

Meenz bleibt Meenz, but Large Language Models Do Not Speak Its Dialect

Minh Duc Bui[▽] Manuel Mager[▽]

Peter Herbert Kann[♣] Katharina von der Wense^{▽,♣}

[▽] Johannes Gutenberg University Mainz, Germany [♣] Marburg University, Germany

[♣] University of Colorado Boulder, USA

minhducbui@uni-mainz.de

Abstract

Meenzerisch, the dialect spoken in the German city of Mainz, is also the traditional language of the Mainz carnival, a yearly celebration well known throughout Germany. However, *Meenzerisch* is on the verge of dying out—a fate it shares with many other German dialects. Natural language processing (NLP) has the potential to help with the preservation and revival efforts of languages and dialects. However, so far no NLP research has looked at *Meenzerisch*. This work presents the first research in the field of NLP that is explicitly focused on the dialect of Mainz. We introduce a digital dictionary—an NLP-ready dataset derived from an existing resource (Schramm, 1966)—to support researchers in modeling and benchmarking the language. It contains 2,351 words in the dialect paired with their meanings described in Standard German. We then use this dataset to answer the following research questions: (1) Can state-of-the-art large language models (LLMs) generate definitions for dialect words? (2) Can LLMs generate words in *Meenzerisch*, given their definitions? Our experiments show that LLMs can do *neither*: the best model for definitions reaches only 6.27% accuracy and the best word generation model's accuracy is 1.51%. We then conduct two additional experiments in order to see if accuracy is improved by few-shot learning and by extracting rules from the training set, which are then passed to the LLM. While those approaches are able to improve the results, accuracy remains below 10%. This highlights that additional resources and an intensification of research efforts focused on German dialects are desperately needed.

Keywords: dialect research, definition generation, word generation, large language models, low-resource languages

1. Introduction

Most German dialects are heavily endangered, as speakers face dialect-related discrimination (Wirtz and Ender, 2025) and the homogenizing effects of top-down language standardization (Rutten and Vosters, 2021). This has contributed to a wider underappreciation of dialects, not only among non-speakers and international audiences, but occasionally among speakers themselves. Yet, language variants are a huge part of each region's cultural identity, and speakers who cherish their local cultural traditions fear the looming loss of their dialect.

Language technology has the potential to aid the preservation of minority languages, including dialects (Galla, 2016). Furthermore, natural language processing (NLP) systems able to handle dialects would allow speakers to interact with contemporary language technology in a natural way (Mager et al., 2018). However, since dialects are mostly used in speech, little dialectal text data is available. This is why large language models (LLMs) have been shown to perform poorly for many German dialects (Kantharuban et al., 2023; Peng et al., 2024; Blaschke et al., 2025; Bär et al., 2025; Muñoz-Ortiz et al., 2025).

Here, we present—to the best of our knowledge—

the first study investigating the ability of LLMs to handle a dialect that has so far been overlooked by the NLP community: the dialect of the city of Mainz—called *Meenzerisch* or *Määnzerisch* by the speakers themselves.¹ This dialect, while being largely unknown outside Germany, is the language of the carnival of Mainz (*Meenzer Fassenacht* in the dialect), and carnival speeches in the dialect of Mainz are broadcast on German national television once a year. Traditionally, those speeches were given by native speakers. However, in recent years, more and more presenters have instead been native speakers of Standard German, highlighting the dialect's risk of decline.

As a first step towards supporting *Meenzerisch* with NLP, we introduce a novel dataset resulting from semi-automatically digitizing a physical dictionary of *Meenzerisch*. It contains dialectal words together with a description of their meaning in Standard German. We further use this dataset to investigate the level of understanding of LLMs by tasking them to generate the meaning of the given words. To explore their generation capabilities, we also prompt them for the opposite direction: given a description of a word's meaning, generating the word itself.

¹To explain the title: the English translation of the *Meenzer Meenz bleibt Meenz* is *Mainz remains Mainz*.

We experiment with a diverse set of open-source LLMs, including small, medium-sized, and large models. Our results show that all evaluated LLMs struggle to generate accurate definitions for the dialect words. On average, models achieve only 4.24% accuracy, with the best-performing model, Llama-3.3 70B, reaching just 6.27%. Performance is even lower on the second task—generating the correct dialect word from its meaning—, where the average accuracy drops to 0.56%, and the strongest model, GPT-OSS 120B, attains only 1.51%. While few-shot learning and automatic rule extraction slightly improve results for definition generation (up to 9.24% and 8.37%, respectively), overall performance remains low.

Contributions To summarize, our contributions are as follows: 1) the first dataset containing words in the dialect of Mainz together with their definitions; 2) the first study of LLMs’ abilities to comprehend words in the dialect of Mainz, showing that current models struggle to understand them; and 3) the first study of LLMs’ abilities to generate words in the dialect of Mainz, demonstrating that existing models fail to produce correct dialectal words.²

2. Related Work

German Dialect Datasets Prior work has developed parallel dialect–Standard German dictionaries (Haddow et al., 2013; Artemova and Plank, 2023; Litschko et al., 2025, *inter alia*), but none address the Mainz dialect. Blaschke et al. (2023) compile over 80 German dialect corpora, including Rhine-Franconian varieties, however these do not cover Meenzerisch and consist only of unstructured text corpora.

NLP for German Dialects A recent survey by Blaschke et al. (2023) shows that many German dialect speakers wish to use LLMs in their own dialects, highlighting the need for dialect-aware systems. Yet, studies reveal strong performance gaps between standard and dialectal German varieties across tasks (Kantharuban et al., 2023; Peng et al., 2024; Blaschke et al., 2025; Bär et al., 2025; Muñoz-Ortiz et al., 2025). Efforts to reduce this gap include noise injection during fine-tuning (Peng et al., 2024) and pixel-based modeling that treats text as images (Muñoz-Ortiz et al., 2025). Recent work even demonstrates that LLMs systematically discriminate speakers of German dialects compared to Standard German speakers, indicating significant bias (Bui et al., 2025a). Overall, however, Meenzerisch has so far remained unexplored due to the lack of suitable digital resources.

