

Language Ideologies in a Multilingual Society: An LLM-based Analysis of Luxembourgish News Comments

Emilia Milano¹, Alistair Plum¹, Yves Scherrer², Christoph Purschke¹

¹University of Luxembourg, Esch-sur-Alzette, Luxembourg,

²University of Oslo, Oslo, Norway

{emilia.milano, alistair.plum, christoph.purschke}@uni.lu

yves.scherrer@ifi.uio.no

Abstract

Detecting language ideologies is a valuable yet complex task for understanding how identities are constructed through discourse. In Luxembourg’s multicultural and multilingual society, language ideologies reflect more than simple preferences: they carry deep cultural and social meanings, shaping identities and social belonging. Following recent developments in applying Natural Language Processing tools to linguistics and social science, this paper explores the potential of large language models to assist in the detection of language ideologies. We manually annotate a corpus of user comments in Luxembourgish with predefined ideological categories and then evaluate the performance of large language models under varying prompt conditions to assess their ability to replicate these human annotations. Since Luxembourgish is a small language and poorly represented in the LLMs’ training data, we also investigate whether machine-translating the data to high-resource languages increases performance on the ideology detection task. Our findings suggest that, while LLMs are not yet fully optimized for a multi-class ideological annotation task, they are practical tools to identify language ideological content.

Keywords: Language Ideologies, Large Language Models, Computational Sociolinguistics

1. Introduction

Automatic detection of language ideologies is a promising but challenging task. Language ideologies, or shared beliefs about language, play a central role in reinforcing and establishing identities and power imbalances (Blommaert, 2010). Although culturally tied to the speakers’ community, similar language ideologies have been found in different contexts. Often, language ideologies are linked to matters of social belonging and discrimination (Grondelaers and van Hout, 2011; Lippi-Green, 2012). This is also found in our case study, Luxembourg. The country is characterized by a complex and dynamic societal multilingualism, in which the choice of language in interaction carries rich social meaning aside from practical requirements or personal preference (Horner and Weber, 2008). This is evidenced in public discourses on the state of multilingualism and the role of Luxembourgish, where language choices and ideological stances are often closely linked. For example, the phrase *En français, s’il vous plaît* (‘In French, please’), used by (French-speaking) service personal as a reaction to requests in Luxembourgish, has become a prominent figure in discourse (Purschke, 2025): It is interpreted (and quoted) by speakers of Luxembourgish as a sign of the unwillingness to integrate in society, linguistically and socially, and linked to questions of demographic development and language vitality (Purschke, 2023).

Hence, being able to automatically and systematically detect language ideologies in large amounts

of text would be beneficial for obtaining a deeper understanding of cultural dynamics and language-related discourse at scale.

Against this backdrop, the increasing use of Large Language Models (LLMs) for linguistic (Manning et al., 2020; Wu et al., 2025) and social science research (Ma et al., 2024) presents a promising opportunity to systematically investigate the automatic detection and classification of language ideologies.

For this study, we base ourselves on a corpus of written Luxembourgish comments provided by the national broadcaster RTL¹ (Ranasinghe et al., 2023). We annotate the comments at two levels: (1) whether they are related to language ideology or not, and (2) which different types of language ideologies they contain (see Section 2.2). We assess a range of open-weights and closed-weights LLMs on their ability to predict these annotations, testing several prompts in both zero-shot and few-shot settings.

Since Luxembourgish is a small language, we presume that it is comparatively poorly represented in the LLMs’ training data and that this results in low classification performance. We therefore follow the *translate-test* paradigm and investigate whether machine translating the Luxembourgish comments into a higher-resourced language (specifically, English, French or German) increases classification accuracy.

¹Our agreement with the data provider prevents us from publicly releasing the corpus. However, we can provide upon request the comment IDs so that the corpus can be recreated.

More specifically, our research questions are:

1. Is language ideology classification a suitable task for the current generation of LLMs?
2. Do the languages of the data and the quality of their translations have an impact on LLM performance on language ideology classification?
3. Can LLMs contribute to inform sociolinguistic research without requiring high technical expertise or task-specific model adaptation?

To explore the feasibility and challenges of the task, we provide LLMs with language ideology categories, descriptions and examples, and instructions to annotate accordingly. Our contributions offer an evaluation of LLMs for language ideologies identification, assessing both the impact of the source language and translations on the same task. In doing so, we provide an overview of how LLMs work with Luxembourgish on cultural tasks, compared to more established languages in NLP. We find that fine-grained language ideology detection is a complex task for LLMs, and that translations do not lead to substantial performance improvement. Integrating NLP tools into sociolinguistic research, we acknowledge limitations and possible improvements of the task setup.

2. Related work

In this section, we discuss work related to Luxembourgish NLP (Section 2.1), as well as language ideologies (Section 2.2). After that, we discuss work on ideologies and LLMs (Section 2.3), as well as cross-lingual transfer in LLMs (Section 2.4).

2.1. Luxembourgish

Luxembourgish is a West Germanic language with about 400.000 speakers (Gilles, 2019). It is one of three official languages in Luxembourg in addition to German and French, but it was only recognized as such in 1984. Historically, Luxembourgish mainly developed as an oral language, acquiring more written domains only recently with the rise of digital and social media. As a multilingual country with an extremely dynamic language regime, Luxembourg is a particularly interesting case for sociolinguistic studies (Purschke and Gilles, 2023). However, in NLP, Luxembourgish should be viewed as a small language, although various resources have been made available in recent years (Plum et al., 2025; Lutgen et al., 2025; Philippy et al., 2024; Plum et al., 2024).

2.2. Language ideologies

We follow Blume (2003)'s description of ideologies and understand them as descriptions of collective, depersonalized normative interpretations used by individuals and groups to justify and evaluate their own and others' actions. As such, ideologies towards and about language (Blommaert, 2010) play a central role in creating and ordering social meaning (Rhodes, 2023), in shaping language policies and perception, in structuring social relations and societal belonging (Woolard, 2020), and in contributing to individual and collective identity constructions (Kroskrity, 2022).

From a methodological perspective, (language) ideologies have also been studied with quantitative tools and methods. Baker and Levon (2015) compares qualitative and corpus-based approaches to the analysis of newspaper articles on masculinity, showing that the two approaches can complement each other, and Vessey (2017) investigates language ideologies in Canada through corpus linguistics approaches. Combining different approaches to corpus study and language ideologies, Purschke (2025) analyzes language ideologies on the RTL.lu corpus as part of public discourse formation.

Building on this foundation, we search for the following five language-ideological motifs about Luxembourgish and multilingualism in the country in the RTL.lu comments.

Identity: statements about national identity established by a common language, culture, and history, or personal preference towards the national language;

Example ID 636: *'1. For us Luxembourgers (250,000 people), Luxembourgish is our first language.'*²

Vitality: statements about the state of a language use, e.g., language endangerment and decay, reasons for these developments and/or means to prevent them;

Example ID 3631: *'Good text, but I'm afraid our language would die out.'*

Belonging: statements about requirements for speakers to integrate in the local society;

Example ID 410165: *'it worked: all residents of Luxembourg have linguistically integrated themselves.'*

Responsibility: statements delegating to an actor or a group of actors responsibility for the linguistic situation in the country (e.g. language decay or support);

Example ID 666: *'And if our cross-border commuters don't speak Luxembourgish, it's our own fault.'*

²Original texts are provided in the appendix.

Recognition: statements about language policies, language status, and how people perceive the languages spoken in the country.

