

# Standard-to-Dialect Transfer Trends Differ across Text and Speech: A Case Study on Intent and Topic Classification in German Dialects

Verena Blaschke<sup>▲</sup><sup>🏠</sup> Miriam Winkler<sup>▲</sup> Barbara Plank<sup>▲</sup><sup>🏠</sup>

<sup>▲</sup> MaiNLP lab, CIS, LMU Munich, Germany

<sup>🏠</sup> Munich Center for Machine Learning, Germany

{verena.blaschke, b.plank} @lmu.de

## Abstract

Research on cross-dialectal transfer from a standard to a non-standard dialect variety has typically focused on text data. However, dialects are primarily spoken, and non-standard spellings cause issues in text processing. We compare standard-to-dialect transfer in three settings: text models, speech models, and cascaded systems where speech first gets automatically transcribed and then further processed by a text model. We focus on German dialects in the context of written and spoken intent classification – releasing the first dialectal audio intent classification dataset – with supporting experiments on topic classification. The speech-only setup provides the best results on the dialect data while the text-only setup works best on the standard data. While the cascaded systems lag behind the text-only models for German, they perform relatively well on the dialectal data if the transcription system generates normalized, standard-like output.

## 1 Introduction

While most natural language processing (NLP) research focuses on highly standardized languages with large amounts of training and evaluation data available, a subfield focusing on dialect NLP has recently been emerging, studying the (lack of) robustness NLP tools exhibit towards dialectal variation (Zampieri et al., 2020; Joshi et al., 2025).

In dialect NLP, oftentimes no training but only evaluation data are available for non-standard dialect varieties. In such cases it is common to train NLP models on a higher-resource, closely related standard language instead. Part of what makes such standard-to-dialect transfer challenging in contrast to standard transfer learning is the large degree of spelling variation exhibited in dialectal writing, as it frequently leads to infelicitous subword tokenizations (Aeppli et al., 2023; Blaschke et al., 2023; Srivastava and Chiang, 2025; Kanjirang et al., 2025).



Figure 1: We compare three evaluation setups for German and dialectal text and speech data.

Such challenges have led to research on processing dialectal data with token-free models, such as for example pixel models (Muñoz-Ortiz et al., 2025). These input representation challenges also bring about the question of how robust speech models are towards dialectal variation, given the continuous nature of their inputs. This has to some extent been investigated for automatic speech recognition (ASR; §2), but not yet for classification tasks. Focusing on speech is especially relevant for processing non-standard dialects, as these varieties are predominantly spoken.

Our work primarily focuses on intent classification in the context of queries for virtual assistants for two reasons: Firstly, dialect-robust virtual assistants are seen as especially desirable in a survey on potential NLP technologies for German dialects (Blaschke et al., 2024). Secondly, this task naturally lends itself to the spoken domain, but most available datasets only include written data. This can partially be explained by the conventional cascaded paradigm of first automatically transcribing audio data (cf. Faruqui and Hakkani-Tür, 2022), and then classifying the transcribed text as a second step. However, this approach ignores the shortcomings of current ASR systems with respect to low-resource and non-standardized languages.

In this study, we extend a text-only intent classification dataset to audio with recordings in German and a related non-standard dialect, Bavarian. We compare three setups across a variety of text and speech encoders and ASR models (Figure 1): models trained and evaluated on the written version of the dataset (“text-only”), on the spoken version (“speech-only”) or trained on text and evaluated on ASR transcripts of the spoken input (“cascaded”). We replicate our findings in a supporting experiment on a conceptually similar task (topic classification) with German and Swiss German (dialect).

We present the following findings, which show that results for standard languages do not necessarily generalize to even closely related non-standard dialects:

1. While the cascaded models lag behind the text-only models for the German data, systems whose ASR step in practice normalizes the data outperform the text-only models for dialectal input (§5.1).
2. Although the speech-only systems perform on par with the cascaded systems on the German data, they greatly outperform the cascaded systems on the dialectal data (§5.2).
3. Although the text-only models perform much better than the speech-only models on the German data, the situation is largely reversed for the dialectal data (§5.3).