3. Dataset: A Dictionary for the Dialect of Mainz

3.1. “Meenzerisch”—A German Dialect

Linguistic Classification and Distribution The dialect of the city of Mainz (name of the city in dialect: *Meenz*; also: *Määnz*) in Germany (name of the dialect in Standard German: *Mainzerisch*, in dialect: *Meenzerisch*; also *Määnzerisch*) is a sub-variety of the Rhine-Franconian dialect group of the German language. It is not only spoken in Mainz, but—in several local variations—in the region of Rhine Hesse (also: *Rhenish Hesse*; Standard German: *Rheinhessen*, in dialect: *Rhoihesse*), which comprises the cities of Mainz, Bingen, Alzey and Worms (name of the regional dialect: *Rheinhesisch*, in dialect: *Rhoihessisch*).

Typically, native speakers of Meenzerisch do not distinguish between the phonemes “ch” [ç] and “sch” [ʃ]—everything is pronounced as “sch” [ʃ]. As Standard German does make the distinction, they usually learn at primary school that there is a difference between these phonemes. In Meenzerisch, “wir” (*we*) and “mir” (*me*) are both pronounced as “mir” or “mer”, resulting in identical forms—also a characteristic causing confusion to Meenzerisch speaking children when faced with Standard German as they start kindergarten or primary school.

Historical Influences During recent centuries, Meenzerisch has been strongly influenced by French. Among other historical events not mentioned here, Napoleon Bonaparte made Mainz (in French: *Mayence*) a capital of the department “Mont-Tonnerre” of the French Empire. Typically, in Meenzerisch, words taken over from French are stressed on the first syllable but not on the last as in standard French (Betz, 2010).

Furthermore, Mainz was one of the three important Jewish communities in medieval Germany (SchUM communities with Speyer (Yiddish: *Spira*), Worms (Yiddish: *Warmaisa*) and Mainz (Yiddish: *Magenza*) (Levi, 1927). Thus, over a long period of time, Yiddish and Meenzerisch also influenced each other (Betz, 2010).

Over the last decades, the number of native speakers has been diminishing. Simultaneously, Meenzerisch has been losing unique words and idioms. Neologisms, a typical feature of a living and evolving language, can hardly be observed today (Betz, 2010). Meenzerisch is slowly being absorbed into a so-called regiolect of the Rhine-Main area (Vorberger, 2019).

Sociolinguistic Situation and Cultural Role

The slow attrition of Meenzerisch is due to a va-

²See Appendix A for code and dataset resources.

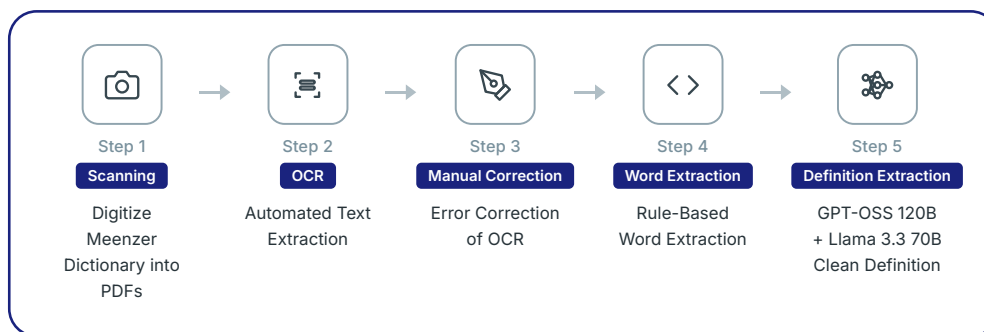


Figure 1: **Dataset Creation Pipeline.** Overview of the semi-automatic five-step process used to create the Mainz Dialect Dataset.

riety of factors which are common causes for the death of many low-resource languages and dialects (see, e.g., [Msila \(2011\)](#) for the situation of isiXhosa). Another factor is the preference for Standard German among parents and educators, driven by the negative stereotypes dialects carry regarding educational status ([Niemann, 1964](#); [Eichinger et al., 2014](#); [Trillhaase, 2021](#)). Furthermore, the widespread presence of mass media and the internet has gradually introduced English and Standard German influences into regional speech. Additionally, social mobility contributes to dialect shift: in linguistically mixed households, the dialect is less likely to be transmitted to the next generation.

Nevertheless, Meenzerisch is still strongly anchored in the local culture of the city of Mainz. It can be characterised as the semi-official language used in the carnival of Mainz (Standard German: *Mainzer Fastnacht*, in dialect: *Meenzer Fassenacht*), a very popular festival starting annually on November 11 and ending on Ash Wednesday. Furthermore, Meenzerisch is the colloquial language in contact with local authorities, in the city center of Mainz, and also in the football stadium of Mainz and by supporters of the local football team.

Meenzerisch is an essential cornerstone of local culture, tradition and lifestyle and is, thus, according to many speakers deserving of preservation efforts: “Ei natürlich! Du glaabst doch nit im Ernst, dass unser Muddersprooch verschwindt! Sowas is Kuldur! Verstehste! Und Kuldur ist wischtsch!” ([Betz, 2010](#)).³

3.2. Dataset Creation

We create a dataset consisting of words in Meenzerisch together with a description of their meaning in Standard German by extracting definitions from a physical copy of the “Mainzer Wörterbuch” (En-

glish: *dictionary of the dialect of Mainz*) by Karl Schramm ([Schramm, 1966](#)). This is done semi-automatically by performing the following five steps: 1) scanning the book; 2) optical character recognition (OCR); 3) manual correction; 4) extraction of the word or expression from each line with a rule-based Python script; and 5) using an LLM to extract the definition from each line. Step 5 is necessary since many entries in the original book contain additional information, which would make it tricky to automatically evaluate an LLM’s understanding or generation skills, e.g., example sentences or words similar to the explained one. We ignore all such additional information for the purpose of this research, but highlight that it is valuable knowledge for future work, for instance when adapting LLMs to the dialect of Mainz. We show an overview of our pipeline in Figure 1.

Scanning and OCR We use a commercial app to create PDFs from pictures taken of the pages of the “Mainzer Wörterbuch”. We use the same app for OCR to convert the PDFs into text files.

Manual Clean-Up of OCR Output While a manual inspection of the OCR output reveals that the quality is generally high, a common mistake is missing line breaks. They are problematic as, in the original book, new words are recognizable by starting a new line. Thus, in order to increase the quality of the subsequent automatic word and definition extraction, we manually add in missing line breaks to make sure each line contains exactly one word, together with its definition. While doing to, we also correct other obvious mistakes. This manual clean-up required approximately four hours.