Example ID 636: ‘2. *Luxembourgish is FIRST OF ALL a spoken language.*’

2.3. Ideologies and LLMs

There is abundant work showing how human stereotypes, biases, and ideologies influence the performance of (large) language models (Lin and Li, 2025). They can be introduced in the (pre-)training data (Hovy and Prabhumoye, 2021), optimization stages (Dahlgren Lindström et al., 2025), and data exposure level (Chen et al., 2024), they can change according to the tested model (Lin et al., 2024) and language (Buyl et al., 2025), and can be shaped by language ideologies. For instance, Gururangan et al. (2022) investigate the role of language ideologies in selecting text data and the subsequent inequalities. LLMs have also been shown to perpetuate in-group solidarity and out-group hostility (Hu et al., 2024) – where in-group refers to the community speakers socially and emotionally identify with, and out-group refers to communities perceived as outsiders. The language material generated by LLMs has been shown to reproduce social bias (Dhamala et al., 2021) as well as language ideologies, both for the standard language ideology in the US (Smith et al., 2025) and for gendered language ideology in English (Watson et al., 2025) and English dialects (Duncan, 2024).

2.4. Cross-lingual transfer in LLMs

Luxembourgish can be considered a small language in a multilingual context characterized by higher resourced languages. Almost all Luxembourgers are fluent in the other official languages, French or German (Fehlen et al., 2023), and despite Luxembourg’s membership in the European Union, Luxembourgish is not an official EU language. Additionally, Luxembourgish is not the language of education, and its domain expansion and societal uptake in writing are relatively recent (Gilles, 2015). For all these reasons, the amount of textual data to train large-scale NLP systems is limited. In such scenarios, it is common to resort to cross-lingual transfer techniques, ensuring that a model’s high performance on a specific task in a high-resource language can be transferred to a low-resource language. A simple cross-lingual transfer technique that can be easily applied in the context of generative LLMs is called *translate-test* (Ponti et al., 2021; Artetxe et al., 2023): the test data (in our case, the user comments written in Luxembourgish) are machine-translated to a high-resource language (in our case, English, French or German) and then fed to the model. While this setup is prone to error

propagation (the translations might contain errors), it is assumed that the performance gains obtained through the model’s better “understanding” of the translated data outweigh the errors introduced in the translation step.

3. Data

We frame ideology detection as a multi-class sentence annotation task over Luxembourgish user comments from the RTL.lu corpus. We operate at sentence level to reduce topic drift within long comments. Our corpus contains 300 comments with 1524 total sentences. Comments were split into sentences using GPT, and then manually corrected where the sentence splitting did not work correctly. The main reason for using GPT for this task is that, for social media data with irregular sentence constructions and punctuation patterns (e.g., including Emoji as sentence boundaries) other sentence splitters usually struggle and GPT, compared to these, includes semantic context parsing for better accuracy. We retained the punctuation for each user and did not carry out any other pre-processing steps such as lexical normalization, as these could blur ideology markers. In a setting without a fully implemented standard, the use of incorrect or correct orthography could be an expression of specific language ideologies.

The corpus is extracted from a larger set of user comments (1,422,759 comments at the time of the study³). It contains two subsets, one subset of 150 comments on language use in the country (from now on, we will refer to this subset as *Lang*), and one subset of 150 comments not related to this topic (referred to as *NotLang*). For the *Lang* subset, the selection of comments was made through keyword matching, using a total of 19 Luxembourgish variants of *Luxembourgish, our language, mother tongue*. Of the comments retrieved (16,015), 150 are then randomly selected and manually annotated.

For the *NotLang* subset, 75 comments were selected following the same list of keywords as for the *Lang* subset, and were carefully read to make sure that they refer to topics other than language use in the country. The remaining 75 comments are randomly selected from the complete set of comments (1,422,759), and do not contain language-related keywords. By diversifying the selection criteria across subsets, we aim to prevent the models from overfitting to surface-level lexical cues and to assess their performance in identifying ideologically charged discourse about language.

The *Lang* subset is annotated by two expert annotators (the first and last authors of this study) with

³For this study, we used the 2025-04 build of the RTL corpus.

the five language ideology categories presented in Section 2.2: IDENTITY, VITALITY, BELONGING, RESPONSIBILITY and RECOGNITION. We use a sixth label NONE for sentences that cannot be assigned one of these specific labels. The observed agreement between the two annotators is 0.78 and Cohen’s Kappa is 0.66. Disagreement cases are discussed by annotators and are labeled in common agreement. The category counts across the dataset are shown in Table 1.

Category	Lang	NotLang
Identity	49	0
Vitality	72	0
Belonging	63	0
Responsibility	121	0
Recognition	125	0
None	273	821
Total	703	821

Table 1: Sentence based gold standard annotation

4. Prompt Engineering

In this first experiment, we focus on evaluating different prompts in both zero- and few-shot scenarios.

Prompts We design four prompt configurations that vary the amount of information that is passed to the model.

- **Prompt 1** lists the labels, instructs the LLM to annotate each sentence with one of the five ideology labels or NONE and to provide a justification for the annotation. Text in quotation marks should not be annotated.
- **Prompt 2** adds three positive and three negative examples per category, and six additional negative examples to prompt 1. The examples are chosen because they are clear illustrations for the categories and are consistent throughout the experiment.
- **Prompt 3** adds detailed descriptions per category to prompt 1.
- **Prompt 4** adds both the examples of prompt 2 and the detailed descriptions of prompt 3 to prompt 1.

The structure of prompt 4, the most complex one, is illustrated in Figure 1. All prompts are written in English, and all comments remain in Luxembourgish. All prompts instruct models to return a single JSON object with fixed keys and value constraints, see Figure 2 below.

Experimental Setup We treat the task as multi-class annotation with sentence level ground truth derived from the human annotation. We report weighted F1, macro F1, and per label F1. We implement quality control to improve robustness without over-correcting model behavior. We validate every response against the JSON schema and re-issue malformed outputs up to 5 retries per sentence.

We test three models of the GPT family: GPT-4o (OpenAI et al., 2023), GPT-4o-mini (OpenAI, 2024), o3 (OpenAI, 2025). We standardize decoding to reduce variance induced by sampling. Temperature is set to 0.2, maximum tokens to 3000, however, we keep defaults for safety-related flags that affect refusal behavior. We batch requests with one comment (5 sentences on average) per call. Also, we enforce per thread rate limits of rps/5 to avoid provider throttling.

Results We observe minimal differences between GPT-4o, GPT-4o-mini, and o3 under prompt 4. This prompt also consistently outperforms other configurations across models, as F1 weighted score improves up to around 0.3 from prompt 1 to prompt 4 – see Table 2. Therefore, we adopt prompt 4 as the final prompting template for all further experiments.

Prompt	o3		GPT-4o		GPT-4o-mini	
	F1w	F1m	F1w	F1m	F1w	F1m
Prompt 1	0.734	0.431	0.491	0.248	0.563	0.280
Prompt 2	0.787	0.510	0.652	0.397	0.459	0.389
Prompt 3	0.806	0.567	0.765	0.427	0.468	0.391
Prompt 4	0.806	0.559	0.783	0.441	0.734	0.427

Table 2: Results per model and prompting conditions (Weighted F1 and Macro F1)

5. Evaluating LLMs and cross-lingual transfer settings

In this main experiment, we evaluate several widely used generative LLMs on the ideology annotation task with Prompt 4. We also vary the language of the comments, comparing annotation performance on the original Luxembourgish comments with their translations to English, French and German.

Machine translation To answer research question 2, we translate the original material into German, French and English via the Google Translate API. The automatic translations receive targeted manual checks⁴ that correct mistranslated country and language mentions. Additionally, 700 sentences evenly split between Lang and NotLang are

⁴Automatic translation quality for Luxembourgish can be uneven, therefore, subtle stance cues may shift.

Prompt Text
Detect one or more of the following categories in the sentences of the comment below, after the instructions and schemas. The categories are: Identity, Vitality, Belonging, Responsibility, Recognition. These categories concern language ideologies and only these exact categories can be used for the annotation. No overlap between categories is allowed, each sentence can have 0 or 1 category. If you cannot identify one of the category assign N/A. Do not take any quoted parts (i.e. text in quotation marks) into account for the labelling. For each sentence, explain why you have assigned that label. Here are the explanations of the five categories: [...]
Recognition: The language related ideology 'recognition' groups opinion on Luxembourgish language and other languages spoken in the country. In this category, different ways of acknowledging languages are involved: [...]
Examples: [...]
Examples where this label is not applicable: [...]

