J. Kersten, F. Klan, Review article: Detection of actionable tweets in crisis events, *Natural Hazards and Earth System Sciences* 21 (2021) 1825–1845. URL: <https://nhess.copernicus.org/articles/21/1825/2021/>. doi:10.5194/nhess-21-1825-2021.
- [10] G. Delétang, A. Ruoss, P.-A. Duquenne, E. Catt, T. Genewein, C. Mattern, J. Grau-Moya, L. K. Wenliang, M. Aitchison, L. Orseau, M. Hutter, J. Veness, Language modeling is compression, 2024. URL: <https://arxiv.org/abs/2309.10668>. arXiv:2309.10668.
- [11] T. Ge, J. Hu, L. Wang, X. Wang, S.-Q. Chen, F. Wei, In-context autoencoder for context compression in a large language model, 2024. URL: <https://arxiv.org/abs/2307.06945>. arXiv:2307.06945.
- [12] T. Schick, H. Schütze, Generating datasets with pretrained language models, in: M.-F. Moens, X. Huang, L. Specia, S. W.-t. Yih (Eds.), *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Online and Punta Cana,

- Dominican Republic, 2021, pp. 6943–6951. URL: <https://aclanthology.org/2021.emnlp-main.555/>. doi:10.18653/v1/2021.emnlp-main.555.
- [13] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, D. Amodei, Language models are few-shot learners, 2020. URL: <https://arxiv.org/abs/2005.14165>. arXiv:2005.14165.
 - [14] R. Bommasani, D. A. Hudson, E. Adeli, R. Altman, S. Arora, S. von Arx, M. S. Bernstein, J. Bohg, A. Bosselut, E. Brunskill, E. Brynjolfsson, S. Buch, D. Card, R. Castellon, N. Chatterji, A. Chen, K. Creel, J. Q. Davis, D. Demszky, C. Donahue, M. Doumbouya, E. Durmus, S. Ermon, J. Etchemendy, K. Ethayarajh, L. Fei-Fei, C. Finn, T. Gale, L. Gillespie, K. Goel, N. Goodman, S. Grossman, N. Guha, T. Hashimoto, P. Henderson, J. Hewitt, D. E. Ho, J. Hong, K. Hsu, J. Huang, T. Icard, S. Jain, D. Jurafsky, P. Kalluri, S. Karamcheti, G. Keeling, F. Khani, O. Khattab, P. W. Koh, M. Krass, R. Krishna, R. Kudritipudi, A. Kumar, F. Ladhak, M. Lee, T. Lee, J. Leskovec, I. Levent, X. L. Li, X. Li, T. Ma, A. Malik, C. D. Manning, S. Mirchandani, E. Mitchell, Z. Munyikwa, S. Nair, A. Narayan, D. Narayanan, B. Newman, A. Nie, J. C. Niebles, H. Nilforoshan, J. Nyarko, G. Ogut, L. Orr, I. Papadimitriou, J. S. Park, C. Piech, E. Portelance, C. Potts, A. Raghunathan, R. Reich, H. Ren, F. Rong, Y. Roohani, C. Ruiz, J. Ryan, C. Ré, D. Sadigh, S. Sagawa, K. Santhanam, A. Shih, K. Srinivasan, A. Tamkin, R. Taori, A. W. Thomas, F. Tramèr, R. E. Wang, W. Wang, B. Wu, J. Wu, Y. Wu, S. M. Xie, M. Yasunaga, J. You, M. Zaharia, M. Zhang, T. Zhang, X. Zhang, Y. Zhang, L. Zheng, K. Zhou, P. Liang, On the opportunities and risks of foundation models, 2022. URL: <https://arxiv.org/abs/2108.07258>. arXiv:2108.07258.
 - [15] T. Schick, H. Schütze, Generating datasets with pretrained language models, 2021. URL: <https://arxiv.org/abs/2104.07540>. arXiv:2104.07540.
 - [16] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. L. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, J. Schulman, J. Hilton, F. Kelton, L. Miller, M. Simens, A. Askell, P. Welinder, P. Christiano, J. Leike, R. Lowe, Training language models to follow instructions with human feedback, 2022. URL: <https://arxiv.org/abs/2203.02155>. arXiv:2203.02155.
 - [17] M. Shanahan, K. McDonell, L. Reynolds, Role-play with large language models, 2023. URL: <https://arxiv.org/abs/2305.16367>. arXiv:2305.16367.
 - [18] J. Li, N. Mehrabi, C. Peris, P. Goyal, K.-W. Chang, A. Galstyan, R. Zemel, R. Gupta, On the steerability of large language models toward data-driven personas, 2024. URL: <https://arxiv.org/abs/2311.04978>. arXiv:2311.04978.
 - [19] T. Ge, X. Chan, X. Wang, D. Yu, H. Mi, D. Yu, Scaling synthetic data creation with 1,000,000,000 personas, 2024. URL: <https://arxiv.org/abs/2406.20094>. arXiv:2406.20094.
 - [20] H. K. Choi, Y. Li, Picle: Eliciting diverse behaviors from large language models with persona in-context learning, 2024. URL: <https://arxiv.org/abs/2405.02501>. arXiv:2405.02501.
 - [21] Y. Zhu, S. Lu, L. Zheng, J. Guo, W. Zhang, J. Wang, Y. Yu, Taxygen: A benchmarking platform for text generation models, 2018. URL: <https://arxiv.org/abs/1802.01886>. arXiv:1802.01886.
 - [22] J. Ye, J. Gao, Q. Li, H. Xu, J. Feng, Z. Wu, T. Yu, L. Kong, Zerogen: Efficient zero-shot learning via dataset generation, 2022. URL: <https://arxiv.org/abs/2202.07922>. arXiv:2202.07922.
 - [23] D. Jurafsky, Speech and Language Processing, Stanford University, 2025.