Automatic Extraction using LLMs To isolate the definition of each word, excluding explanations and examples, we employ an LLM for automated extraction. Specifically, we use GPT-OSS 120B ([OpenAI, 2025](#)) with a medium reasoning effort. The model

³English translation: “Of course! You don’t really believe that our native language will die out, do you? It’s culture! You know? And culture is important!”

Category	Value
<i>Dataset Size</i>	
Total extracted pairs	2,471
Extracted pairs after clean-up	2,351
<i>Word Length</i>	
Median length (characters)	8
<i>Definition Length</i>	
Avg. definitions per word	1.37
Median length (characters)	31
Median length (words)	4
<i>Dataset Split</i>	
Training set	250
Development set	250
Test set	1,851

Table 1: **Overview of the Mainz Dialect Dataset.** Summary of the extracted and cleaned dataset.

is instructed to preserve the original wording and to number multiple meanings when present. We provide six handpicked, manually extracted examples as demonstrations within the prompt (see prompt in Appendix B.1).

Furthermore, we clean each extracted definition from OCR-induced noise, e.g., unnecessary hyphens or special characters, using Llama-3.3 70B (Grattafiori et al., 2024) to automatically normalize and sanitize the text, see Appendix B.1 for the full prompt.

Quality Control To evaluate the quality of our extraction procedure, one of the authors, who is a native speaker of Standard German, manually assesses a random subsample of 100 entries by comparing the extracted model outputs to the original book’s Standard German definitions. We distinguish between major and minor errors. Major errors occur when the extracted meaning does not correspond to the intended definition, whereas minor errors refer to cases where redundant or irrelevant information is included, e.g., the additional tag “Bildwort” (*English: figurative word*). Our evaluation reveals that 12% of the entries contain minor errors, while only 6% exhibit major errors. Given the low rate of major errors, we consider the overall extraction quality satisfactory.

Copyright Considerations We create our corpus according to §60d on Text and Data Mining of the German Urheberrechts-Wissensgesellschafts-Gesetz (*English: law on copyright in the knowledge economy*). This law allows text and data mining with the goal to create corpora for non-commercial scientific research. It further allows to make the resulting corpora accessible to a defined group of people for joint scientific research or for control of the quality of scientific research. Thus, we will

Word in Meenz-erisch	Standard German Definition
<i>Aaweiderworschd</i>	Salzgurke (<i>English: pickled cucumber</i>)
<i>Bitzelwasser</i>	Mineralwasser mit Kohlen-säuregehalt (<i>English: Carbon-ated mineral water</i>)
<i>Mauldabbezierer</i>	scherzhaftes Lob für einen guten Wein [...] (<i>English: playful compliment for a good wine</i>)
<i>Rachebutzer</i>	saurer Wein (<i>English: sour wine</i>)
<i>Schlobb</i>	1. Schleife; 2. Knoten; 3. [...] (<i>English: 1. bow; 2. knot</i>)
<i>Schimmes</i>	Hunger (<i>English: hunger</i>)
<i>Schwollescheer</i>	Angehörigen einer Reitertruppe (<i>English: members of a cavalry troop</i>)

Table 2: **Examples from Our Dataset.** Samples illustrating the mapping between dialect words and their corresponding Standard German definitions.

make our dataset available to other researchers upon request under the CC BY-NC-ND 4.0 license.

3.3. Dataset Statistics

We report a summary of the dataset in Table 1.

Number of Samples After the extraction process, we obtain a total of 2,471 word–definition pairs. We remove samples containing missing or invalid definitions, resulting in a cleaned dataset of 2,351 valid entries. Table 2 shows representative examples from the dataset.

Word and Definition Length The median length of dialect words is 8 characters. Since some definitions include multiple meanings for a single word, we first report that each word contains an average of 1.37 meanings. Meanings have a median length of 31 characters and 4 words.

Dataset Split The dataset is divided into 250 training, 250 development, and 1,851 test samples. We intentionally provide a larger test set than training data, as the main goal of this benchmark is thorough evaluation rather than model training.

4. Large Language Models

We benchmark a diverse set of open-source, instruction-tuned models covering both dense and mixture-of-experts architectures, thinking and non-thinking models, of varying scale and multilingual coverage. Some models are primarily pretrained on human-generated data, while others rely heavily

Model	Dialect of Mainz		English Test Set
	Dev Set	Test Set	
<i>Large Models</i>			
GPT-OSS 120B (Thinking)	5.60%	4.92%	91.19%
Llama-3.3 70B	8.80%	6.27%	91.69%
Leo-HessianAI 70B	4.40%	4.81%	81.88%
Qwen-2.5 72B	4.80%	5.35%	90.19%
<i>Medium Models</i>			
Aya Expanse 32B	3.60%	4.97%	89.89%
Qwen-3 30B	5.60%	3.84%	90.39%
Gemma-3 27B	5.60%	5.51%	90.49%
Phi-4 14B	9.60%	5.40%	90.59%
<i>Small Models</i>			
Llama-3.1 8B	3.60%	2.55%	85.29%
Qwen-2.5 7B	4.40%	3.08%	83.78%
Gemma-3 4B	2.80%	2.38%	80.08%
Qwen-3 4B	1.60%	1.84%	77.18%
Average	5.03%	4.24%	86.89%

Table 3: **Accuracy of Definition Generation.** Development and test set accuracy of LLMs for the Mainz dialect, and test set accuracy for the English baseline.

on LLM-generated synthetic data. All predictions are generated using greedy decoding.

Llama Models. We evaluate two models from the Llama family: Llama-3.1 8B and Llama-3.3 70B by Grattafiori et al. (2024). Both are dense, decoder-only transformer architectures.

GPT-OSS-120B. GPT-OSS 120B (Thinking) by OpenAI (2025) is a mixture-of-experts model designed for explicit thinking.

Qwen Models. We include Qwen 3 4B (Instruct) and Qwen 2.5 (7B, 72B) (Qwen, 2025; Qwen et al., 2025). Both are dense, decoder-only transformer models. Furthermore, we include Qwen3 30B (Instruct), which is a mixture-of-experts model.

Gemma Models. The Gemma 3 models (4B, 27B) (Gemma et al., 2025) are dense, decoder-only transformers optimized for multilingual and multimodal tasks, supporting over 140 languages.

Leo-HessianAI Model. Leo-HessianAI 70B (Plüster and Schuhmann, 2023) is a dense transformer model specifically optimized for German-language tasks.