Figure 1: Prompt 4 comprehensive of instructions, definitions, and examples

```
"sentence_id": "<string>",
"text": "<original sentence>",
"labels": ["Identity" | "Vitality"
| "Belonging" | "Responsibility" |
"Recognition" | "None"],
"rationale": "<justification>"
```

Figure 2: JSON representation of the required schema.

manually reviewed to improve translation quality; correcting mistakes due to orthographic variation and lack of capitalization, and adapting ideology bearing content to the target languages. The distribution of categories within the two translation sets is described in Table 3. We run the same prompting and inference pipeline on each language set. Therefore, any observed changes can be attributed to language rather than prompt drift.

Labels	Post-Edited	Automatic
Identity	18	31
Vitality	22	50
Belonging	16	47
Responsibility	64	57
Recognition	56	69
None	525	569

Table 3: Translation sets

Models We evaluate widely used frontier and open generative LLMs for applied annotation workflows. The frontier group includes GPT-5 (OpenAI, 2025), while the open group includes Aya Expansive-32B (Aya; Dang et al. 2024), Llama 4 Scout 17B Instruct (Llama; AI 2025a), Mistral Magistral Small 24B (Magistral; AI 2025b), Qwen3-Next-80B-A3B-Thinking (Qwen; Yang et al. 2025), DeepSeek-V3 (DeepSeek; Liu et al. 2024), and gpt-oss-20b (GPT-OSS; OpenAI et al. 2025). All models follow the experimental setup described in Section 4.

Lang.	DeepSeek		GPT-5		Magistral		Qwen	
	F1w	F1m	F1w	F1m	F1w	F1m	F1w	F1m
Luxemb.	0.537	0.478	0.558	0.497	0.436	0.351	0.500	0.380
German	0.518	0.469	0.551	0.500	0.450	0.381	0.529	0.399
French	0.507	0.474	0.585	0.525	0.461	0.406	0.532	0.404
English	0.490	0.446	0.596	0.542	0.462	0.346	0.514	0.392

Table 4: Comparison of models per language (Weighted F1 and Macro F1)

Results Of the models tested, GPT-OSS is discarded because it is unable to produce consistent output. Aya and Llama are also discarded, as they return up to 10 and 9 labels respectively even when explicitly stating the number of labels in the prompt – probably because of their limited reasoning capabilities compared to the other tested LLMs. Although Qwen and Magistral also sporadically return additional categories, we keep them in our experiments as the categories are justified by the model-generated explanations as a subset of the main ones (e.g., PERSONAL IDENTITY is mapped to IDENTITY).

Table 4 shows the performance of the remaining models broken down by language. While DeepSeek gets the highest scores with the original Luxembourgish data, the remaining models work slightly better with the translations. GPT-5 outperforms the other models in all languages, and the best overall performance is obtained with the English translations.

5.1. Translation effects

Zooming in on the two different sets of translations, we see that the quality of the translations does not affect model performance. For both translation sets, the best performing model is GPT-5 in all languages. As shown in Table 5, the GPT-5 performance gap between English and Luxembourgish is larger with automatic translations than with post-edited ones. For DeepSeek, Luxembourgish seems to work better than English both for automatic translations and post-edited ones.

Lang.	DeepSeek		GPT-5		Magistral		Qwen	
	Auto	PE	Auto	PE	Auto	PE	Auto	PE
Luxemb.	0.762	0.771	0.797	0.780	0.699	0.737	0.755	0.661
German	0.747	0.783	0.803	0.767	0.728	0.738	0.774	0.753
French	0.731	0.781	0.806	0.800	0.721	0.744	0.768	0.766
English	0.743	0.750	0.829	0.783	0.724	0.746	0.767	0.740

Table 5: Comparison of models on Automatic (Auto) and Post-Edited (PE) translations across languages and on the original Luxembourgish corpus with the same split (Weighted F1)

Discussion To answer research question 2, translating to more-represented languages in NLP produces a limited performance improvement, whereas we find no improvement in the task between using automatic and post-edited translation sets. We consider two potential explanations to this: cross-lingual transfer in LLMs and the challenge of transferring cultural and social meanings from the source text to the target languages.

On the one hand, Luxembourgish shares some grammatical and lexical features with high-resource languages – mainly German and French – which are represented in the models’ training data and officially supported by them. Therefore, it is quite likely that the tested models generalize from shared multilingual representations (Zhang et al., 2024).

On the other hand, the way sociocultural meanings are treated in the translation process may also negatively impact the classification performance. For the post-edited set, we tried to find a middle ground between preserving the same language ideologies, reporting appropriate linguistic cues to each target language, and avoiding introducing noise in the translations. For instance, capitalization for emphasis is kept, while lack of orthographic capitalization is corrected in the translations. In the non-post-edited set, only mistranslated country and language mentions are corrected. However, our results do not show an increase in annotation performance that could be ascribed to the different translation conditions. Similarly, there is no significant performance improvement between the original Luxembourgish comments and the translations. This is attributable to the failure to eliminate noise and reproduce ideology in target languages.

Due to these challenges, and to the LLMs’ relatively even performances across languages, keeping the original text seems more appropriate.

5.2. Label-specific results

Table 6 shows F1 scores averaged on the four languages per category and model. All analyzed models perform better for discriminating between items with and without language ideologies than for distinguishing between individual categories. In fact, the absence of language ideologies – the label NONE – is the category with higher F1 scores com-

pared to the others (around 0.9). Compared to our gold standard, RESPONSIBILITY and RECOGNITION are under-detected, while VITALITY, IDENTITY, and BELONGING are over-predicted.

DeepSeek and GPT-5 achieve the highest per-label F1 scores among the evaluated models. Table 7 shows per-label F1 scores for GPT-5 and DeepSeek on the original data, compared to the highest F1 score overall. GPT-5 has the highest scores in English for BELONGING (0.571) and VITALITY (0.555), and in French for RECOGNITION (0.565). DeepSeek has highest scores in French for IDENTITY (0.475) and in Luxembourgish for RESPONSIBILITY (0.398). These values are not significantly higher than the F1 scores obtained for Luxembourgish for DeepSeek and GPT-5, except for RECOGNITION for DeepSeek (0.330).

Analysis Here, we examine the annotation categories in detail, focusing on the original Luxembourgish data and the Lang subset. We discuss consistent overlaps among categories, as well as potential causes of mislabeling. To understand the cause of systematic overlaps, we refer both to the content of the sentence and to the models-generated explanations. The benefit of models-generated explanations is discussed in the following section.

False negatives are observed for the category RECOGNITION in all other categories. Reading the explanations, the models seem to consider the hierarchy among languages – which we subsume under RECOGNITION – as a personal or national preference towards one or more languages. In the Luxembourg ideological landscape, the use of the pronoun *our* (*eis*) can indicate a reference to the national in-group. This information is provided as part of the IDENTITY category description. Therefore, also when not carrying an IDENTITY ideology, first-person plural pronouns and mentions of the ingroup are used as markers of national identity (Magistral and DeepSeek). Overlap with VITALITY arises when sentences describe the limited use of Luxembourgish in the country (Magistral, GPT-5, DeepSeek, Qwen). RECOGNITION is then mislabeled as (social) BELONGING when Luxembourgish is stated as the main language of the country (DeepSeek and Qwen). Furthermore, comments perceived as blaming the current language situation on actors mentioned in the text are labeled as RESPONSIBILITY (Qwen).⁵

For RESPONSIBILITY, it is mislabeled as BELONGING when the responsible actor is an out-group member (Qwen, GPT-5), and as RECOGNITION when it is an official institution or a private person (Qwen). It is then mislabeled as VITALITY when the implications of language decay are considered

⁵Further examples and statistics are provided in the appendix.

