TailNLG: A Multilingual Benchmark Addressing Verbalization of Long-Tail Entities

Anonymous ACL submission

Abstract

The automatic verbalization of structured knowledge is a key task for making knowledge graphs accessible to non-expert users and for supporting retrieval-augmented generation systems. Although recent advances in RDF-to-text generation have improved multilingual coverage, little attention has been paid to potential biases in the verbalization of rare entities, frequently known as long-tail entities. In this work, we present the first systematic study of long-tail entities in RDF-to-text generation. We introduce TailNLG[1], a new multilingual benchmark in English, Italian, and Spanish, built from Wikidata and covering entities with varying levels of popularity. We evaluate three different families of large language models in zero-shot settings and compare their performance on rare *versus* common entities, as well as against the established WebNLG benchmark. Our results reveal a consistent bias against long-tail entities: embedding-based scores are lower, and model uncertainty is higher for rare entities. We further show that the impact of long-tail entities varies across models and languages, and that existing evaluation metrics do not consistently capture these differences, highlighting the need for more reliable evaluation frameworks.

1 Introduction

The automatic conversion of structured information into human-readable text is a crucial task for making knowledge accessible to a broad audience. Knowledge graphs (KGs), catalogues, and taxonomies are important sources of information that are typically exploitable only by domain experts, whereas the general public requires this knowledge to be transformed into natural language. This need also extends to language modeling, as retrieval-augmented generation (RAG) technologies leverage textual databases to improve the safety and accountability of large language models (LLMs).

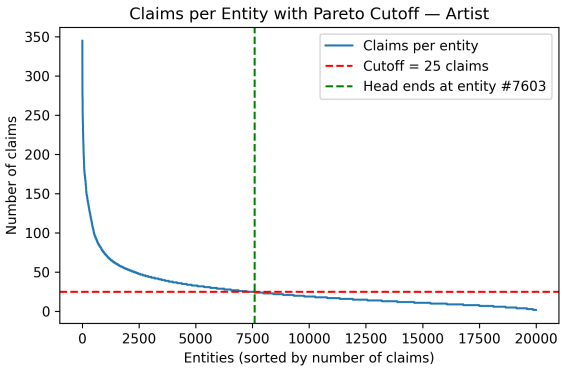

Figure 1: Distribution of entities by claims in a Wikidata category ("Artist"). A small number of entities are associated with many claims, giving rise to a long-tail distribution characterized by many entities with few relations. The red dashed line indicates the Pareto cut-off, while the green dashed line marks the head–tail threshold

Despite advancements in this field and the growing number of multilingual resources (Ferreira et al., 2018; Shimorina et al., 2019; Cripwell et al., 2023; Oliverio et al., 2025; Ramón-Ferrer et al., 2025), there is a lack of work aimed at uncovering biases in the verbalization of structured knowledge (Blodgett et al., 2020). Specifically, an open challenge is to gain a deeper understanding of how LLMs handle rare entities, frequently known as *long-tail* entities, and whether their performance differs when compared to well-known entities. While it has been demonstrated that LLMs exhibit performance drops when handling rare entities in tasks such as entity linking (Boscariol et al., 2024), representation of factual knowledge (Sun et al., 2024), and question answering (Hogan et al., 2025), no prior work has conducted an in-depth investigation of how different models handle long-tail

[1]Full code and Benchmark available here https://anonymous.4open.science/r/TailNLG-benchmark-B339/README.md

entities in RDF-to-text generation.

In this work, we tackle this gap by presenting the first systematic study of long-tail entities in RDF-to-text verbalization. Acknowledging the socio-cultural factors that can determine the rarity of an entity (Adams et al., 2019; Stranisci et al., 2023), we adopt a multilingual approach. Our study introduces **TailNLG**, a novel multilingual RDF-to-text benchmark covering three languages (English, Italian, and Spanish). The dataset contains triples extracted from Wikidata (Vrandečić and Krötzsch, 2014), including entities from different categories and spanning varying levels of popularity, from top-head (most well-known) to long-tail (less well-known) entities.

Through the creation of TailNLG, our work addresses the following research questions:

- **[RQ1]** Do LLMs perform differently in the verbalization of rare entities, and does language impact the verbalization output?
- **[RQ2]** What is the impact of LLMs' prior knowledge on the verbalization of rare *versus* common entities, and how does model performance differ when comparing a long-tail benchmark (TailNLG) with a well-established resource such as WebNLG?

To this end, three families of LLMs are assessed on the verbalization of long-tail entities in two settings: (i) comparing them with head entities within the TailNLG corpus, and (ii) comparing them with WebNLG (Gardent et al., 2017b), the established benchmark for the RDF-to-text task. We assess performance with a diverse set of automatic measures, capturing semantic similarity, surface-level overlap with reference texts, and perplexity to quantify model uncertainty. Our results show that LLMs have a systematic bias against long-tail entities in both evaluation settings. The embedding-based metrics tend to be lower for long-tail entities and models' uncertainty in text generation of triple containing long-tail entities is almost always higher.

In addition, language-specific bias plays also a role in the verbalization of long-tail entities. Analysis of perplexity scores reveals that embedding-based and overlap-based metrics fail to systematically quantify their impact on verbalization, demonstrating that a more reliable evaluation framework of fairness in RDF-to-Text is needed. Finally, we observe notable variability across models, indicating that different LLMs may be affected by different biases.

2 Related Work

2.1 Long-tail

Long-tail entities are generally understood as those that occur with low frequency across large scale data sources, including training corpora (Kandpal et al., 2023), Wikipedia (Mallen et al., 2023), and large knowledge bases such as Wikidata (Kumar et al., 2024). Recent work has shown that LLMs exhibit a marked decline in performance when processing long tail inputs (Graciotti et al., 2025; Hogan et al., 2025; Li et al., 2024; Sun et al., 2024) and that LLMs memorization is highly influenced by the frequency of information in the pretraining data (Kandpal et al., 2023). Several task-specific evaluations further confirm this pattern. Boscariol et al. (2024) find that entity linking is particularly difficult for rare entities when comparing various LLMs with ReLiK (Orlando et al., 2024), a state-of-the-art (SoTA) entity linking and relation extraction system. GRADES, an evaluation framework for graph-based question answering (Draetta et al., 2025), shows that SoTA models struggle with rare entities at multiple stages; in particular, verbalization suffers from limited semantic understanding of long-tail entities. In knowledge extraction, Graciotti et al. (2025) introduce KEMISTO, a benchmark centered on low-popularity entities, and demonstrate that LLMs perform significantly worse on it and exhibit systematic biased failures. A similar pattern emerges in question answering: using a benchmark spanning head, torso, and tail entities, Sun et al. (2024) show that accuracy consistently declines as entity frequency decreases, confirming that limited training exposure hampers LLMs' knowledge of long-tail entities.

Despite increasing attention to long-tail evaluation, and some insight emerged from previous studies no existing work specifically examines how well LLMs verbalize information when required to handle rare entities.

2.2 RDF-to-Text

The task of verbalizing RDF triples in NLP has evolved substantially, driven both by template-based NLG systems and more recent neural approaches. Early work focused on hand-crafted or automatically induced templates, which provided high precision and strong control over the output,

but suffered from limited scalability and reduced the ability to generalize beyond predefined patterns (Duma and Klein, 2013).