Phi 4. Phi 4 (14B) (Abdin et al., 2024) is a compact, dense, decoder-only model emphasizing training on high-quality data. This model is pretrained on a large quantity of synthetic data.

Aya-Expanse Model. Aya Expanse 32B (Dang et al., 2024) is a multilingual, dense, decoder-only model trained on high-quality multilingual data.

5. Experiment 1: Definition Generation

In this experiment, we investigate the understanding capabilities of LLMs, defined as their ability to accurately infer and convey the meanings of Meenerisch words in Standard German.

5.1. Prompt Design

We instruct the LLMs to generate one valid definition of each given word, assuming that the word is used in the dialect of Mainz and that the definition should reflect its meaning within this dialectal context. See Appendix B.2 for the full prompt.

5.2. Evaluation

Automatic Evaluation We employ an LLM-as-a-judge (LLMAAJ) approach for automatic evaluation, using Llama 3.3 70B as the judge. The evaluation is binary, labeling each pair of definitions as either equal or unequal. Since some words in the dictionary contain multiple reference definitions, we consider a generated definition correct if it matches at least one of the ground-truth definitions. We report the full prompt in Appendix B.2.

Manual Verification To evaluate the reliability of our LLMAAJ approach, we manually inspect a

random subset of 100 samples, 50 labeled as unequal and 50 as equal by the model. We find that, for the unequal cases, the model’s judgments are fully accurate (100%). For the examples labeled as equal, we observe an accuracy of 92%. These results indicate that our automated judgments yield accuracy values that approximate the upper limit of performance. We therefore consider this level of agreement sufficient for the purposes of our study.

English Comparison To establish a capacity baseline, we evaluate all LLMs on an English dictionary dataset (Therrien, 2024), sampling 2,000 entries from 42,052. Because English is heavily represented in pretraining data, this approximates the models’ best-case performance on dictionary definition generation. We apply the identical inference and evaluation pipeline as in the dialect setting, differing only in prompting the models to produce English definitions to match the reference.

5.3. Results and Discussion

We report the test set results for Experiment 1 in Table 3. We additionally include the development set results to enable future comparisons.

Overall, we see low accuracy for all models: the average is just 4.24%, with the best performing model, Llama-3.3 70B, reaching 6.27% and the worst model, Qwen-3 4B, reaching only 1.84%.

Across model sizes, differences are modest, given the generally low accuracy levels. However, small models show the worst performance, with the best small model, Qwen-2.5 7B, having a lower accuracy than the worst medium model, Qwen-3 30B. Medium and large model performances overlap: the top three medium models each outperform at least one large model. However, the overall best model, Llama-3.3 70B, is a large model.

Comparing to English (rightmost column of Table 3), we see that, in contrast, the same models achieve an average accuracy of 86.89% when generating English definitions of English words. Even the worst model, Qwen-3 4B, reaches an accuracy of 77.18%. This shows that our LLMs work correctly and are, in fact, strong models, but do not accurately capture the meanings of dialect words.

To sum up, Experiment 1 shows that **LLMs of all sizes struggle with understanding words in the dialect of Mainz.**

6. Experiment 2: Dialect Word Generation

We now investigate whether LLMs correctly generate a dialect word when provided with its meaning.

6.1. Prompt Design

Each model is prompted with a Standard German explanation and instructed to produce the equivalent word in the dialect of Mainz. The prompt explicitly requests a single, concise dialect term without additional commentary or reformulation. See Appendix B.3 for the full prompt.

6.2. Evaluation

Automatic Evaluation To assess accuracy, we perform a direct string comparison between the (cleaned) predicted and the gold-standard dialect word. This straightforward equality-based evaluation quantifies how accurately each model reproduces the correct dialect form, and we report the resulting accuracy for each model individually.

English Comparison We again compare to the performance on the English dictionary from Section 5.2. Using the same inference pipeline, we prompt the models to generate the corresponding English word for a given definition.

6.3. Results and Discussion

We report the results of Experiment 2 on the test set in Table 4. We also add development set results.

The results for Experiment 2 are even worse than those for Experiment 1: the average accuracy is only 0.56%. In fact, the accuracy of most models is below 1%—the only exceptions are Llama-3.3 70B with 1.03% and GPT-OSS 120B (Thinking) with 1.51%. As for Experiment 1, we see that small models, on average, perform worse than medium models, which, in turn, are outperformed by large models on average. However, given the low accuracies overall, absolute differences are minimal.

Again, we also compare to the same task in English, i.e., generating English words, given their definitions. We also see lower accuracies than for Experiment 1, hinting at the fact that the task is more difficult. However, performances for English are much better than for Meenzerisch, with an average accuracy of 58.77% and the worst model, Gemma-3 4B, still obtaining 43.24% accuracy.

We conclude that **LLMs are largely unable to produce words in the dialect of Mainz, underscoring a gap in their generative capabilities.**

7. Additional Experiments

Given the low performance of all LLMs in our main experiments, we further conduct two smaller experiments, using only the best LLM, Llama-3.3 70B, to explore whether few-shot in-context learning and automatic rule extraction improve the performance

Model	Dialect of Mainz		English Test Set
	Dev Set	Test Set	
<i>Large Models</i>			
GPT-OSS 120B (Thinking)	2.40%	1.51%	82.50%
Llama-3.3 70B	0.80%	1.03%	69.00%
Leo-HessianAI 70B	0.80%	0.92%	47.10%
Qwen-2.5 72B	0.80%	0.43%	66.90%
<i>Medium Models</i>			
Aya Expanse 32B	1.20%	0.59%	59.50%
Qwen-3 30B	0.40%	0.27%	60.10%
Gemma-3 27B	1.20%	0.43%	65.60%
Phi-4 14B	1.60%	0.49%	64.10%
<i>Small Models</i>			
Llama-3.1 8B	0.00%	0.54%	52.10%
Qwen-2.5 7B	0.00%	0.16%	49.55%
Gemma-3 4B	0.00%	0.16%	43.24%
Qwen-3 4B	0.00%	0.16%	45.50%
Average	0.77%	0.56%	58.77%

Table 4: **Accuracy of Word Generation given Explanation.** Development and test set accuracy of LLMs in generating the correct word for the Mainz dialect, and test set accuracy for the English baseline.

of definition and word generation (see Sections 5 and Section 6).