Label	DeepSeek			GPT-5			Magistral			Qwen		
	Precision	Recall	F1	Precision	Recall	F1	Precision	Recall	F1	Precision	Recall	F1
Belonging	0.468	0.603	0.527	0.541	0.492	0.516	0.466	0.329	0.386	0.365	0.468	0.410
Identity	0.338	0.526	0.411	0.504	0.362	0.421	0.216	0.612	0.319	0.341	0.459	0.391
Recognition	0.502	0.218	0.304	0.479	0.608	0.536	0.447	0.092	0.153	0.458	0.348	0.395
Responsibility	0.603	0.242	0.345	0.689	0.211	0.323	0.667	0.190	0.296	0.368	0.295	0.328
Vitality	0.554	0.375	0.447	0.481	0.611	0.538	0.435	0.382	0.407	0.278	0.625	0.385
None	0.858	0.952	0.903	0.908	0.950	0.929	0.854	0.950	0.899	0.926	0.864	0.894

Table 6: Comparison of models per category

Label	DeepSeek	GPT-5	Highest
Belong.	0.487	0.512	0.571 (EN)
Identity	0.376	0.386	0.475 (FR)
Recogn.	0.330	0.520	0.565 (FR)
Respons.	0.398	0.299	0.398 (LU)
Vitality	0.448	0.500	0.555 (EN)
None	0.910	0.927	0.934 (EN)

Table 7: Per-class F1-scores of DeepSeek and GPT-5. The first two columns show F1-scores for DeepSeek and GPT-5 respectively, on the **Luxembourgish** dataset. The third column reports the absolute best F1 scores per label, regardless of the language of the dataset.

to be the main ideology of the item (Qwen, GPT-5, Magistral). Additionally, the use of first-person pronouns is reported in the model-generated explanations to justify the `IDENTITY` label (Magistral, DeepSeek).

`VITALITY` is mislabeled as `IDENTITY` when language decay is associated with national decay (DeepSeek), and through the use of first-person plural pronouns (Qwen). According to the model-generated explanations, items encouraging the use of the national language as a preservation mechanism are annotated as `RECOGNITION` (GPT-5).

`BELONGING` is mislabeled as `IDENTITY` when Luxembourgish is stated as the language of society (Magistral, Qwen). The presence of the in-group in these utterances makes it difficult for the models to discriminate between the national language as a characteristic of the nation self-determination – which we subsume under `IDENTITY` – and as a means of social exclusion – which we subsume under `BELONGING`.

Finally, `IDENTITY` is mislabeled as `RECOGNITION` when the model (GPT-5) fails to disambiguate between the use of language to create and reinforce a national feeling – which we subsume under `IDENTITY` – and the description of language use in the country – which we subsume under `RECOGNITION`.

5.3. Examining model-generated explanations

The language ideology definitions developed by the two annotators were given to the LLMs as part of the prompt. However, these guidelines were not always efficient, showing that defining language ideologies for LLMs introduces different challenges compared to human annotators. For instance, since language ideologies might share both topics and linguistic features to express them, the annotators already faced (and solved) cases of overlapping categories. Despite providing specific guidelines, these cases are still error sources for LLMs – together with unexpected ones. In this context, the explanations provided by the LLMs are particularly valuable, as they help us to understand how definitions might cause mislabeling, and how clearly delineated the categories are.

Table 8 provides an example of a model-generated explanation. The annotated item is shown in its English translation, and is followed by the labels assigned by the evaluated models and the explanations. The `IDENTITY` label is justified by the use of the first-person plural pronoun *we* (*mir*) as a reference to the national group, and by the use of Luxembourgish in opposition to German and French. The `VITALITY` explanations show that inciting people to use more Luxembourgish words is considered a strategy against language decay. This example was annotated as `RECOGNITION` as it shows a desired hierarchy of language use opposed to the perceived status quo. However, ambiguity is induced by the exhortation of using Luxembourgish *more*, that adds a hint to another ideological layer – `VITALITY` – to the language hierarchy. Similar examples caused label overlaps also for human annotators, but are already disentangled in the prompt given to the models.

Language Ideology Definitions In general, clues which seem helpful for human annotators resulted in confusion between categories for the models. For instance, the use of specific pronouns and determiners is a useful hint for human annotators to discriminate between `IDENTITY` and other categories, but it causes an over-prediction of this category in the models – as shown in the example

Dear – We should use more Luxembourgish words, and not German words, nor French words.
Gold standard: RECOGNITION
Magistral : IDENTITY 'The sentence refers to the national group using the first-person plural pronoun ('mir') and emphasizes the use of Luxembourgish language in opposition to other languages (German and French). This aligns with the national identity category, as it highlights the importance of Luxembourgish language for Luxembourgers.'
DeepSeek: VITALITY 'The sentence emphasizes the need to use Luxembourgish words instead of German or French words, which aligns with the 'vitality' category as it addresses the preservation of Luxembourgish language.'
Qwen: VITALITY 'The sentence expresses concern about the use of German and French words instead of Luxembourgish words, which is a typical indicator of language decay and threat to the vitality of Luxembourgish. The advocacy for using more Luxembourgish vocabulary over foreign alternatives directly relates to preserving the language's integrity and resisting linguistic erosion.'
GPT-5: VITALITY 'The speaker urges using more Luxembourgish words instead of German or French ones, proposing a way to safeguard the language against lexical replacement; this fits concerns and remedies typical of vitality.'

Table 8: Examples of models-generated explanations. The first line reports the annotated sentence, followed by labels and explanations by each model

above. Then, for RESPONSIBILITY, human annotators identified a subject of responsibility, as well as an object. Three different subjects of responsibility have been found – in-group, politician, out-group – and given to the model together with the description of the category. From the explanations, it seems that the presence of the three different responsibility agents leads the tested models to mislabel the category as RECOGNITION – official institutions as subject of responsibility – and BELONGING – the out-group as subject of responsibility.⁶

Language Ideology Labels The analysis in the previous paragraphs suggests that the five language ideology labels might be too fine-grained for reliable prediction with simple in-context learning. Therefore, to answer research question 3: LLMs require expert support or task-specific model adaptation to adequately inform sociolinguistic research. This raises the question whether a more coarse-grained classification scheme would be beneficial. At the same time, the example also demonstrates how LLMs could serve as a valuable tool to improve human annotations and improve them, especially in cases of doubt.

For instance, both IDENTITY and BELONGING, although referring to two different social groups, convey feelings of social belonging obtained and constructed through language use. However, these ideologies highlight a power imbalance that will be lost if the categories are merged, i.e., the dominant

⁶See Gal and Irvine (2019) for more information about language ideologies definitions.

group constructs its own identity – IDENTITY – and shapes the role of other people in society – BELONGING. Similarly, statements about language preference could be grouped under one label, but it will cause losing implications between IDENTITY – personal preference – and RECOGNITION – reference to language status and hierarchy among languages. Therefore, restructuring the categorization scheme could mean losing important aspects of societal power dynamics for sociolinguistic analyses.

6. Conclusions

This paper offers an evaluation of LLMs for language ideology detection. It shows that Luxembourgish does not need to be translated to higher-resourced languages in order to obtain satisfactory classification performance with generative LLMs. It highlights the crucial role of model-generated explanations for this task and that the evaluated models successfully discriminate items with language ideologies from those without. However, discriminating among the different categories of ideologies turns out to be more challenging. The first research question of the paper has thus both a positive and a negative answer: current generations of LLMs can perform a binary language ideology classification, but a fine-grained classification seems to be a more challenging task. Although it might not help with further understanding language ideologies and language discourse in sociolinguistic research, a binary classification between content with and without language ideologies already seems to be a quite useful task, mainly for selecting material to

be analyzed. Our approach could be repurposed for other less-represented languages in NLP, to test if translations of original material could benefit the automatic analysis of language ideologies on the source language. However, given the complexity of the ideology classification task, this would quite likely result in an overview of the presence/absence of language ideologies in a discourse and, hence, require an in-depth analysis from trained (socio)linguists to understand how the motifs are implemented and whether additional language ideologies characterize the discourse in question.