With the advent of neural encoder–decoder architectures, particularly LSTM-based triple encoders and later transformer models, RDF-to-text generation became considerably more flexible, enabling fluent and information-rich verbalizations that better capture lexical and syntactic variation (Distiawan et al., 2018; Oliverio et al., 2024). Nevertheless, challenges related to factual consistency and robustness to unseen entities have motivated a series of enhancements to transformer-based systems. For example, integrating reward signals derived from information-extraction models has been shown to improve factual accuracy (Gao et al., 2021), while leveraging large external corpora, pretraining noise processes, and data augmentation strategies leads to more resilient generalization, particularly in low-resource or zero-shot scenarios (Montella et al., 2020; Zhang et al., 2023).

To support systematic development and evaluation of this task, the WebNLG benchmark has emerged as one of the primary datasets for RDF-to-text, providing multiple versions, multilingual extensions, and low-resource configurations that have enabled extensive comparative studies across approaches and languages (Gardent et al., 2017b; Castro Ferreira et al., 2020).

3 The TailNLG Benchmark

Since no benchmark currently targets long-tail triple–verbalization pairs, we introduce the first dataset specifically designed for long-tail verbalization. Inspired by Gardent et al. (2017b), we created the TailNLG benchmark following the same procedure the authors employed, from the entities extraction strategy to the verbalization, aiming at facilitating the exploitation of the data adhering to a standard structure. Constructing the benchmark in three languages adds another layer of complexity, as many rare entities are not consistently represented across multilingual resources. We prioritized high data quality by incorporating human evaluation and correction at every stage of the pipeline.

3.1 Entities Selection

To be consistent with previous resources, we followed the extraction methodology proposed by Gardent et al. (2017a) by selecting the same entity categories. Although there is not yet a widely accepted and unambiguous definition of long-tail entities (Jiang and Joshi, 2024; Boscariol et al., 2024; Kumar et al., 2024), following Sun et al. (2024) we took into consideration two criteria in defining long-tail entities: (i) the number of claims (Wikidata relations) linked to a certain entity and (ii) the number of Wikipedia pages an entity has.

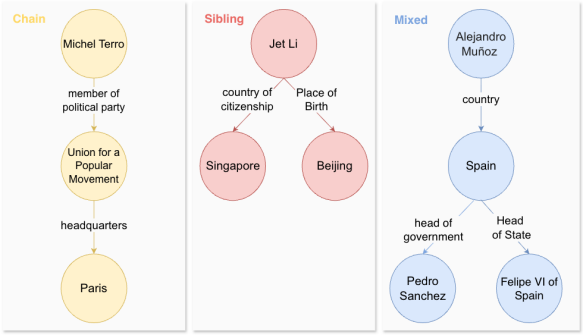

Figure 2: Examples of Chain, Sibling and Mixed configuration of Data units.

Firstly, we set a threshold based on the number of claims linked to each entity. As noted by Kumar et al. (2024), large-scale resources such as Wikidata typically follow a Zipf-like distribution, where a small number of entities are associated with many claims while the vast majority with few. Based on a subset of $20,000$ entities for each category (full list of categories in Table 4 in Appendix B) we define the long-tail threshold using a Pareto cut-off: we compute the minimum number of claims required for an entity to fall within the top 20% of entities that account for 80% of all claims (the head). Entities below this threshold are classified as long-tail. Since claim distributions differ across categories, we compute the cut-off separately for each category.

Figure 1 shows an example of the threshold for the categories Artists: a small number of entities exhibit a very high number of claims, while the vast majority fall into a long-tail distribution.

To further refine the set of long-tail entities below the Pareto cut-off, we additionally select those that have a Wikipedia page only in one of the three target languages (e.g., entities with wikipedia page in English but not in Italian and Spanish).

In summary, we categorize entities into: (i) **Long-tail entities**, whose number of claims falls below the Pareto cut-off, such as *Doğan Güzel* (Q268285), a Turkish cartoonist; and (ii) **Head entities**, whose number of claims exceeds the Pareto cut-off, such as *London* (Q84) or *Google* (Q95).

3.2 Triples Extraction

Once the sets of long-tail and head entities were defined, the next step was to extract data units for each entity. We define a data unit as a set of RDF triples, each consisting of a subject, a predicate, and an object (e.g., Jet Li, place of birth, Beijing). To obtain diverse data units that vary in relation type and size, we followed Perez-Beltrachini et al. (2016) and extracted units with different relational configurations (Figure 2):

- **Chain**: the object of one triple serves as the subject of another;
- **Sibling**: the triples share the same subject;
- **Mixed**: the data unit contains both sibling and chain relations.

For each QID (Wikidata entity ID), we automatically extracted from Wikidata, via a SPARQL query, a set of 10 data units per entity in the three languages. Since not all triples were available in all three languages, we further filtered the data units to ensure the alignment across languages, resulting in a set of parallel data units as exemplified below:

EN: `[Solar, creator, Paul S. Newman], [Paul S. Newman, place of birth, Manhattan], [Paul S. Newman, award received, Inkpot Award]`

ES: `[Solar, creator, Paul S. Newman], [Paul S. Newman, lugar de nacimiento, Manhattan], [Paul S. Newman, premio recibido, Premio Inkpot]`

IT: `[Solar, creatore, Paul S. Newman], [Paul S. Newman, luogo di nascita, Manhattan], [Paul S. Newman, premio ricevuto, Premio Inkpot]`

We finally obtain 689 data units per language, divided into 408 long-tail and 218 head entities. The complete distribution per category and type is in Appendix B.

3.3 Verbalization

To build the benchmark, we aimed to produce *<Data_Unit, Verbalization>* pairs for all data in all languages. Because manual verbalization is time consuming we adopted a two step approach: we first manually verbalized a subset of data units for each language, then cross-translated them into the remaining languages using SoTA machine translation models in a human-in-the-loop setup.

The verbalization process was carried out on a subset of 689 Data units engaging nine volunteers native speakers (one for Spanish, one for English, and seven for Italian). Annotators were instructed to verbalize the triples as naturally as possible, producing grammatically correct and fluent sentences while following guidelines designed to promote naturalness and ensure data homogeneity. All guidelines and some example of verbalization are reported in Appendix A.1.

Annotators could skip a data unit if they judged it incorrect or incomplete. After verbalization, a validation phase followed in which the annotators cross-checked each other's output, confirming or revising the initial sentences (full guidelines in Appendix A.2). Ambiguous cases were discussed collectively before a final decision was made. Following this verbalization and cross-check process, 73 data units were excluded from the dataset.

3.4 Translation

The manual verbalization phase produced a total of 616 gold verbalizations, distributed as 130 in English, 142 in Spanish, and 344 in Italian[2]. To create a fully parallel dataset we adopted an approach based on Machine Translation (MT) and Automatic Post-Editing (APE), as proposed by Aditya Hari et al. (2023).

Drawing on earlier studies on cross-lingual extensions of triples-to-text datasets (Oliverio et al., 2025; Ramón-Ferrer et al., 2025), we leveraged DeepL for Automatic Translation[3]. For each manually created verbalization in a source language (X) corresponding translations were automatically generated in the other two target languages (Y and Z), resulting in a total of 1,232 automatic generated translations.

To mitigate known issues in MT, such as omissions and hallucinations (Jurafsky and Martin, 2025), and to ensure data quality, we introduced an APE phase aimed at correcting errors generated during the MT step. To identify errors in the translations, we leveraged the 3.5B version of xCOMET Guerreiro et al. (2024) [4], a multilingual model for quality estimation and error span extraction. Given the absence of reference translations, we evaluated the model in a reference-free setting. For each input, the model outputs a sentence-level score (ranging from *Excellent* to *Weak*), as well as an error spans detected field, which identifies potential er-

[2]This linguistic imbalance is due to the varying availability of expert native-speaking annotators.