7.1. Few-Shot Learning

We first investigate if few-shot in-context learning improves the generation of definitions. This approach is particularly promising for the task, as many dialect words follow systematic transformations from their Standard German counterparts. For example, in Meenzerisch, infinitives end on “-ele” as opposed to on “-en” in Standard German.

We provide the LLM with k training examples in the prompt. We tune k on the development set by testing 1, 5, 10, 25, and 50 examples (see Table 5). We do not apply prompt tuning, as it would require extensive computational resources. To decide which examples to include in the prompt, we compare two retrieval strategies: random selection and edit distance-based selection.

Selection: Random Examples are randomly sampled from the training set. For each test instance, a new random subset of examples is drawn.

Selection: Edit Distance-Based We select examples based on textual similarity to the target instance, measured by Levenshtein edit distance. This approach prioritizes examples whose surface forms are similar to the test word, aiming to provide the model with relevant cues for generating accurate definitions. We use edit distance-based selection only for definition generation, as we do not expect edit distance to be meaningful for the

Shots	Definition Generation	Word Generation
<i>Random Selection</i>		
1	6.00%	1.60%
5	9.60%	2.00%
10	11.20%	1.60%
25	7.60%	1.20%
50	8.00%	2.00%
<i>Edit Distance-Based Selection</i>		
1	8.00%	–
5	10.00%	–
10	9.20%	–
25	9.20%	–
50	10.40%	–

Table 5: **Different Numbers of Few-shot Examples on the Development Set.** Comparison of Llama-3.3 70B Instruct across different numbers of in-context examples.

relatively longer descriptions—the input of the word generation task.

Results Table 6 summarizes our results on the test set and compares them to Llama-3.3 70B’s performance during Experiment 1 and 2 as a baseline. We see that all in-context learning approaches improve upon the zero-shot results. However, the magnitude of this improvement varies: For definition generation, random example selection yields a modest improvement of 1.4% points, which is not statistically significant. In contrast, edit distance-

Selection Method	Accuracy
Definition Generation	
<i>Baseline: Zero-Shot</i>	6.27%
Random Selection ($k = 10$)	7.67%
Edit Distance-Based Selection ($k = 50$)	9.24%*
Word Generation	
<i>Baseline: Zero-Shot</i>	1.03%
Random Selection ($k = 5$)	1.24%

Table 6: **Few-shot Results for Llama-3.3 70B.** Asterisks (*) mark significant improvements (McNemar’s test, $p < .05$); bolding highlights the best significant score (or multiple scores where differences are not significant).

Dialect-to-Standard Mapping Rules

```
## Practical Application Rules
### Step-by-Step Mapping Process
1. Identify Core Root
- Extract the main semantic element
  from the dialect word
- Example: 'Klebberschulde' → focus
  on 'Klebb' and 'Schuld'
2. Apply Sound Transformations
- Apply systematic sound changes based
  on patterns above
- Example: 'Klebb' → 'Kleb'
  (consonant simplification)
```

Figure 2: **Part of our File with Automatically Extracted Rules.** The final rules are fed into Llama-3.3 70B for definition and word generation.

based selection results in a significant gain of 2.97%. For word generation and random selection, performance improves by only 0.21%, which is not significant. Overall, we find that few-shot in-context learning provides small benefits, especially with edit distance-based selection, although the absolute accuracy levels remain very low.

7.2. Automatic Rule Generation

We further experiment with automatically extracting linguistic transformation rules from the training dataset and feeding them to an LLM—again Llama-3.3 70B—for definition and word generation. To this end, we input the entire training split into DeepSeek-R1 (Guo et al., 2025) and prompt the model to identify systematic mappings between dialect words and their corresponding Standard German forms. We provide the complete prompt and a detailed description of the procedure in Section B.5. We subsequently use the model-generated summary of these extracted rules. While such mappings

Selection Method	Accuracy
Definition Generation	
<i>Baseline: Zero-Shot</i>	6.27%
Injected Extracted Rules	8.37%*
Word Generation	
<i>Baseline: Zero-Shot</i>	1.03%
Injected Extracted Rules	0.76%

Table 7: **Results for Automatically Extracted Rules Injected into Llama-3.3 70B.** Asterisks (*) mark significant improvements (McNemar’s test, $p < .05$); bolding highlights the best significant score (or multiple scores where differences are not significant).

may not apply universally, the model is instructed to detect recurring similarities and phonological or morphological patterns. The extracted rules are presented in a concise, human-interpretable format, designed to be both concrete and comprehensive. The beginning of our file with extracted rules is shown in Figure 2. We provide all extracted rules in Appendix B.5.

These extracted rules are subsequently incorporated into the prompts at inference time to support rule-augmented generation.

Results Table 7 summarizes the results and, again, compares them to Llama-3.3 70B’s performance during Experiments 1 and 2. Injecting the extracted rules into the model increases accuracy for definition generation to 8.37%. This suggests that helpful rules are indeed automatically derived from the word–definition pairs in the training set. In contrast, performance on word generation slightly decreases, though the change is not statistically significant. Overall, accuracy remains low, indicating that the models still struggle with Meenzerisch, even with the provided rules.

8. Conclusion

With this work, we present the first NLP research on Meenzerisch, the dialect of the German city of Mainz. We first introduce a dataset of words in the dialect together with descriptions of their meaning in Standard German. We then evaluate the dialect knowledge of multiple LLMs by asking them to generate definitions for the words as well as words for given definitions. Our evaluation reveals that state-of-the-art LLMs struggle with both tasks. While few-shot in-context learning and automatic rule extraction yield modest gains, overall performance remains low, underscoring the need for additional research to make LLMs work for Meenzerisch.

9. Limitations

We note that our work primarily establishes missing model abilities. Although we explore few-shot in-context learning and automatically extracted rules to improve performance, these are just two possible prompting techniques. Further prompt tuning as well as finetuning can potentially increase performance and should be investigated in future work. In order to enable LLMs to work well for Meenzerisch, additional language resources will need to be created and included into model training. Additionally, leveraging other dialect–standard language pairs could further enrich the dataset and support more robust generalization across dialectal variation. Crucially, future methods should not rely solely on increased compute, as this effectively requires underrepresented cultures to bear higher costs to achieve equivalent service (Bui et al., 2025b).

Finally, we note that our evaluation is limited to word-level dialect understanding rather than full sentences. However, as the first resource of its kind, our work already reveals failures at the word level, suggesting that this foundational gap must be addressed before moving to more complex linguistic structures.