As for the nature of the task, language ideologies do not have a unique linguistic encoding: they may share topical vehicles, the context of utterance, as well as grammatical and lexical features, and are quite challenging for human annotators to discriminate. As including human annotators guidelines in the prompt design turned out to be of limited usefulness for LLMs, new paths should be considered. One option, discussed above, could be the simplification of the annotation scheme. Another option could consist in collecting more manually annotated data to enable fine-tuning. The challenge will be to find a balance between sociolinguistic interests, LLM capabilities and human annotation efforts.

Although LLMs show promise as tools for sociolinguistic inquiry, their potential for the study of language ideologies remains limited without the guidance of human experts. Mainly, language ideologies cannot be studied only using LLMs with no previous understanding of the discourse in analysis – as demonstrated in the trial phase. Mixed methods seem to be a promising way to enrich sociolinguistic research with quantitative methods and NLP advancements (Kircher and Hawkey, 2022; Nguyen et al., 2016; Purschke, 2025).

In future experiments, we will focus on providing more general descriptions instead of detailed, example-based ones. We will improve them to highlight both core and specific elements per category, to give space to infrequent ideological expressions, and to reframe features frequently leading the models to mislabel certain categories.

Limitations

As the project aimed to understand how two disciplines – sociolinguistics and natural language processing – can be mutually informative for each other, we are aware of limitations from both fields. As per the goal of the project, keyword searching should be adjusted to include language ideologies concerning Luxembourgish. Expanding the keyword search to include the different languages spoken in the country would help to get a more comprehensive view of language ideologies in the country. Additionally, we refer to the LLM-generated annota-

tions to understand what could be improved in the prompt design, but we are aware of their limitations in accurately describing the reasoning behind the classification (Randl et al., 2025; Di Bonaventura et al., 2024). Concerning the choice of models, two main limitations can be reported. First, models with limited reasoning capabilities proved to be not robust enough for this task and will be discarded for future experiments. Second, since we only had access to a small amount of annotated data, we were unable to train or fine-tune models on the task. Future work should explore this direction, as improved model performance could support language discourse understanding, as well as its societal implications. However, although a fine-tuned model could achieve high performances, human annotators and linguistic experts remain essential to understand the consistency of model outputs.

We noticed a grammatical mistake in the prompt used to preform this classification: for IDENTITY we state that it is expressed by the use of the first-person plural pronoun *eis*. However, *eis* is the dative form of the first-person plural pronoun *mir*, and its possessive determiner. As this could be the cause of some labeling errors, we will implement the task further with the correct version of the prompt.

Ethical Statement

Data used for the project are provided and anonymized by the data owner RTL.lu.

Acknowledgments

This research was supported by the Luxembourg National Research Fund (Project code: C22/SC/117225699). We would like to thank the members of the TRAVOLTA project and the Language and Technology Group of the University of Oslo for their invaluable advice on the project, and Raul Ian Sosa for his precious feedback.

Bibliographical References

- Meta AI. 2025a. [Llama 4 scout 17b instruct](#).
- Mistral AI. 2025b. [Magistral](#). *arXiv preprint arXiv:2506.10910*.
- Mikel Artetxe, Vedanuj Goswami, Shruti Bhosale, Angela Fan, and Luke Zettlemoyer. 2023. [Revisiting machine translation for cross-lingual classification](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6489–6499, Singapore. Association for Computational Linguistics.

- Paul Baker and Erez Levon. 2015. [Picking the right cherries? A comparison of corpus-based and qualitative analyses of news articles about masculinity](#). *Discourse & Communication*, 9(2):221–236.
- Jan Blommaert, editor. 2010. *Language Ideological Debates*, originally published 1999 edition. Number 2 in Language, Power and Social Process. De Gruyter, Berlin.
- Thomas Blume. 2003. *Ideologie*. Vandenhoeck & Ruprecht.
- Maarten Buyl, Alexander Rogiers, Sander Noels, Guillaume Bied, Iris Dominguez-Catena, Edith Heiter, Iman Johary, Alexandru-Cristian Mara, Raphaël Romero, Jeffrey Lijffijt, and Tijn De Bie. 2025. [Large Language Models Reflect the Ideology of their Creators](#). ArXiv:2410.18417 [cs].
- Kai Chen, Zihao He, Jun Yan, Taiwei Shi, and Kristina Lerman. 2024. [How Susceptible are Large Language Models to Ideological Manipulation?](#)
- Adam Dahlgren Lindström, Leila Methnani, Lea Krause, Petter Ericson, Íñigo Martínez de Rituerto de Troya, Dimitri Coelho Mollo, and Roel Dobbe. 2025. [Helpful, harmless, honest? Sociotechnical limits of AI alignment and safety through reinforcement learning from human feedback](#). *Ethics and Information Technology*, 27(2):28.
- John Dang, Shivalika Singh, Daniel D'souza, et al. 2024. [Aya Expanse: Combining research breakthroughs for a new multilingual frontier](#). *arXiv preprint arXiv:2412.04261*.
- Jwala Dhamala, Tony Sun, Varun Kumar, Satyapriya Krishna, Yada Pruksachatkun, Kai-Wei Chang, and Rahul Gupta. 2021. [BOLD: Dataset and Metrics for Measuring Biases in Open-Ended Language Generation](#). In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 862–872. ArXiv:2101.11718 [cs].
- Chiara Di Bonaventura, Lucia Siciliani, Pierpaolo Basile, Albert Meroño-Peñuela, and Barbara McGillivray. 2024. Is explanation all you need? an expert survey on llm-generated explanations for abusive language detection. In *Proceedings of the 10th italian conference on computational linguistics (CLiC-it 2024)*, pages 280–288.
- John Duncan. 2024. [Does ChatGPT have sociolinguistic competence?](#) *Journal of Computer-Assisted Linguistic Research*, 58:123–145.
- Fernand Fehlen, Peter Gilles, Louis Chauvel, Isabelle Pigeron-Piroth, Yann Ferro, and Etienne Le Bihan. 2023. *Une diversité linguistique en forte hausse*. STATEC, Luxembourg, Luxembourg.
- Susan Gal and Judith T. Irvine. 2019. *Signs of Difference: Language and Ideology in Social Life*. Cambridge University Press.
- Peter Gilles. 2015. *From Status to Corpus: Codification and Implementation of Spelling Norms in Luxembourgish*. Palgrave Macmillan UK.
- Peter Gilles. 2019. [39. Komplexe Überdachung II: Luxemburg. Die Genese einer neuen Nationalsprache](#). In Joachim Herrgen and Jürgen Erich Schmidt, editors, *Sprache und Raum - Ein internationales Handbuch der Sprachvariation. Volume 4 Deutsch*, pages 1039–1060. De Gruyter Mouton, Berlin, Boston.
- Stefan Grondelaers and Roeland van Hout. 2011. The standard language situation in the Netherlands. In Tore Kristiansen & Nik Coupland, editor, *Standard languages and language standards in a changing Europe*, page 113–118. Novus.
- Suchin Gururangan, Dallas Card, Sarah K. Dreier, Emily K. Gade, Leroy Z. Wang, Zeyu Wang, Luke Zettlemoyer, and Noah A. Smith. 2022. [Whose Language Counts as High Quality? Measuring Language Ideologies in Text Data Selection](#). ArXiv:2201.10474 [cs].
- Kristine Horner and Jean Jacques Weber. 2008. [The language situation in luxembourg](#). *Current Issues in Language Planning*, 9:69–128.
- Dirk Hovy and Shrimai Prabhumoye. 2021. [Five sources of bias in natural language processing](#). *Lang. Linguistics Compass*, 15(8).
- Tiancheng Hu, Yara Kyrychenko, Steve Rathje, Nigel Collier, Sander Van Der Linden, and Jon Roozenbeek. 2024. [Generative language models exhibit social identity biases](#). *Nature Computational Science*, 5(1):65–75.
- Ruth Kircher and James W. Hawkey. 2022. [Mixed-methods approaches to the study of language attitudes](#). In Ruth Kircher and Lenore A. Zipp, editors, *Research Methods in Language Attitudes*, pages 330–345. Cambridge University Press.
- Paul V. Kroskrity. 2022. [5 Language ideologies and social identities](#). In Svenja Völkel and Nico Nassenstein, editors, *Approaches to Language and Culture*, pages 101–126. De Gruyter.
- Luyang Lin, Lingzhi Wang, Jinsong Guo, and Kam-Fai Wong. 2024. [Investigating Bias in LLM-Based](#)