[3]https://www.deepl.com/translator

[4]https://huggingface.com/Unbabel/XCOMET-XL

rors and their positions within the sentence.

For the error correction phase, we adopted xTOWER[5], a multilingual decoder-only model proposed by Treviso et al. (2024), designed to provide natural-language explanations of errors present in a translation and suggest a corrected version of the sentence. An example of the prompt and the corresponding output are provided in the Appendix C.1. We performe error detection and correction on all the MT translated data; the resulting output is a parallel dataset consisting of data units and verbalizations in Italian, Spanish and English as follows:

```
<lex quality="gold" lid="Id1" lang="en">
    Themisthocles was born in Turkey and
    he was a citizen of the Classical Athens.
    Themistocles's spouse is Archippe.
</lex>
<lex quality="silver" lid="Id2" lang="es">
    Temístocles nació en Turquía y fue
    ciudadano de la Atenas clásica.
    El cónyuge de Temístocles es Arquippe.
</lex>
<lex quality="silver" lid="Id3" lang="it">
    Temistocle è nato in Turchia ed era un
    cittadino dell'Atene classica.
    La consorte di Temistocle era Archippe.
</lex>
```

To ensure the high quality of the dataset, we conducted a manual evaluation on a subset, selecting those translations whose xCOMET *Sentence-level score* was rated below *good* (i.e., *moderate* or *weak*). We sampled 50 verbalizations for each source-target language pair and involved eight volunteer native speakers (two with high competence in Spanish and English, two in Italian and Spanish, and two in Italian and English). Each source-target language pair was annotated by two different annotators following a widely adopted taxonomy for the evaluation of Accuracy and Fluency (Lommel et al., 2014) (full evaluation guidelines in Appendix C.2). In addition, we asked the annotators to propose a corrected version of the automatically translated verbalization to increase the number of gold references in TailNLG. We use Spearman (Spearman, 1961) and Persons (Pearson, 1895) to measure the Inter annotator agreement (IAA) after aggregating Accuracy and Fluency. The results for each language pair are reported in Appendix C.3. Although inter-annotator agreement (IAA) shows a moderate positive correlation overall, agreement is low for IT→EN and ES→IT pair. This indicates a high degree of subjectivity in the task, with annotators showing divergent interpretations. These findings further underline that even tasks such as machine translation can reflect differences in human judgment. The final verbalization distribution is available in Table 1.

Language	Gold	Silver	Total
English	168	447	615
Spanish	173	442	615
Italian	384	231	615
Total	725	1,120	1,845

Table 1: Distribution of gold and silver verbalizations by language in TailNLG.

4 Methodology

To study multilingual triple-to-text verbalization, we leverage three families of open state-of-the-art multilingual instruction-tuned models, each evaluated at two comparable sizes: Llama-3.2-3B-Instruct [6], Llama-3.1-8B-Instruct[7], Qwen2.5-3B-Instruct[8], Qwen2.5-7B-Instruct[9], Gemma-3-4B-IT[10], and Gemma-3-12B-IT[11]. We focus on these model families because they are multilingual, widely used, and well established in the literature, openly available, and offered in moderate parameter ranges that enable efficient and reproducible experimentation.

4.1 Experimental Setup

Given a set of RDF-style triples (s, p, o), the system produces a single-paragraph natural language description in English, Spanish, or Italian. Inputs are serialized by explicitly listing each triple as a structured item (subject, predicate, object) in the prompt. We run experiments on two datasets: (i) WebNLG, used as a baseline benchmark, and (ii) TailNLG, used to analyze performance on data unit containing long-tail entities. Since our goal is to study biases and examine how LLMs handle rare entities, we deliberately adopt a zero-shot approach for both the configurations to avoid influencing

[5]https://huggingface.com/sardinelab/xTower13B

[6]https://huggingface.co/meta-llama/Llama-3.2-3B-Instruct

[7]https://huggingface.co/meta-llama/Llama-3.1-8B-Instruct

[8]https://huggingface.co/Qwen/Qwen2.5-3B-Instruct

[9]https://huggingface.co/Qwen/Qwen2.5-7B-Instruct

[10]https://huggingface.co/google/gemma-3-4b-it

[11]https://huggingface.co/google/gemma-3-12b-it

model outputs with external examples. Prompts are written English, Spanish, or Italian depending on the target verbalization. Full prompts are available in Appendix E.

For each input we generate three candidate verbalizations to capture decoding variability. We use stochastic decoding with temperature sampling (temperature $= 0.7$), which increases output diversity compared to greedy decoding while keeping generations coherent. Generations are bounded by a fixed maximum length (max_new_tokens $= 256$) to prevent overly long outputs and to standardize comparisons across models and settings. All outputs are stored together with the corresponding prompts and metadata. Experiments were conducted on a server with two Intel Xeon Gold 5418Y CPUs (48 physical cores each, 96 threads total), around 504 GB RAM, and two NVIDIA L40 GPUs with approximately 46 GB VRAM each.

4.2 Evaluation

In evaluating the results, we compare multiple metrics designed to capture three characteristics of the generated output: (i) semantic similarity (embedding-based metrics), (ii) textual similarity (overlap-based metrics), and (iii) model uncertainty in text generation[12]. We use BERTScore (Zhang et al., 2019) to evaluate semantic similarity, computing pairwise cosine similarity between candidate and reference texts with contextual embeddings. Specifically, we adopt the rescaled-with-baseline variant (BERTScore$_r$) for more interpretable scores, using a multilingual BERT model[13] to ensure comparability across languages. To evaluate textual similarity, we adopt: BLEU (Papineni et al., 2002), chrF++ (Popović, 2017), and three variants of ROUGE (Lin, 2004), namely ROUGE$_1$, which measures unigram (word-level) overlap; ROUGE$_2$, which measures bigram overlap; and ROUGE$_L$, which captures sentence-level structural similarity by identifying the longest common subsequences. To quantify models uncertainty in triple verbalization, we computed the average Perplexity (PPL) (Jelinek et al., 1977) interpreting lower perplexity values as a proxy of the higher confidence of a model in text generation.

To assess the statistical significance of each metric, we applied the Wilcoxon–Mann–Whitney test (Fagerland and Sandvik, 2009) on triple verbalization of the same model on different populations. We repeated the test ten times on random samples of 500 predictions and averaged the results. We performed the test in two settings: long-tail *versus* head verbalization in the TailNLG corpus and long-tail verbalization in the TailNLG corpus *versus* WebNLG verbalization.

[12]Evaluations are performed using Hugging Face packages

[13]https://huggingface.co/google-bert/bert-base-multilingual-cased

5 Results

In this section, we present results comparing long-tail and head entities, as well as across languages (RQ1). We also examine the impact of LLMs' prior knowledge on verbalizing rare and well-known entities by comparing PPL scores on TailNLG and WebNLG, assessing whether models show greater uncertainty for rare entities (RQ2). Following (Du et al., 2024), we define prior knowledge as the information acquired during pretraining.

5.1 Do LLMs perform differently in the verbalization of rare entities, and does language impact the verbalization output?