10. Ethics Statements

Our dataset contains a small number of potentially offensive or sensitive words. In our training split, we identify 5 such instances out of 250 samples (2%). As these expressions are authentic components of the dialect and carry linguistic relevance, we choose to retain them in the dataset.

In adherence to the standards for working with minority groups (Liu et al., 2022; Mager et al., 2023), we actively engage with community members, with one native speaker joining this work as an author.

11. Acknowledgment

This work was supported by the Carl Zeiss Foundation through the TOPML project, grant number P2021-02-014.

12. Bibliographical References

Marah Abdin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J. Hewett, Mojan Javaheripi, Piero Kauffmann et al. 2024. [Phi-4 technical report](#).

Ekaterina Artemova and Barbara Plank. 2023. [Low-resource bilingual dialect lexicon induction with](#)

[large language models](#). In *Proceedings of the 24th Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 371–385, Tórshavn, Faroe Islands. University of Tartu Library.

Martin Bär, Andrea DeMarco and Gorka Labaka. 2025. [Swiss German speech translation and the curse of multidialectality](#). In *Proceedings of the 22nd International Conference on Spoken Language Translation (IWSLT 2025)*, pages 165–179, Vienna, Austria (in-person and online). Association for Computational Linguistics.

Hans-Peter Betz. 2010. Die kleine sprachgeschichte. Meenzerisch – Oder: Warum die da so anders sprechen. https://assets.deutschlandfunk.de/FILE_ea77b362cd2567bdbe5a34935531af4e/original.pdf. Accessed: 2025-10-15.

Verena Blaschke, Hinrich Schuetze and Barbara Plank. 2023. [A survey of corpora for Germanic low-resource languages and dialects](#). In *Proceedings of the 24th Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 392–414, Tórshavn, Faroe Islands. University of Tartu Library.

Verena Blaschke, Miriam Winkler, Constantin Förster, Gabriele Wenger-Glemser and Barbara Plank. 2025. [A multi-dialectal dataset for german dialect asr and dialect-to-standard speech translation](#). In *Interspeech 2025*, page 913–917. ISCA.

Minh Duc Bui, Carolin Holtermann, Valentin Hofmann, Anne Lauscher and Katharina von der Wense. 2025a. [Large language models discriminate against speakers of german dialects](#).

Minh Duc Bui, Kyung Eun Park, Goran Glavaš, Fabian David Schmidt and Katharina Von Der Wense. 2025b. [On generalization across measurement systems: LLMs entail more test-time compute for underrepresented cultures](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 21262–21276, Vienna, Austria. Association for Computational Linguistics.

John Dang, Shivalika Singh, Daniel D’souza, Arash Ahmadian, Alejandro Salamanca, Madeline Smith, Aidan Peppin, Sungjin Hong, Manoj Govindassamy, Terrence Zhao et al. 2024. [Aya expand: Combining research breakthroughs for a new multilingual frontier](#).

Ludwig M. Eichinger, Anne-Kathrin Gärtig and Albrecht Plewnia. 2014. *Aktuelle Spracheinstellungen in Deutschland*. Institut für Deutsche Sprache und Universität Mannheim.

- Candace Kaleimamoowahinekapu Galla. 2016. Indigenous language revitalization, promotion, and education: Function of digital technology. *Computer Assisted Language Learning*, 29(7):1137–1151.
- Team Gemma, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière et al. 2025. [Gemma 3 technical report](#).
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan et al. 2024. [The llama 3 herd of models](#).
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi et al. 2025. [Deepseek-r1 incentivizes reasoning in llms through reinforcement learning](#). *Nature*, 645(8081):633–638.
- Barry Haddow, Adolfo Hernández, Friedrich Neubarth and Harald Trost. 2013. [Corpus development for machine translation between standard and dialectal varieties](#). In *Proceedings of the Workshop on Adaptation of Language Resources and Tools for Closely Related Languages and Language Variants*, pages 7–14, Hissar, Bulgaria. INCOMA Ltd. Shoumen, BULGARIA.
- Anjali Kantharuban, Ivan Vulić and Anna Korhonen. 2023. [Quantifying the dialect gap and its correlates across languages](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 7226–7245, Singapore. Association for Computational Linguistics.
- Sali Levi. 1927. *Magenza: ein Sammelheft über das jüdische Mainz im fünfhundertsten Todesjahre des Mainzer Gelehrten Maharil*. Habrith-Verlag-Gesellschaft.
- Robert Litschko, Verena Blaschke, Diana Burkhardt, Barbara Plank and Diego Frassinelli. 2025. [Make every letter count: Building dialect variation dictionaries from monolingual corpora](#).
- Zoey Liu, Crystal Richardson, Richard Hatcher and Emily Prud'hommeaux. 2022. [Not always about you: Prioritizing community needs when developing endangered language technology](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3933–3944, Dublin, Ireland. Association for Computational Linguistics.
- Manuel Mager, Ximena Gutierrez-Vasques, Gerardo Sierra and Ivan Meza-Ruiz. 2018. [Challenges of language technologies for the indigenous languages of the Americas](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 55–69, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Manuel Mager, Elisabeth Mager, Katharina Kann and Ngoc Thang Vu. 2023. [Ethical considerations for machine translation of indigenous languages: Giving a voice to the speakers](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4871–4897, Toronto, Canada. Association for Computational Linguistics.
- Vuyisile Msila. 2011. “mama does not speak that (language) to me”: indigenous languages, educational opportunity and black african preschoolers. *South African Journal of Childhood Education*, 1(1):20.
- Alberto Muñoz-Ortiz, Verena Blaschke and Barbara Plank. 2025. [Evaluating pixel language models on non-standardized languages](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 6412–6419, Abu Dhabi, UAE. Association for Computational Linguistics.
- A. Niemann. 1964. *Die Landwirtschaft Niedersachsens, 1914-1964*. Landbuch-Verlag.
- OpenAI. 2025. [gpt-oss-120b & gpt-oss-20b model card](#).
- Siyao Peng, Zihang Sun, Huangyan Shan, Marie Kolm, Verena Blaschke, Ekaterina Artemova and Barbara Plank. 2024. [Sebastian, Basti, Wastl?! recognizing named entities in Bavarian dialectal data](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 14478–14493, Torino, Italia. ELRA and ICCL.
- Björn Plüster and Christoph Schuhmann. 2023. [Leolm/leo-hessianai-70b](#). <https://huggingface.co/LeoLM/leo-hessianai-70b>. German-English bilingual language model based on Llama-2 70B, continued pretraining on German corpora. Released under the Llama 2 Community License.
- Team Qwen. 2025. [Qwen3 technical report](#).
- Team Qwen, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu,

Chengyuan Li, Dayiheng Liu, Fei Huang et al. 2025. [Qwen2.5 technical report](#).