- [Bias Detection: Disparities between LLMs and Human Perception](#). ArXiv:2403.14896 [cs].
- Xinru Lin and Luyang Li. 2025. Implicit bias in LLMs: A survey. *arXiv preprint arXiv:2503.02776*.
- Rosina Lippi-Green. 2012. *English with an accent: Language, ideology and discrimination in the United States*. Routledge.
- Aixin Liu, Bei Feng, Bing Xue, et al. 2024. [Deepseek-v3 technical report](#). *arXiv preprint arXiv:2412.19437*.
- Anne-Marie Lutgen, Alistair Plum, Christoph Purschke, and Barbara Plank. 2025. [Neural text normalization for Luxembourgish using real-life variation data](#). In *Proceedings of the 12th Workshop on NLP for Similar Languages, Varieties and Dialects*, pages 115–127, Abu Dhabi, UAE. Association for Computational Linguistics.
- Bolei Ma, Xinpeng Wang, Tiancheng Hu, Anna-Carolina Haensch, Michael A. Hedderich, Barbara Plank, and Frauke Kreuter. 2024. [The Potential and Challenges of Evaluating Attitudes, Opinions, and Values in Large Language Models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 8783–8805, Miami, Florida, USA. Association for Computational Linguistics.
- Christopher D. Manning, Kevin Clark, John Hewitt, Urvashi Khandelwal, and Omer Levy. 2020. [Emergent linguistic structure in artificial neural networks trained by self-supervision](#). *Proceedings of the National Academy of Sciences*, 117(48):30046–30054.
- Dong Nguyen, A. Seza Dođruöz, Carolyn P. Rosé, and Franciska de Jong. 2016. [Computational sociolinguistics: A survey](#). *Computational Linguistics*, 42(3):537–593.
- OpenAI. 2024. [Gpt-4o mini explained: Fast, affordable ai model with 128k context](#).
- OpenAI. 2025. [Gpt-5 system card](#). Technical report, OpenAI. Accessed: 2025-10-20.
- OpenAI. 2025. [Gpt-o3 model overview](#).
- OpenAI et al. 2023. [Gpt-4 technical report](#). *arXiv preprint arXiv:2303.08774*.
- OpenAI et al. 2025. [gpt-oss-120b & gpt-oss-20b model card](#). *arXiv preprint arXiv:2508.10925*.
- Fred Philippy, Shohreh Haddadan, and Siwen Guo. 2024. [Forget NLI, use a dictionary: Zero-shot topic classification for low-resource languages with application to Luxembourgish](#). In *Proceedings of the 3rd Annual Meeting of the Special Interest Group on Under-resourced Languages @ LREC-COLING 2024*, pages 97–104, Torino, Italia. ELRA and ICCL.
- Alistair Plum, Caroline Döhmer, Emilia Milano, Anne-Marie Lutgen, and Christoph Purschke. 2024. [LuxBank: The first Universal Dependency treebank for Luxembourgish](#). In *Proceedings of the 22nd Workshop on Treebanks and Linguistic Theories (TLT 2024)*, pages 30–39, Hamburg, Germany. Association for Computational Linguistics.
- Alistair Plum, Tharindu Ranasinghe, and Christoph Purschke. 2025. [Text generation models for Luxembourgish with limited data: A balanced multilingual strategy](#). In *Proceedings of the 12th Workshop on NLP for Similar Languages, Varieties and Dialects*, pages 93–104, Abu Dhabi, UAE. Association for Computational Linguistics.
- Edoardo Maria Ponti, Julia Kreutzer, Ivan Vulić, and Siva Reddy. 2021. [Modelling latent translations for cross-lingual transfer](#).
- Christoph Purschke. 2023. Diskurs-Figuren. Wie Politik und Öffentlichkeit in Luxemburg über Sprache sprechen. *Hemecht: Zeitschrift für Luxemburger Geschichte*, (2023/3):313–330.
- Christoph Purschke. 2025. [Discourse figures in the luxembourg language debate \(2015–2020\)](#). *Journal of Language Variation and Sociolinguistics*, 1(2):37–53.
- Christoph Purschke and Peter Gilles. 2023. Sociolinguistics in Luxembourg. In *The Routledge handbook of sociolinguistics around the world*, pages 542–549. Routledge.
- Tharindu Ranasinghe, Alistair Plum, Christoph Purschke, and Marcos Zampieri. 2023. [Publish or hold? Automatic comment moderation in Luxembourgish news articles](#). In *Proceedings of the 14th International Conference on Recent Advances in Natural Language Processing*, pages 968–978, Varna, Bulgaria. INCOMA Ltd., Shoumen, Bulgaria.
- Korbinian Randl, John Pavlopoulos, Aron Henriksson, and Tony Lindgren. 2025. Mind the gap: from plausible to valid self-explanations in large language models. *Machine Learning*, 114(10):220.
- Catherine R Rhodes. 2023. [Language ideologies](#). In *Oxford Research Encyclopedia of Anthropology*. Oxford University Press.
- Genevieve Smith, Eve Fleisig, Madeline Bossi, Ishita Rustagi, and Xavier Yin. 2025. [Standard Language Ideology in AI-Generated Language](#). ArXiv:2406.08726 [cs].

Rachelle Vessey. 2017. [Corpus approaches to language ideology](#). *Applied Linguistics*, 38(3):277–296.

Julia Watson, Sophia S. Lee, Barend Beekhuizen, and Suzanne Stevenson. 2025. [Do language models practice what they preach? Examining language ideologies about gendered language reform encoded in LLMs](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 1201–1223, Abu Dhabi, UAE. Association for Computational Linguistics.

Kathryn A. Woolard. 2020. [Language Ideology](#). In James Stanlaw, editor, *The International Encyclopedia of Linguistic Anthropology*, 1 edition, pages 1–21. Wiley.

A Wu, Ryan Liu, Kerem Oktar, T Sumers, and T Griffiths. 2025. Are large language models sensitive to the motives behind communication? *Advances in Neural Information Processing Systems (NeurIPS)*, 38.

An Yang, Anfeng Li, Baosong Yang, et al. 2025. [Qwen3 technical report](#). *arXiv preprint arXiv:2505.09388*.

Zhihao Zhang, Jun Zhao, Qi Zhang, Tao Gui, and Xuanjing Huang. 2024. [Unveiling linguistic regions in large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6228–6247, Bangkok, Thailand. Association for Computational Linguistics.

Appendix A

In this appendix, we report the original Luxembourgish version of the examples in Section 2.2:

- **Example ID 636:** 1. Lëtzebuergesch ass fir eis Lëtzebuurger (250'000 Leit) eis éischt Sprooch.
- **Example ID 3631:** Gudden Text, mee ech färten, eis Sprooch wärt ausstiewen.
- **Example ID 410165:** et ass geschafft: All Awunner vu Lëtzebuerg hu sech sproochlech integréiert.
- **Example ID 666:** A wann eis Frontalieren kee Letzebuergesch schwaetzen dann ass daat eis eege schold.
- **Example ID 636:** 2. Lëtzebuergesch ass VIRUN ALLEM eng geschwaate Sprooch.