Table 2 reports the metric scores for triples containing long-tail and head entities, broken down by language (English, Italian and Spanish), for the bigger size models. Complete results are provided in Appendix F. In all cases, the embedding-based metric BERTScore$_r$, scores higher on head triples. Conversely, results from chrF++, which correlates with morphological accuracy, and ROUGE are less consistent: in some cases they score higher on long-tail triples, while in others on head triples. In addition, when the results for BERTScore$_r$ are statistically significant, the scores are consistently higher for head entities than for long-tail entities, which may be linked to greater semantic awareness when verbalizing triples involving head entities. In contrast, the overlap-based metrics exhibit the opposite trend: when the results are statistically significant, higher scores are observed for triples containing long-tail entities. The average PPL provides clear insights, particularly for Qwen and Llama models. The results suggest that verbalizing triples with long-tail entities leads to higher PPL, indicating that those entities constitute a substantially more challenging prediction setting for these models. In contrast, the PPL scores for Gemma models are less clear and always lower for long-tail entities. Investigating whether language influences the generated output, we observe that across all settings performance is higer for English, for both data units

Type	Lang	BERT$_r$ ↑	BLEU ↑	chrF ↑	R-1 ↑	R-2 ↑	R-L ↑	PPL ↓
Llama-8B								
Long-tail	en	0.61	0.25	60.39	0.66	0.43	0.56	93.36
	es	0.58	0.17	53.30	0.58	0.35	0.49	222.16
	it	0.46	0.13	47.73*	0.49*	0.29*	0.43*	88.06
Head	en	0.65*	0.26	61.20	0.67	0.44	0.57	59.03*
	es	0.60	0.19	53.23	0.58	0.36	0.48	94.52*
	it	0.45	0.13	45.58	0.46	0.27	0.40	68.84*
Gemma-12B								
Long-tail	en	0.63	0.26	61.65	0.67	0.44	0.58	330.02*
	es	0.62	0.21	58.80	0.62	0.40	0.53	611.42*
	it	0.61	0.24*	60.68*	0.64*	0.42*	0.55*	425.27*
Head	en	0.66*	0.27	61.72	0.68	0.45	0.57	397.46
	es	0.64*	0.22	57.51	0.61	0.39	0.51	926.65
	it	0.61	0.20	56.97	0.59	0.38	0.50	1068.31
Qwen-7B								
Long-tail	en	0.65	0.29	62.47	0.69	0.46	0.58	134.13
	es	0.60	0.16	53.42	0.57	0.34	0.49	181.75
	it	0.44	0.11	46.82	0.47	0.26	0.41	100.26
Head	en	0.69*	0.29	63.06	0.70	0.47	0.59	98.05*
	es	0.61	0.16	51.83	0.55	0.32	0.46	102.36*
	it	0.46	0.10	45.35	0.45	0.25	0.39	63.48*

Table 2: Results on TailNLG long-tail vs head entitities verbalizations for Llama-3.1-8B, Gemma2.5-12B 12B, Qwen2.5-7B by data type and language. Metrics include BERTScore Rescaled (BERT$_r$), BLEU, chrF++ (chrF), ROUGE$_1$ (R-1) , ROUGE$_2$ (R2), ROUGE$_L$ (R-L) and perplexity (PPL). Asterisks (*) denote a statistically significant difference ($p < 0.05$).

containing long-tail and head entities and consistently lower for Italian. This suggests that English content is handled more effectively by the models. In addition, for Llama and Qwen, average PPL scores for Italian are consistently higher than those for English and Spanish, indicating that the models could be affected by language bias and that outputs in more widely represented languages tend to be more accurate.

5.2 [RQ2] How does model performance differ when comparing the TailNLG benchmark with WebNLG?

Table 3 reports the average performance of each model on the WebNLG test set and on long-tail triples from TailNLG. All statistically significant results indicate that embedding- and overlap-based metrics achieve higher scores on WebNLG than in the long-tail setting. This trend suggests that models generate more accurate and lexically aligned verbalizations for well-known entities, whereas performance tends to degrade for long-tail entities.

Across all models, PPL is even lower for WebNLG verbalizations, with statistical significance, supporting the hypothesis that WebNLG entities and triples are more familiar to the models, potentially due to exposure during pre-training.

Overall, these results highlight the impact of LLMs' prior knowledge on the verbalization task. When entities and triples are well represented in the models' pre-training data, LLMs tend to produce verbalizations that are more closely aligned with the ground truth. Conversely, for long-tail entities, limited prior exposure leads to a degradation in both surface-level quality and model confidence, suggesting that prior knowledge plays a crucial role in the effective verbalization of structured data.

In addition, since models tend to achieve higher scores when synthetic data are used as ground truth

Type	BERT$_r$ ↑	BLEU ↑	chrF ↑	R-1 ↑	R-2 ↑	R-L ↑	PPL ↓
Llama-8B							
Long-tail	0.55	0.18	53.81	0.58	0.36	0.50	134.53
WebNLG	0.53	0.21*	53.77	0.65*	0.43*	0.53*	68.23*
Gemma-12B							
Long-tail	0.62	0.24	60.38	0.64	0.42	0.55	455.57
WebNLG	0.65*	0.30*	63.63*	0.73*	0.50*	0.59*	290.10*
Qwen-7B							
Long-tail	0.56	0.19	54.24	0.58	0.35	0.49	138.71
WebNLG	0.57	0.24*	56.49	0.66*	0.44*	0.54*	95.54*

Table 3: Performance comparison of models between long-tail triples from TailNLG and the WebNLG test set. Metrics include BERTScore Rescaled (BERT$_r$), BLEU, chrF++ (chrF), ROUGE-1 (R-1), ROUGE-2 (R-2), ROUGE-L (R-L), and perplexity (PPL). Asterisks (*) denote a statistically significant difference ($p < 0.05$).

(Sarti and Nissim, 2024), our findings are further reinforced. Even when comparing model performance on the English WebNLG test set (entirely gold) with the long-tail portion of TailNLG (containing both gold and silver references), models still achieve higher scores on WebNLG. For further details, we refer the reader to Appendix F.

6 Discussion

Results show that, across all models, verbalization of triples that contains long-tail entities is more challenging (RQ1). This suggests that LLMs suffer from entity bias, with widely represented entities being processed more easily and producing more accurate outputs. The fact that in some cases overlap-based metrics are higher for data units containing long-tail entities may suggest that the verbalization of head entities contains more data and key content relative to the reference, and that prior knowledge could play a role in shaping the generated output adding information. For example, the Italian verbalization of the data unit `Karl Marx | spouse | Jenny von Westphalen; Karl Marx | country of citizenship | statelessness` is "Karl Marx è stato apatico alla cittadinanza, preferendo lo status di apolidia. Karl Marx ebbe come consorte Jenny von Westphalen" (English translation: Karl Marx was apathetic toward citizenship, preferring to remain stateless. Karl Marx's spouse was Jenny von Westphalen), which results in a more verbose output that differs from the target sentence. The results concerning language strongly confirm the hypothesis that models also suffer from language bias, and that verbalizations in English are more accurate. These results may also be influenced by the fact that both Italian and Spanish have richer morphologies than English. For example, the Italian verbalization of the triples `Amanda Ammann | country of citizenship | Switzerland; Amanda Ammann | occupation | model` results in "Amanda Ammann è una cittadina svizzera e lavora come modello" (English translation: Amanda Ammann is a Swiss citizen and works as a model), which exhibits a gender disagreement between the first part of the sentence ("cittadina svizzera") and the final noun ("modello"). Generally Italian output tend to be less accurate and suffer from some additions and omissions.

Another important finding emerging from the results is the high variability across evaluation metrics. In particular, reference-based metrics are not always able to detect the presence of bias or provide reliable scores, as becomes evident from the comparison with PPL, which yields more consistent results. In addition, the analysis suggests that different models may be affected by different types of bias; in particular, Gemma exhibits a higher PPL for triples containing long-tail entities.

The prior-knowledge experiment confirms our hypothesis that models perform better on a benchmark such as WebNLG compared to TailNLG (RQ2), as the models may have been exposed during pre-training to information related to the WebNLG benchmark or to the entities it contains. This finding is supported, with statistical significance, by both reference-based evaluation metrics and PPL.