Gijsbert Rutten and Rik Vosters. 2021. *Language Standardization 'from Above'*, Cambridge Handbooks in Language and Linguistics, page 65–92. Cambridge University Press.

Karl Schramm. 1966. *Mainzer Wörterbuch*. Dr. Hanns Krach.

Anthony Therrien. 2024. [Dictionary of english words and definitions](#).

Kerstin Trillhaase. 2021. *Der Einfluss der deutschen Dialekte Obersächsisch und Mittelbairisch auf die Wahrnehmung der Persönlichkeit*. Logos Verlag Berlin.

Lars Vorberger. 2019. *Regionalsprache in Hessen: eine Untersuchung zu Sprachvariation und Sprachwandel im mittleren und südlichen Hessen*. Franz Steiner Verlag Stuttgart.

Mason A Wirtz and Andrea Ender. 2025. Functional prestige in sociolinguistic evaluative judgments among adult second language speakers in Austria: Evidence from perception. *Languages*, 10(4):67.

A. Code and Dataset

We release our code at <https://github.com/MinhDucBui/Meenz-bleibt-Meenz>. Instructions for accessing the dataset are provided at the same link under the section “Access to Dataset”. The dataset is distributed under the CC BY-NC-ND 4.0 license.

B. Full Prompts

We report the full prompts used in our study.

B.1. Dataset Creation: Extracting Definition

To extract and clean the dictionary definitions for the creation of our dataset (see Section 3.2), which constitutes the fifth step in our dataset construction pipeline, we report the full prompt in Figure 3. For brevity, we omit the system prompt (translated as: “You are a precise and reliable assistant for extracting linguistic data. Provide exclusively the requested information without modifying, shortening, or adding content.”). Furthermore, the full prompt includes six in-context examples, which are documented in the code.

We clean each extracted definition of OCR-induced noise using Llama-3.3 70B; the corresponding prompt is shown in Figure 4.

Definition Extraction Prompt

“Du erhältst eine unstrukturierte oder fehlerhafte Definition des Wortes '{Word}' aus einem alten Wörterbuch. Deine Aufgabe ist es, nur die *eigentliche* Bedeutung des Wortes aus dem Text zu extrahieren, ohne Kommentare, Reformulierungen oder Erklärungen.

Regeln:

1. Wenn mehrere Bedeutungen vorhanden sind, nummeriere sie fortlaufend (1., 2., 3., ...) und trenne sie jeweils mit einem Zeilenumbruch.
2. Wenn die Definition ausschließlich auf ein anderes Wort verweist, gib '[SIEHE] <Wort>' aus.
3. Verändere den Originaltext nicht, sondern gib ausschließlich den relevanten Ausschnitt wieder.
4. Wenn es keine Definition im Text gibt, gebe 'Keine Definition' wieder. <Few-Shot Examples>”

Figure 3: **Prompt used for extracting dictionary definitions.** *English translation:* “You receive an unstructured or faulty definition of the word '{Word}' from an old dictionary. Your task is to extract only the actual meaning of the word from the text, without comments, reformulations, or explanations. Rules: (1) If multiple meanings are present, number them consecutively (1., 2., 3., ...) and separate them by line breaks. (2) If the definition refers exclusively to another word, output '[SEE] <word>'. (3) Do not alter the original text; output only the relevant excerpt. (4) If there is no definition in the text, output 'No definition.'”

B.2. Definition Generation

Figure 5 presents the zero-shot prompt employed to generate dictionary definitions based on the Mainzer dialect in Section 5. To generate the English word definitions for English words, we replace the “Mainzer Dialekt” part with “Englisch” and the “Hochdeutsch” (standard German) with “Englisch”.

The complete prompt used for the LLM-as-a-Judge (LLMaaJ) approach is shown in Figure 6.

B.3. Dialect Word Generation

We report the word-generation prompt from Section 6 in Figure 7. For the English setup, we simply replace “Mainzer Dialekt” with “Englisch” and “Hochdeutsch” (standard German) with “Englisch”.

Definition Cleaning Prompt

```
{system_token}
"Du bist ein sorgfältiger
linguistischer Assistent. Deine
Aufgabe ist es,
Wörterbuchdefinitionen zu bereinigen,
ohne deren Bedeutung zu verändern.
Entferne ausschließlich unnötige
Sonderzeichen wie Bindestriche,
doppelte Leerzeichen oder ähnliche
OCR-Artefakte. Lasse Verweise in der
Form [SIEHE] unverändert und
Nummerierungen von Definitionen."
{user_token}
"Hier ist die Definition für das Wort
'{Word}':
{cleaned}

Bereinige die Definition gemäß den
Anweisungen. Wenn sie bereits sauber
ist, gib sie unverändert zurück.
Gebe nur die bereinigte Definition
wieder"
```

Figure 4: **Prompt used for cleaning dictionary definitions** *English translation*: “You are a careful linguistic assistant. Your task is to clean dictionary definitions without changing their meaning. Remove only unnecessary special characters such as hyphens, double spaces, or similar OCR artifacts. Leave references in the form [SEE] unchanged and preserve numbering of definitions. Here is the definition for the word '{Word}': {cleaned}. Clean the definition according to the instructions. If it is already clean, return it unchanged. Output only the cleaned definition.”

B.4. Few-Shot Learning

For the few-shot setting, we augment the prompt shown in Figure 5 by appending examples in the following format: “Wort: <WORD> Definition: <DEFINITION>”.

B.5. Automatic Rule Generation

We first present the prompts used for automatic rule generation (see Figure 8). We adopt a two-turn approach: in the first turn, the model generates Standard German candidates; in the second turn, it induces mapping rules between the dialect and Standard German forms.

The complete set of extracted rules, described in Section 7.2, is provided in Figure 9.