Appendix B

Keywords used to retrieve the comments:

'eis sprooch', 'eis sproch', 'ons sprooch', 'ons sproch', 'lëtzebuergesch', 'lëtzeboiesch', 'letzebuergesch', 'lëtzeburgesch', 'lëtzbuergesch', 'letzeburgesch', 'letzeboiesch', 'lëtzebuergesch', 'eiser sprooch', 'onser sprooch', 'eiser sproch', 'onser sproch', 'lëtzebuergesch', 'mammesprooch', 'mammesproch'.

Appendix C

In this appendix, we add the complete version of Prompt 4:

Detect one or more of the following categories in the sentences of the comment below, after the instructions and schemas. The categories are: Identity, Vitality, Belonging, Responsibility, Recognition. These categories concern language ideologies and only these exact categories can be used for the annotation. No overlap between categories is allowed, each sentence can have 0 or 1 category. If you cannot identify one of the category assign N/A. Do not take any quoted parts (i.e. text in quotation marks) into account for the labeling. For each sentence, explain why you have assigned that label. Here are the explanations of the five categories:

Identity: The language related ideology 'identity' groups both national and personal identity. National identity: it refers to Luxembourgish national identity and it only concerns Luxembourgish people, born in Luxembourg with Luxembourgish as native language. It is characterized by a strong emotional attachment to the language and culture of Luxembourg, in opposition to people with different languages and cultures. It is found in clauses or part of clauses underling:

- the tied connection between national identity and national language (example: Luxembourgishers unified by a common national language)
- the tied connection between Luxembourgish culture and Luxembourgish language (example: importance of Luxembourgish language and culture)
- the formation of the Luxembourgish national identity during the Second World War (Luxembourgishers explicitly chose Luxembourgish language as national language for the first time)
- the opposition between the national group (= born and raised in Luxembourg, by family of Luxembourgish origins, with Luxembourgish culture and Luxembourgish as native language) and the group seen as external to the

country (= not born in Luxembourgish families, or coming from abroad, or living across the borders but going to Luxembourg for work. These people are grouped by lack of knowledge of Luxembourgish language and refusal to learn it)

Grammatically, this category is expressed by the use of the first-person plural pronoun ('eis') when referring:

- to the national group
- to what belongs to the national group
- to what characterizes the national group.

'eis' is used in opposition to the third person plural pronoun ('si'). The third person plural pronoun refers to the group seen external to the country. Personal identity = personal preference towards a language with no reference to any group-shared ideology or feeling (examples = speaking Luxembourgish as personal preference and not as a national identity trait)

examples:

1. lëtzebuergesch ass fir eis lëtzebuerger (250'000 leit) eis éischt sprooch.
2. et gött leit de zu dachau agespaart ware wéint eiser sprooch an e vollek dat seng sprooch verléiert verléiert och seng identitéit. . .
3. ech perséinlech maan mer emmer en spaass draus vir mat de verschiddenen wierder ze spillen wann ech eppes schreiwen

examples where this label is not applicable:

1. tschüss!
2. mee mir kann et dee moment egal sinn, well ech da wahrscheinlech an däitschland wunnen... tschüss!
3. patricia courtois

Vitality: The language related ideology 'vitality' groups opinions about Luxembourgish language considered:

- endangered
- threatened
- in decay
- soon to be replaced by other languages

Threatening factors are:

- languages other than Luxembourgish spoken in the country

- impossibility of speaking Luxembourgish in everyday contexts (shops, hospitals, ...)
- languages other than Luxembourgish used at official levels
- lack of use or incorrect use of Luxembourgish orthographical and grammatical norms
- use of non-Luxembourgish words instead of available Luxembourgish words
- scarcity of Luxembourgish speaking people

Proposing ways to preserve and safeguard Luxembourgish is also grouped in this category. The focal element of this category is always the language situation. Even when proposing possible ways to save the language and reporting the role policy-makers have in this, the focus always stays on the language situation and does not shift on secondary elements. There are a few semantic fields typical of this category: decay, endangerment, violence (towards the language).

examples:

1. eng dagesmamm déi lëtzebuergesch schwätzt ze fannen, ass quasi onméiglech. awer eng babá fënnt ee séier
2. soss gi mer all frankophon gemaach
3. ech denken sou seier geht dat net verluer

examples where this label is not applicable:

1. bonjour a vill gléck a neie joer.
2. all guddes an deem sënn.
3. ech verstinn eent net, léif fra.

Belonging: The language related ideology 'belonging' groups answers to the following questions:

- Do people born outside Luxembourg belong in Luxembourg?
- Do people speaking Luxembourgish at a basic level belong to Luxembourg?
- Do people not learning Luxembourgish belong to Luxembourg?

This category is about the integration in Luxembourgish society of people coming from abroad or people of non-Luxembourgish descents. It always involves the presence of foreigners, or of people not speaking Luxembourgish or not being Luxembourgish native speakers. Two main factors are significant in this category: 1) Luxembourgish language, 2) foreigners. These two factors intertwine as follows:

- Luxembourgish language is seen as the only means of integration in Luxembourgish society for foreigners
- Foreigners are expected to learn Luxembourgish
- Foreigners refuse to learn Luxembourgish language
- Luxembourgish language has to be learned at school for foreigners not to be disadvantaged in society
- The required level of Luxembourgish to obtain the citizenship is too low
- Foreigners who learn or try to learn/speak Luxembourgish are welcomed in Luxembourgish society
- Foreigners adapting to Luxembourgish language and culture are part of Luxembourgish society

examples:

1. eischt offiziell sprooch awer onbedengt fir e kloert zeeche fir d'integratioun ze setzen, fir eisen auslännesche matbierger kloer ze verstoen ze ginn dass mir eng eegestänneg denkweis, sprooch a kultur hunn an si sech eis unzepassen hunn andeems si net just 1 joer lëtzebuergesch léieren mee bis si et sou kënnen dass si et och kënne schwätzen a verstoen
2. ween lëtzebuenger wëll ginn muss och lëtzebuergesch schwätzen
3. ganz traureg daat do. wann et engem schlecht geht muss een sech kennen an senger mamesprooch ausdrecken an et huet een et nët néideg sech vun schlecht gelaanten franséischen infirmièren ungranzten ze loossen. déi däitsch maachen jo nach éischer en effort an probéieren lëtzebuergesch ze schwätzen

examples where this label is not applicable:

1. ierch all een schéinen owend
2. där meenung sin ech och.
3. wou soll daat dann hin feieren?

Responsibility: The language related ideology 'responsibility' groups opinions about who is considered responsible for Luxembourgish decay. This category has two main elements: 1) responsible agent, 2) subject of responsibility. The subject of responsibility is the Luxembourgish language decay (everything defined in 'vitality' category) Possible responsible agents are:

- Politicians and policymakers: because they don't defend Luxembourgish language through adequate language policies and requires a low Luxembourgish language level to acquire Luxembourgish citizenship
- Luxembourgers: because they never impose themselves to speak Luxembourgish with foreigners, they never ask foreigners to learn Luxembourgish, but adapt to foreigners with English, French, or German
- Foreigners: because they refuse to learn/speak Luxembourgish and impose the languages they can speak (mainly French and English)

These responsible agents are also seen as responsible for a possible solution to the subject of responsibility. Other expressions fitting in this category involve Luxembourgers defending themselves from being called being racist, like 'we are not racist, but...', 'they call us racist when we advocate for Luxembourgish.'

examples:

1. dann sollen dei ausländesch elteren emol ufänken eis sprooch ze leieren mee dofir sin mer ze faul
2. et ass wéi am palais, do gët och keen lëtzebuergesch geschwaat.sie hun problemer vir sech an eiser sprooch ze artikuléieren
3. et as leider net einfach ze verlaangen, dass soll letzebuergesch geschwaat soll gin, well soss get een nämlech ganz ganz ganz seier asl rassist ugesin. wann dir am geschäft engem soot e soll letzebuergesch schwetzen, da get een einfach lenks leihen gelos an färdeg

examples where this label is not applicable:

1. also muss daat och hei goen.
2. schummt iech!
3. as daat lo wierklech den sprengenden punkt vun desem artikel?