7 Conclusion and Future Works

In this work, we investigated the performance of three state-of-the-art language models' families in depth for verbalizing long-tail entities in a multilingual setting. To this end, we introduced TailNLG, the first benchmark composed of triple–verbalization pairs covering entities that range from rare to well-represented. Our results show that the leveraged models suffer from both language and entity biases, and that they perform better on a widely exploited benchmark such as WebNLG compared to TailNLG. Consistent with prior work, our findings highlight that current models struggle to handle long-tail entities, revealing limitations that make tasks involving less-represented entities more difficult and less reliable.

As future work, we plan to first enrich the TailNLG benchmark by adding more languages and data units. To further refine the dataset, we also aim to distinguish entity types not only based on their popularity (i.e., number of claims) but also by incorporating culturally specific data units to explore the models' cultural bias. Finally, we plan to enrich the dataset with a gold standard that takes into account the diverse perspectives of annotators, enabling the investigation of human label variation in the verbalization task.

Limitations

We acknowledge that building a new high-quality benchmark is challenging, as it requires achieving an appropriate balance between human annotation and automatic processing steps. In this work, we consistently paired the automatic phase with hu-

man assessment in an effort to ensure the highest possible quality. Nevertheless, the presence of *silver* verbalizations in the benchmark constitutes a limitation, since machine translation tools, although reliable, may not fully align with human judgments.

Second, by leveraging a well-known resource such as Wikidata, the variability of available claims is limited. As a result, models may already be familiar with many of the relations, which also recur across different data units.

Finally, in defining long-tail entities, we adopted a simplified approach for experimental purposes, considering only a quantitative parameter (number of claims). However, we are aware that entity underrepresentation is a multifaceted issue that involves multiple dimensions, both vertical (numerosity) and horizontal (cultural aspects). In future work, it would be valuable to consider additional perspectives and to define long-tail entities using a more comprehensive characterization.

Ethical considerations

Our work relies on knowledge extracted from Wikidata, whose community of contributors has recognized the underrepresentation of non-Western and woman contributors. This has an impact on how knowledge is shaped and which entities suffer underrepresentation. Adopting a purely statistical approach to the definition of *long-tail* entities might have resulted in overlooking these cultural factors. Nevertheless, this approach was preferred to the adoption of theoretically grounded models of underrepresentation, as the latter could have introduced additional research bias in corpus creation.

All annotations have been performed voluntarily by the members of our research group without relying on external annotation platforms.

During the preparation of this work, the authors used AI tools in order to: Grammar and spelling check. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

Acknowledgments

We thank all the individuals who contributed to this project by supporting the verbalization phase and the evaluation of translations. Without their help, we would not have been able to achieve the desired level of data quality.

References

Julia Adams, Hannah Brückner, and Cambria Naslund. 2019. Who counts as a notable sociologist on wikipedia? gender, race, and the "professor test". *Socius*, 5:2378023118823946.

Kancharla Aditya Hari, Bhavyajeet Singh, Anubhav Sharma, and Vasudeva Varma. 2023. WebNLG challenge 2023: Domain adaptive machine translation for low-resource multilingual RDF-to-text generation (WebNLG 2023). In *Proceedings of the Workshop on Multimodal, Multilingual Natural Language Generation and Multilingual WebNLG Challenge (MM-NLG 2023)*, pages 93–94, Prague, Czech Republic. Association for Computational Linguistics.

Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. Language (technology) is power: A critical survey of "bias" in nlp. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5476.

Marta Boscariol, Luana Bulla, Lia Draetta, Beatrice Fiumanò, Emanuele Lenzi, and Leonardo Piano. 2024. Evaluation of llms on long-tail entity linking in historical documents (short paper). In *EKAW (Companion)*.

Thiago Castro Ferreira, Claire Gardent, Nikolai Ilinykh, Chris van der Lee, Simon Mille, Diego Moussallem, and Anastasia Shimorina, editors. 2020. *Proceedings of the 3rd International Workshop on Natural Language Generation from the Semantic Web (WebNLG+)*. Association for Computational Linguistics, Dublin, Ireland (Virtual).

Liam Cripwell, Anya Belz, Claire Gardent, Albert Gatt, Claudia Borg, Marthese Borg, John Judge, Michela Lorandi, Anna Nikiforovskaya, and William Soto Martinez. 2023. The 2023 webnlg shared task on low resource languages. overview and evaluation results (webnlg 2023). In *Association for Computational Linguistics*, page 55–66.

Bayu Distiawan, Jianzhong Qi, Rui Zhang, and Wei Wang. 2018. Gtr-lstm: A triple encoder for sentence generation from rdf data. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1627–1637.

Lia Draetta, Marco Antonio Stranisci, Flaviana Corallo, Pier Felice Balestrucci, Michael Oliverio, Rossana Damiano, and Alessandro Mazzei. 2025. Beyond the metrics: an investigation into the reliability of evaluation metrics for domain specific graph-based question answering.

Kevin Du, Vésteinn Snæbjarnarson, Niklas Stoehr, Jennifer White, Aaron Schein, and Ryan Cotterell. 2024. Context versus prior knowledge in language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13211–13235, Bangkok, Thailand. Association for Computational Linguistics.

Daniel Duma and Ewan Klein. 2013. Generating natural language from linked data: Unsupervised template extraction. In *Proceedings of the 10th International Conference on Computational Semantics (IWCS 2013)–Long Papers*, pages 83–94. ASSOC COMPUTATIONAL LINGUISTICS-ACL.

Morten W Fagerland and Leiv Sandvik. 2009. The wilcoxon–mann–whitney test under scrutiny. *Statistics in medicine*, 28(10):1487–1497.

Thiago Castro Ferreira, Diego Moussallem, Emiel Krahmer, and Sander Wubben. 2018. Enriching the webnlg corpus. In *Association for Computational Linguistics*, page 171–176.

Hanning Gao, Lingfei Wu, Po Hu, Zhihua Wei, Fangli Xu, and Bo Long. 2021. Rdf-to-text generation with reinforcement learning based graph-augmented structural neural encoders. *CoRR*.

Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. 2017a. Creating training corpora for nlg micro-planning. In *55th Annual Meeting of the Association for Computational Linguistics, ACL 2017*, pages 179–188. Association for Computational Linguistics (ACL).

Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. 2017b. The webnlg challenge: Generating text from rdf data. In *10th International Conference on Natural Language Generation*, pages 124–133. ACL Anthology.

Arianna Graciotti, Leonardo Piano, Nicolas Lazzari, Enrico Daga, Rocco Tripodi, Valentina Presutti, and Livio Pompianu. 2025. Ke-mhisto: Towards a multilingual historical knowledge extraction benchmark for addressing the long-tail problem. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 20316–20339.

Nuno M. Guerreiro, Ricardo Rei, Daan van Stigt, Luisa Coheur, Pierre Colombo, and André F. T. Martins. 2024. xCOMET: Transparent machine translation evaluation through fine-grained error detection. *Transactions of the Association for Computational Linguistics*, 12:979–995.

Aidan Hogan, Xin Luna Dong, Denny Vrandečić, and Gerhard Weikum. 2025. Large language models, knowledge graphs and search engines: A crossroads for answering users' questions. *arXiv preprint arXiv:2501.06699*.

Fred Jelinek, Robert L Mercer, Lalit R Bahl, and James K Baker. 1977. Perplexity—a measure of the difficulty of speech recognition tasks. *The Journal of the Acoustical Society of America*, 62(S1):S63–S63.