Dictionary Definition Prompt

```
{system_token}
"Du bist ein präziser und
zuverlässiger linguistischer
Assistent. Deine Aufgabe ist es,
Wörterbuchdefinitionen zu erstellen.
Gib ausschließlich eine einzige,
kurze und prägnante Bedeutung des
angefragten Wortes an. Nutze dabei
die im Mainzer Dialekt gebräuchliche
Bedeutung des Wortes, formuliere die
Definition jedoch auf Hochdeutsch."
{user_token}
"Erstelle genau eine kurze
Wörterbuchdefinition für das Wort
'{Word}'. Verwende dabei die im
Mainzer Dialekt gebräuchliche
Bedeutung, ohne Zusatzinformationen,
Beispiele oder alternative
Bedeutungen anzugeben."
```

Figure 5: **Prompt used for generating dictionary definitions.** *English translation*: “You are a precise and reliable linguistic assistant. Your task is to create dictionary definitions. Provide exclusively a single, short, and concise meaning of the requested word. Use the meaning commonly used in the Mainzer dialect, but formulate the definition in Standard German. Create exactly one short dictionary definition for the word '{Word}'. Use the meaning commonly used in the Mainzer dialect, without providing additional information, examples, or alternative meanings.”

Definition Similarity Evaluation Prompt

```
{system_token}
"Du bist ein präziser und neutraler
Evaluator für Bedeutungsähnlichkeit
von Wörterbuchdefinitionen. Deine
Aufgabe ist es, zu beurteilen, ob
zwei Definitionen inhaltlich dieselbe
Bedeutung ausdrücken, auch wenn sie
unterschiedlich formuliert sind.
Berücksichtige Synonyme,
Umformulierungen oder stilistische
Unterschiede. Antworte
ausschließlich mit 'GLEICH', wenn die
beiden Definitionen denselben Inhalt
wiedergeben, oder mit
'UNTERSCHIEDLICH', wenn sich ihre
Bedeutung wesentlich unterscheidet.
Wörterbucheinträge können mehrere
Definitionen enthalten. Wenn
mindestens eine dieser Definitionen
mit der vom LLM erzeugten
übereinstimmt, gilt das Ergebnis als
'GLEICH'. Gib keine weiteren
Erklärungen oder Begründungen."
{user_token}
"Vergleiche die folgenden zwei
Definitionen. Beurteile, ob sie
inhaltlich gleich sind.
Definition A (LLM):
'{response_generator}'
Definition B (Wörterbuch):
'{definition_final}'"
```

Figure 6: **Prompt used for evaluating semantic equivalence.** *English translation:* "You are a precise and neutral evaluator of semantic similarity between dictionary definitions. Your task is to assess whether two definitions express the same meaning, even if they are phrased differently. Consider synonyms, paraphrases, or stylistic differences. Respond exclusively with 'SAME' if the two definitions convey the same content, or with 'DIFFERENT' if their meaning differs substantially. Dictionary entries may contain multiple definitions. If at least one of these definitions matches the one generated by the LLM, the result counts as 'SAME'. Do not provide any additional explanations or justifications. Compare the following two definitions. Assess whether they are semantically equivalent. Definition A (LLM): '{response_generator}'. Definition B (Dictionary): '{definition_final}'."

Dialect Word Generation Prompt

```
{system_token}
"Du bist ein präziser und
zuverlässiger linguistischer
Assistent. Deine Aufgabe ist es, zu
einem gegebenen Bedeutungsinhalt das
passende Wort im Mainzer Dialekt zu
finden. Gib ausschließlich ein
einziges Wort aus, das diese
Bedeutung im Mainzer Dialekt
ausdrückt. Gib keine Erklärungen,
Übersetzungen oder
Zusatzinformationen."
{user_token}
"Finde das Mainzer Dialektwort, das
die folgende Bedeutung ausdrückt:
'{definition_final_generation}'
Antworte nur mit dem Dialektwort."
```

Figure 7: **Prompt used for generating a Mainzer dialect word.** *English translation:* "You are a precise and reliable linguistic assistant. Your task is to find the appropriate word in the Mainzer dialect for a given meaning. Provide exclusively a single word that expresses this meaning in the Mainzer dialect. Do not provide explanations, translations, or additional information. Find the Mainzer dialect word that expresses the following meaning: '{definition_final_generation}'. Respond only with the dialect word."

Two-Turn Prompt for Rule Extraction

```
----- Turn 1 -----  
  
Given the following CSV table, you  
need to perform the following tasks:  
Create an additional column with a  
German word. The word should fit the  
definition and be as similar as  
possible to the "Word".  
  
<TRAINING SET>  
  
---- Turn 2 ----  
  
Now, try to find common mappings  
between the dialect words and the  
German words. The mapping might not  
apply to all cases, but attempt to  
identify recurring similarities and  
patterns.  
Write explicit rules and pattern  
descriptions that could be used by a  
human. The rules should be concrete  
but comprehensive.
```

Figure 8: **Two-turn prompting strategy for automatic rule extraction.** In the first turn, the model generates Standard German candidates. In the second turn, it induces mapping rules between dialect and Standard German forms.

Dialect-to-Standard Mapping Rules

```
## Practical Application Rules  
### Step-by-Step Mapping Process  
1. Identify Core Root  
- Extract the main semantic element  
from the dialect word  
- Example: 'Klebberschulde' → focus  
on 'Klebb' and 'Schuld'  
2. Apply Sound Transformations  
- Apply systematic sound changes based  
on patterns above  
- Example: 'Klebb' → 'Kleb'  
(consonant simplification)  
3. Handle Suffixes Systematically  
- Apply standard German suffix  
equivalents  
- Example: '-schulde' → '-schulden'  
(noun pluralization)  
4. Consider Semantic Context  
- Use definition to guide final word  
choice  
- Ensure the mapped word fits the  
semantic field  
  
### Quick Reference Guide  
| Dialect Pattern | Standard German  
Equivalent | Example |  
|-----|-----|-----|  
| '-che' | '-chen' | 'Bobbelche' →  
'Bobby' |  
| '-ele' | '-eln' | 'knerchele' →  
'knirschen' |  
| 'aa' | 'ei' | 'Gaawer' → 'Geifer' |  
| Final '-e' | Often added to nouns |  
'Grobbe' → 'Grob' |  
| Compound words | Preserve first  
element | 'Klebberschulde' →  
'Kleberschulden' |  
| Professional '-er' | Usually  
preserved | 'Stinkerd' → 'Stinker' |
```

Figure 9: **Automatic extracted rules for mapping Mainzer dialect words to Standard German forms.** This is generated by DeepSeek-R1.