Recognition: The language related ideology 'recognition' groups opinion on Luxembourgish language and other languages spoken in the country. In this category, different ways of acknowledging languages are involved:

- Luxembourgish considered as an oral language or/and a dialect
- Luxembourgish considered as an official, standardized language, whose norms have to be followed by speakers/writers

- Luxembourgish has to be recognized at the European level as one of the official languages of the European Union
- Hierarchy among languages

examples:

1. mee ech verstinn haut nach ëmmer net, wéisou alles, awer och wierklech alles hei zu lëtzebuerg op franséisch leeft.
2. jo mir hun eng nationalsprooch, an déi solle mir och héich halen, mee ons stäerkt war ëmmer, dass mir dräisproocheg waren, an och houfreg drop waren
3. ass wuel net esou einfach, en dialekt forcement zu enger sprooch wëllen ze forméieren...

examples where this label is not applicable:

1. ganz richtig
2. gudd geschriwen
3. bai allem respekt awer dier sidd mengen ech am joer 1945 henken bliwen.

examples where none of the labels described are applicable:

- Ech kennen 2 lëtzebuergesch Mammen déi aleng mat hire Kanner do stin, déi kruten null Ennerstëtzung fir eng Wunneng ze fannen.
- Letzebuerg huet keen mobiliteitsproblem ausser moies an owes zu den Spetzenstonnen.
- Kommt mir lossen eis net mei manipuleiren!
- Do huet de Fotograf awer gutt oppgepass fir emmer Leit opt Bild ze krein ...
- gott sein dank dass mir keng vum "adr" an der spëtz hun soss hätten mir awer wirklech "FOLKLOR"
- Am EU parlament gouf gestemmt: 39 Jo- 478 neen-165 enthalung-72 abwesend nozekucken ob dem eu parlameter vun zdf. do kann een och kucken wien wei ofgestemmt hued. Letzebuergesch Parlamentarier hun all neen gestemmt ausser Mme Lulling.

Appendix D

In this appendix, we add further examples of annotations mislabeled by all the models per category. The example is first reported in Luxembourgish followed by the English translation, its gold standard, and the mislabeled annotation by

the models and model generated explanations.

Example 586

Mir sinn schon dat eenzegt Land wou d'Kanner mussen an enger anerer Sprooch alphabetiséiert ginn wéi hier Mammensprooch, fir den net francophonien=Lëtzebuerg et lo nach méi schwéier ze machen, sollen d'Kanner lo och nach op franséisch alphabetiséiert ginn??

We are already the only country where the children have to be alphabetized in a language other than their mother tongue, to make it even more difficult for non-francophone=Luxembourgers, should the children be alphabetized in French??

Gold standard: RESPONSIBILITY

Magistral : BELONGING

'The sentence discusses the integration of children from different linguistic backgrounds into Luxembourgish society, focusing on the role of Luxembourgish language in this process. It involves the presence of foreigners and the expectation to learn Luxembourgish.'

DeepSeek: VITALITY

'The sentence emphasizes the need to use Luxembourgish words instead of German or French words, which aligns with the 'vitality' category as it addresses the preservation of Luxembourgish language.'

Qwen: VITALITY

'The sentence expresses concern that children in Luxembourg are being alphabetized in a language other than their mother tongue (Luxembourgish), and further that non-French-speaking Luxembourgish children are being forced to learn in French, which is portrayed as making the situation even harder. This reflects a perception of Luxembourgish being threatened by the dominance of other languages in education — a core concern of the 'vitality' category, as it highlights endangerment through institutional language use.'

GPT-5: RECOGNITION

'The sentence critiques alphabetizing children in a language different from their mother tongue and questions doing it in French, framing Luxembourgers as non-Francophone; this is an opinion about which languages should be used in official/educational contexts and the dominance of French.'

Example 339

Si sollen mol zu alleréischt all Lëtzebuergesch léieren.

First of all, they should all learn Luxembourgish.

Gold standard: BELONGING

Magistral : IDENTITY

'The sentence emphasizes the importance of learning Luxembourgish first, which is a personal identity trait.'

DeepSeek: RESPONSIBILITY

'The sentence emphasizes the responsibility of

teaching Luxembourgish first, highlighting the importance of prioritizing Luxembourgish language education.'

Qwen: IDENTITY

'The sentence emphasizes that children should learn Luxembourgish 'at first' ('zu alleréischt'), framing it as the foundational and primary language of the national group. The use of 'si sollen' (they should) implies a collective norm tied to national identity Luxembourgish as the essential, first language of Luxembourgers, reinforcing the tied connection between national identity and national language.'

GPT-5: VITALITY

'Advocates that all children first learn Luxembourgish, emphasizing increased acquisition and priority of Luxembourgish as a way to strengthen and preserve the language.'

Example 356

Jo mir hun eng Nationalsprooch, an déi solle mir och héich halen, mee ons Stäerkt war ëmmer, dass mir dräisproocheg waren, an och houfreg drop waren.

Yes, we have a national language, and we should hold it in high esteem, but our strength has always been that we were trilingual, and we were also proud of it.

Gold standard: IDENTITY

Magistral : RECOGNITION

'The sentence acknowledges Luxembourgish as a national language and discusses the strength of being multilingual, which falls under the recognition category.'

DeepSeek: RECOGNITION

'The sentence acknowledges Luxembourgish as the national language and emphasizes the importance of maintaining it, fitting the 'Recognition' category.'

Qwen: RECOGNITION

'The sentence acknowledges Luxembourgish as a 'Nationalsprooch' that should be 'héich gehalten' (held high), while also recognizing the value of trilingualism ('dräisproocheg'). This reflects a hierarchical acknowledgment of languages — affirming Luxembourgish as official and worthy of prestige, while accepting coexistence with other languages. This fits the 'recognition' category, which deals with status, standardization, and hierarchy of languages.'

GPT-5: RECOGNITION

'It explicitly discusses the status of Luxembourgish as a national language and values Luxembourg's trilingualism, i.e., views on languages and their place in the country.'

Example 678

Gleichzaiteg gehéiert Lëtzebuergesch zu onser Kultur a muss eis erhaale bleiwen!

GLikewise Luxembourgish belongs to our culture and must be preserved!

Gold standard: VITALITY

Magistral : IDENTITY

'The sentence emphasizes the importance of Luxembourgish to the national culture and identity, which is characteristic of the 'Identity' category.'

DeepSeek: IDENTITY

'The sentence emphasizes the importance of Luxembourgish to Luxembourg's culture, fitting the 'Identity' category as it highlights a connection between language and national identity.'

Qwen: IDENTITY


'The sentence asserts that Luxembourgish belongs to 'our' culture and must be preserved, using the first-person plural 'eis' (our) to express collective emotional attachment to the language as a core element of national identity. This reflects the tied connection between national identity and national language, fitting the 'identity' category.'

GPT-5: IDENTITY


'It explicitly links Luxembourgish to 'our culture' and uses first-person plural to assert a cultural-national attachment, a hallmark of the identity category.'

Appendix E


In this appendix, we add row normalized confusion matrices for the Lang subset and for the whole dataset in analysis (Lang and NotLang) in its original language (Luxembourgish) for the four analyzed models: GPT-5, DeepSeek, Magistral, and Qwen.



(a) Lang subset



(b) Lang and NotLang subsets




(a) Lang subset




(b) Lang and NotLang subsets

Figure 3: Confusion matrix per annotation with GPT-5

Figure 4: Confusion matrix per annotation with DeepSeek



(a) Lang subset




(b) Lang and NotLang subsets

Figure 5: Confusion matrix per annotation with Magistral



(a) Lang subset



(b) Lang and NotLang subsets

Figure 6: Confusion matrix per annotation with Qwen