Ming Jiang and Mansi Joshi. 2024. CPopQA: Ranking cultural concept popularity by LLMs. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pages 615–630, Mexico City, Mexico. Association for Computational Linguistics.

Daniel Jurafsky and James H. Martin. 2025. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition, with Language Models*, 3rd edition. Online manuscript released August 24, 2025.

Nikhil Kandpal, Haikang Deng, Adam Roberts, Eric Wallace, and Colin Raffel. 2023. Large language models struggle to learn long-tail knowledge. In *Proceedings of the 40th International Conference on Machine Learning*, pages 15696–15707.

Rohan Kumar, Youngmin Kim, Sunitha Ravi, Haitian Sun, Christos Faloutsos, Ruslan Salakhutdinov, and Minji Yoon. 2024. Automatic question-answer generation for long-tail knowledge. *arXiv preprint arXiv:2403.01382*.

Huihan Li, Yuting Ning, Zeyi Liao, Siyuan Wang, Xiang Lorraine Li, Ximing Lu, Wenting Zhao, Faeze Brahman, Yejin Choi, and Xiang Ren. 2024. In search of the long-tail: Systematic generation of long-tail inferential knowledge via logical rule guided search. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 2348–2370, Miami, Florida, USA. Association for Computational Linguistics.

Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.

Arle Lommel, Aljoscha Burchardt, Maja Popović, Kim Harris, Eleftherios Avramidis, and Hans Uszkoreit. 2014. Using a new analytic measure for the annotation and analysis of MT errors on real data. In *Proceedings of the 17th Annual Conference of the European Association for Machine Translation*, pages 165–172, Dubrovnik, Croatia. European Association for Machine Translation.

Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. When not to trust language models: Investigating effectiveness of parametric and non-parametric memories. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9802–9822.

Sebastien Montella, Betty Fabre, Tanguy Urvoy, Johannes Heinecke, and Lina M Rojas Barahona. 2020. Denoising pre-training and data augmentation strategies for enhanced rdf verbalization with transformers. In *Proceedings of the 3rd International Workshop on Natural Language Generation from the Semantic Web (WebNLG+)*, pages 89–99.

Michael Oliverio, Pier Felice Balestrucci, Alessandro Mazzei, and Valerio Basile. 2025. Webnlg-it: Construction of an aligned rdf-italian corpus through machine translation techniques. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 12073–12083.

Michael Oliverio, Pier Felice Balestrucci, Alessandro Mazzei, Valerio Basile, and 1 others. 2024. Dipinfounito at the gem'24 data-to-text task: Augmenting llms with the split-generate-aggregate pipeline. In *Proceedings of the 17th International Natural Language Generation Conference: Generation Challenges*, pages 59–65. Association for Computational Linguistics.

Riccardo Orlando, Pere-Lluís Huguet Cabot, Edoardo Barba, and Roberto Navigli. 2024. Relik: Retrieve and link, fast and accurate entity linking and relation extraction on an academic budget. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 14114–14132.

Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.

Karl Pearson. 1895. Vii. note on regression and inheritance in the case of two parents. *proceedings of the royal society of London*, 58(347-352):240–242.

Laura Perez-Beltrachini, Rania Sayed, and Claire Gardent. 2016. Building RDF content for data-to-text generation. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 1493–1502, Osaka, Japan. The COLING 2016 Organizing Committee.

Maja Popović. 2017. chrf++: words helping character n-grams. In *Proceedings of the second conference on machine translation*, pages 612–618.

Virginia Ramón-Ferrer, Carlos Badenes-Olmedo, and Oscar Corcho. 2025. Spanish triple-to-text benchmark on low-resource large language models.

Gabriele Sarti and Malvina Nissim. 2024. IT5: Text-to-text pretraining for Italian language understanding and generation. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 9422–9433, Torino, Italia. ELRA and ICCL.

Anastasia Shimorina, Elena Khasanova, and Claire Gardent. 2019. Creating a corpus for russian data-to-text generation using neural machine translation and post-editing. In *Association for Computational Linguistics*, pages 44–49.

Charles Spearman. 1961. The proof and measurement of association between two things.

Marco Antonio Stranisci, Eleonora Bernasconi, Viviana Patti, Stefano Ferilli, Miguel Ceriani, and Rossana Damiano. 2023. The world literature knowledge graph. In *International Semantic Web Conference*, pages 435–452. Springer.

Kai Sun, Yifan Xu, Hanwen Zha, Yue Liu, and Xin Luna Dong. 2024. Head-to-tail: How knowledgeable are large language models (LLMs)? A.K.A. will LLMs replace knowledge graphs? In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 311–325, Mexico City, Mexico. Association for Computational Linguistics.

Marcos V Treviso, Nuno M Guerreiro, Sweta Agrawal, Ricardo Rei, José Pombal, Tania Vaz, Helena Wu, Beatriz Silva, Daan Van Stigt, and Andre Martins. 2024. xTower: A multilingual LLM for explaining and correcting translation errors. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 15222–15239, Miami, Florida, USA. Association for Computational Linguistics.

Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: a free collaborative knowledgebase. *Communications of the ACM*, 57(10):78–85.

Fan Zhang, Meishan Zhang, Shuang Liu, Yueheng Sun, and Nan Duan. 2023. Enhancing rdf verbalization with descriptive and relational knowledge. *ACM Transactions on Asian and Low-Resource Language Information Processing*, 22(6):1–18.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.

A Verbalization guidelines

The annotators were asked to verbalize the data-units following guidelines:

A.1 Verbalization

- Verbalize the triples as naturally as possible, aiming to create grammatically correct and fluent sentences. You may choose whether or not to follow the original order of the triples. Please indicate this in the "Order preserved" column by writing YES if the triple is verbalized in the same order, or NO if it is not.
- You are free to break down the information conveyed by each set of triples in one or more sentences according to your preferences.
- If a triple makes no sense (e.g., "Black Racer", "different from", "Black Racer"), you may choose not to verbalize it. In this case, write NA in the annotation column.
- If a triple is very difficult or uncertain to interpret, you may propose a verbalization and

leave comments in the Notes column. A colleague will review uncertain verbalizations later.

- If you have the necessity to check the meaning or the translation of certain entities of claim you can use WikiData `https://www.wikidata.org/wiki/Wikidata:Main_Page`.
- Named entities
  - For named entities, if an equivalent exists in Italian or Spanish, please use the correct translation in your verbalization.
  - If you have doubts about the translation of a claim, you can find the corresponding version in Spanish or Italian on Wikidata.
  - If you cannot find an equivalent in Italian or Spanish, do not translate it. Instead, flag the triple as problematic in the "Notes" column.

A.2 Validation

In this phase, the annotators will cross-check and review colleagues' annotations, evaluating whether each verbalization is acceptable by answering the following questions:

- Does the text sound fluent and natural?
- Does the text include all and only the information from the data?
- Is the text written in good language (no spelling or grammatical mistakes)?

B Data Unit distribution per Category

Category	Count
Politician	47
Airport	46
Model	30
SportsTeam	37
Building	47
Company	44
Artist	41
City	47
MeanOfTransportation	20
Food	35
CelestialBody	17
Astronaut	33
Monument	47
ComicsCharacter	48
River	30
University	47
WrittenWork	26
Athlete	47
Type	**Count**
Long-tail	408
Head	218
Total entries	**689**

Table 4: Distribution of data entries across categories and entity types.

C Translation

C.1 Error correction

You are provided with a Source, Translation, Translation quality analysis, and Translation quality score (weak, moderate, good, excellent). The Translation quality analysis contain a translation with marked error spans with different levels of severity (minor, major or critical). Additionally, we may provide a **reference translation**. Given this information, generate an explanation for each error and a fully correct translation.
Italian source: Osvaldo il congilio è un personaggio dell'opera Lupo Alberto
English translation: Osvaldo the congil is a character from the opera Wolf Albert.
Translation quality analysis: Osvaldo the con<error severity="minor">gil</error> is a character from the opera<error severity="minor">Wolf Albert</error>.
Translation quality score: good

The corresponding output is:

Explanation for error: The term "congilio" is a typographical error and should be corrected to "coniglio," which means "rabbit" in Italian. The correct translation of "Lupo Alberto" is "Wolfy Albert," as "Lupo" means "wolf" and "Alberto" is a common name in Italian. The translation also misses the gender of "Osvaldo," which is masculine, and the correct translation should reflect that. Additionally, the translation of "opera" as

"opera" is correct, but the translation of "Wolf Albert" as "Wolf Albert" is incorrect because it does not convey the correct name of the work. The correct translation should reflect the proper name of the work, which is "Lupo Alberto."

C.2 Guidelines Human MT Evaluation

Annotators were instructed to strictly follow the evaluation guidelines and, when necessary, to correct the MT output by providing a correct or more accurate version.

Annotate as **1** if the error is present and **0** if not, respecting the following taxonomy:

Accuracy

Issues related to whether the information content of the target is equivalent to the source. Subcategories:

- **Terminology:** Issues related to the use of domain-specific terms.
- **Mistranslation:** Issues related to the improper translation of content, including erroneous translation of named entities.
- **Omission:** Content present in the source is missing in the target.
- **Addition:** Content not present in the source has been added to the target.
- **Untranslated:** Text inappropriately appears in the source language.

Fluency

Issues related to the linguistic properties of the target without relation to its status as a translation. Subcategories:

- **Grammar:** Issues related to grammatical properties of the text.
- **Style:** The text shows stylistic problems.
- **Spelling:** The text is spelled incorrectly.
- **Typography:** Problems related to typographical conventions.
- **Unintelligible:** Text is garbled or otherwise unintelligible; indicates a major breakdown in fluency.

Corrected Translation

If the target translation is incorrect, you may propose a new translation in the "Corrected translation" column.

Guidelines for Annotators

- Evaluate whether translating a Named Entity (NE) makes sense in context, taking into account cultural and linguistic conventions; some NEs may be left untranslated.
- Named entities should be considered correctly translated when it is common in the target language to refer to the entity in that way; otherwise, mark them as incorrect. Examples:
  - Crist Redemptor → Cristo Redentore → Christ the Redeemer [Correct]
  - Lupo Alberto → Wolfy Albert [Incorrect]
  - Louis the Springer → Luigi il Precursore [Incorrect]

  Wrong NEs must be annotated in the column "Mistranslation".
- Assess when it is appropriate to keep a term in English (or another source language).
- Pay attention to the correct use of verb tenses.
- Ensure agreement between singular and plural.
- The translation should be read naturally in the target language.
- Ensure that the meaning is accurately preserved.
- Annotators should mark **1** for each subcategory when an issue occurs. If you label **1**, specify why in the "note" column.

C.3 Manual assessment of translation

In the table below are reported the Spearman and Pearson correlation between the two annotators.

D TailNLG Benchmark

Each entry contains the following metadata labels:

- **Category**: domain/class of the subject (e.g., Artist, City, Monument).
- **Eid**: internal entry id (e.g., `Id1`).

Source	Target	ρ	r
es	en	0.23	0.25
en	es	0.32*	0.27
it	en	0.15	0.02
en	it	0.56*	0.64
es	it	0.26	0.22
it	es	0.53*	0.49*

Table 5: Spearman (ρ) and Pearson (r) correlation coefficients across language pairs. Source and Target denote the source language and the target language into which the source is translated, respectively. Asterisks (*) denote a statistically significant difference ($p < 0.05$).

- **Shape_type**: data Units configuration (chain, sibling, mixed) that controls structural complexity.
- **Size**: number of triples for Data Units.
- **Qid**: Wikidata Id.
- **Type**: Distinction between head and long-tail entities
- **Sub_type**: Distinction between subtypes (e.g. head entities with rare claims).
- **triplesets**: `originaltripleset` (as collected) and `modifiedtripleset` (normalized input).
- **lexicalizations**: one or more `lex` elements with attributes `quality` (gold/silver) and `lang` (en/es/it).

E Prompts

English version In English, structured data is commonly represented as triples, with the format [subject, predicate, object]. Based on these triples, generate a single-paragraph text composed of complete, grammatically correct, and natural sentences.
INSTRUCTIONS:
- Generate the text solely from the input triples
- Return the final verbalization with this format: The final verbalization is: [verbalization output]
- Insert the verbalization of the input triples within square brackets [], without adding anything else.

Spanish version En español, los datos estructurados se representan comúnmente como tríos, con el formato [sujeto, predicado, objeto]. Basándose en estos tríos, genere un texto de un solo párrafo compuesto por oraciones completas, gramaticalmente correctas y naturales.
INSTRUCCIONES:
- Genere el texto únicamente a partir de las tripletas de entrada.
- Devuelva la verbalización final con este formato: La verbalización final es: [salida de verbalización]
- entre corchetes [] inserte la verbalización de las tripletas de entrada, sin añadir nada más

Italian version In italiano, i dati strutturati sono comunemente rappresentati come triple, con il formato [soggetto, predicato, oggetto]. Sulla base di queste triple, genera un testo di un solo paragrafo composto da frasi complete, grammaticalmente corrette e naturali.
INSTRUZIONI:
- Genera il testo esclusivamente dalle triple in input
- Restituisci la verbalizzazione finale con questo formato: La verbalizzazione finale è: [output di verbalizzazione]
- tra le parentesi quadre [] inserisci la verbalizzazione delle triple in input, senza aggiungere altro.

F Results

Model	Type	BERT_r	BLEU ↑	chrF ↑	R-1 ↑	R-2 ↑	R-L ↑	PPL ↓
Qwen-3B	Long-tail	0.50	0.13	50.10	0.53	0.30	0.45	182.87
	Top-head	0.51	0.13	49.12	0.52	0.29	0.44	70.45*
Qwen-7B	Long-tail	0.56	0.19	54.27	0.58	0.35	0.49	138.70
	Top-head	0.58	0.18	53.41	0.57	0.35	0.48	87.96*
Gemma-4b	Long-tail	0.53	0.17	54.32	0.57	0.35	0.48	4724.94
	Top-head	0.56	0.17	53.48	0.56	0.34	0.46	3488.00*
Gemma-12b	Long-tail	0.62	0.24	60.38	0.64	0.42	0.55	455.57*
	Top-head	0.64	0.23	58.73	0.63	0.40	0.53	797.47
Llama-3B	Long-tail	0.49	0.13	49.75	0.51	0.30	0.43	102.40
	Top-head	0.50	0.14	48.19	0.49	0.29	0.42	72.55*
Llama-8B	Long-tail	0.55	0.18	53.81	0.58	0.36	0.50	134.53
	Top-head	0.57	0.19	53.34	0.57	0.36	0.48	74.13*

Table 6: Comparison of models by type. Asterisks (*) denote a statistically significant difference with $p < 0.05$.

Model	Lang	BERT_r	BLEU ↑	chrF ↑	R-1 ↑	R-2 ↑	R-L ↑	PPL ↓
Llama-3B	en	0.58*	0.21*	56.67*	0.63*	0.40*	0.53*	73.51*
	es	0.46	0.11	46.40	0.47	0.26	0.39	114.99
	it	0.44	0.08	44.26	0.42	0.22	0.35	81.98
Llama-8B	en	0.63*	0.25*	60.72*	0.66*	0.43*	0.56*	79.37*
	es	0.59	0.18	53.27	0.58	0.35	0.49	170.14
	it	0.45	0.13	46.85	0.48	0.28	0.42	80.23
Qwen-3B	en	0.50	0.17*	52.79*	0.60*	0.36*	0.51*	84.29
	es	0.54*	0.12	49.10	0.51	0.28	0.43	227.52
	it	0.48	0.09	47.19	0.46	0.24	0.39	100.23
Qwen-7B	en	0.67*	0.29*	62.71*	0.69*	0.46*	0.59*	119.43
	es	0.60	0.16	52.78	0.56	0.33	0.48	149.16
	it	0.45	0.11	46.25	0.46	0.26	0.40	85.39*
Gemma-4b	en	0.62*	0.23*	60.32*	0.66*	0.42*	0.55*	3798.99
	es	0.51	0.14	50.76	0.53	0.31	0.44	6068.41
	it	0.50	0.14	50.85	0.50	0.30	0.43	2794.93*
Gemma-12b	en	0.64*	0.27*	61.68*	0.68*	0.44*	0.57*	357.51
	es	0.63	0.21	58.27	0.62	0.39	0.52	739.90
	it	0.61	0.22	59.17	0.62	0.40	0.53	687.37

Table 7: Average scores for different languages for each model (zero-shot). Asterisks (*) denote a statistically significant difference with $p < 0.05$.

Model	Quality	BERT_r ↑	BLEU ↑	chrF ↑	R-1 ↑	R-2 ↑	R-L ↑	PPL ↓
Llama-3B	gold	0.46	0.12	48.04	0.48	0.27	0.40	95.30
	silver	0.52*	0.14	49.80	0.52*	0.30*	0.44*	86.86*
Llama-8B	gold	0.52	0.18	52.68	0.50	0.55*	0.48	90.94
	silver	0.58*	0.19	54.22	0.59*	0.36	0.50*	122.10*
Qwen-3B	gold	0.50	0.13	50.62	0.53	0.30	0.44	94.12
	silver	0.51	0.13	49.08	0.52	0.29	0.45	165.11
Qwen-7B	gold	0.53	0.17	52.59	0.55	0.33	0.47	117.53
	silver	0.60*	0.20	54.75	0.59*	0.36	0.50	118.35
Gemma-4b	gold	0.52	0.17	53.67	0.55	0.34	0.46	1143.13
	silver	0.55*	0.18	54.18	0.57	0.35	0.48	6197.47
Gemma-12b	gold	0.61	0.23	59.88	0.63	0.41	0.54	406.83
	silver	0.64*	0.24	59.59	0.64	0.41	0.54	715.74

Table 8: Average scores for different quality levels (silver, gold) for each model (zero-shot). Asterisks (*) denote a statistically significant difference with $p < 0.05$.

Type	Lang	BERT_r	BLEU ↑	chrF ↑	R-1 ↑	R-2 ↑	R-L ↑	PPL ↓
Llama-3B								
Long-tail	en	0.57	0.21	56.58	0.62	0.39	0.53	86.17
	es	0.45	0.11	47.07	0.48	0.26	0.40	144.99
	it	0.44	0.09	45.61	0.44	0.23	0.37	75.80
	total	0.49	0.13	49.75	0.51	0.30	0.43	102.32
WebNLG	en	0.56	0.23	57.17	0.71*	0.47*	0.56	53.67*
	es	0.49*	0.16*	49.20	0.58*	0.35*	0.45*	68.28*
	it	0.48*	0.13*	47.75*	0.53*	0.30*	0.42*	41.18*
	total	0.51	0.17*	51.37	0.60	0.37	0.48*	54.38*
Llama-8B								
Long-tail	en	0.61	0.25	60.39	0.66	0.43	0.56	93.36
	es	0.58*	0.17	53.30	0.58	0.35	0.49	222.16
	it	0.46*	0.13	47.73	0.49	0.29	0.43	88.06
	total	0.55	0.18	53.81	0.58	0.36	0.50	134.53
WebNLG	en	0.61	0.30	65.53	0.74*	0.51*	0.59*	42.56*
	es	0.55	0.19*	52.43	0.66*	0.43*	0.53*	90.42*
	it	0.43	0.15*	45.35	0.54*	0.34*	0.46*	71.73*
	total	0.53	0.21*	53.77	0.65*	0.43*	0.53*	68.23*
Gemma-4B								
Long-tail	en	0.60	0.23	59.85	0.66	0.42	0.55	4886.76
	es	0.50	0.14	51.49	0.54	0.32	0.45	8074.87
	it	0.49	0.15	51.62	0.51	0.31	0.44	1213.20
	total	0.53	0.17	54.32	0.57	0.35	0.48	4724.94
WebNLG	en	0.61	0.28*	62.43*	0.74*	0.50*	0.59*	930.34*
	es	0.52	0.20*	53.79*	0.61*	0.39*	0.48*	1493.20*
	it	0.49	0.18*	53.72	0.57*	0.36*	0.46	601.97*
	total	0.54	0.22*	56.65*	0.64*	0.42*	0.51	1008.50*
Gemma-12B								
Long-tail	en	0.63	0.26	61.65	0.67	0.44	0.58	330.02*
	es	0.62	0.21	58.80	0.62	0.40	0.53	611.42
	it	0.61	0.24	60.68	0.64	0.42	0.55	425.27
	total	0.62	0.24	60.38	0.64	0.42	0.55	455.57
WebNLG	en	0.66*	0.33*	65.97	0.77*	0.53*	0.62*	333.93
	es	0.65	0.30*	62.68	0.71*	0.49*	0.57*	210.04*
	it	0.64*	0.28*	62.24*	0.70*	0.47*	0.57	326.33*
	total	0.65*	0.30*	63.63	0.73*	0.50*	0.59*	290.10*
Qwen-3B								
Long-tail	en	0.49	0.17	52.66	0.59	0.35	0.51	100.79
	es	0.53	0.12	49.47	0.52	0.29	0.44	325.03
	it	0.48	0.10	48.08	0.48	0.25	0.41	124.49
	total	0.50	0.13	50.07	0.53	0.30	0.45	183.44
WebNLG	en	0.55*	0.24*	57.05	0.72*	0.47*	0.57*	51.40*
	es	0.56*	0.17*	51.54	0.63*	0.38*	0.49*	150.84*
	it	0.55*	0.15*	51.75	0.59*	0.34*	0.47*	55.81*
	total	0.55*	0.19*	53.44	0.64*	0.40*	0.51*	86.02*
Qwen-7B								
Long-tail	en	0.65*	0.29	62.47	0.69	0.46	0.58	134.13
	es	0.60	0.16	53.43	0.57	0.34	0.49	181.75
	it	0.44	0.11	46.82	0.47	0.26	0.41	100.26
	total	0.56	0.19	54.24	0.58	0.35	0.49	138.71
WebNLG	en	0.68	0.34*	66.33	0.79*	0.55*	0.64*	46.52*
	es	0.60	0.24	56.19	0.66*	0.43*	0.53*	117.41*
	it	0.44	0.15	46.96	0.54*	0.33*	0.45*	122.70*
	total	0.57	0.24	56.49	0.66*	0.44*	0.54*	95.54*

Table 9: Performance comparison of models on long-tail triples from TailNLG and WebNLG test set. Asterisks (*) denote a statistically significant difference with $p < 0.05$.