Furthermore, we contribute a spoken intent dataset for German and Bavarian (§3.2). We make the dataset<sup>1</sup> and our code and model predictions<sup>2</sup> publicly available.

## 2 Related Work

**Intent classification for written dialects** Several works have focused on standard-to-dialect transfer for written dialects, starting with the xSID dataset (van der Goot et al., 2021), and including multiple shared tasks (Aeppli et al., 2023; Malaysha et al., 2024; Scherrer et al., 2025). A common theme in these works is the gap between the performance on the standard variety and the non-standard dialects. Intent classification systems have been shown to be somewhat affected by syntactic variation (Artemova et al., 2024). Other works have focused on improving intent classification by mitigating the

effect of infelicitous subword tokenizations for dialectal text by manipulating the subword tokenization of the standard language as well (Srivastava and Chiang, 2023; Blaschke et al., 2025a) or by using token-free models (Muñoz-Ortiz et al., 2025).

**Spoken dialect processing** Most work on spoken dialect processing has focused on ASR, covering a variety of languages like Arabic (Djanibekov et al., 2025), Dutch (ten Bosch, 2000), English (Torgbi et al., 2025), German (Beringer et al., 1998), Greek (Vakirtzian et al., 2024), Frisian (Amooie et al., 2025), Irish (Lonergan et al., 2023), Mandarin (Tang et al., 2021), Telugu (Yadavalli et al., 2022), Wenzhounese (Gao et al., 2025), and Yorùbá (Ahia et al., 2024). Common themes in these works are performance gaps between standard and non-standard varieties, the effect of dialectal pronunciation differences, as well as lexical and syntactic differences between the dialect and the standard.

The most relevant works for our study analyze the performance of current ASR models on German dialects. In a study on automatically transcribing German dialects, Blaschke et al. (2025b) find that ASR models differ in how much they normalize their output towards standard German. In a similar vein, Dolev et al. (2024) and Gorisch and Schmidt (2024) find that Whisper ASR models (Radford et al., 2023) tend to replace colloquial or regional German syntactic constructions with constructions that are more common in written German, although this is not always successful (Sicard et al., 2023).

Focusing on a different task – question answering – Faisal et al. (2021) evaluate cascaded models on different nation-level varieties of five languages. They find performance gaps between varieties of the same language and a positive correlation between transcription quality and question answering performance.

**Spoken intent and topic classification** A range of recent works has focused on spoken intent and topic classification. Bastianelli et al. (2020) present an English-language intent classification dataset that was later extended into the multilingual Speech-MASSIVE by Lee et al. (2024). Similarly, Schmidt et al. (2025) and Keller and Glavaš (2025) present large-scale multilingual speech classification datasets. Rajaa et al. (2022) present an Indian English intent classification dataset. Two datasets include multiple varieties of the same language: MInDS-14 (Gerz et al., 2021) includes standard English accents from three countries, with no

<sup>1</sup><https://doi.org/10.5281/zenodo.19554427>

<sup>2</sup>[github.com/mainlp/dialects-text-vs-speech](https://github.com/mainlp/dialects-text-vs-speech)

eng	Remind me to call Stephanie on Tuesday		
deu	Erinnere mich	Stephanie am Dienstag	anzurufen
bar	Erinner mi d	Stephanie am Dienstag	ozumruafa
	<i>remind me the Stephanie on Tuesday to.call</i>		
eng	Because many services that we use today...		
deu	Denn viele Dienste,	die wir heute nutzen,	
gsw-be	Denn vili Dienschtä,	womer hüt nutzä,	
gsw-lu	Well vell Dienst,	womer höt bruched,	
	<i>because many services that we today use</i>		
eng	...only work with the help of satellites.		
deu	funktionieren nur	mit Hilfe von Satelliten.	
gsw-be	funktionierä numä	mit Hiuf vo Satellitä.	
gsw-lu	fonktionierted nor	met Helf vo Satellite.	
	<i>work only with help of satellites</i>		

Table 1: **Standard German and dialectal sentences from the datasets we use** (top: German and Bavarian sentences from xSID; bottom: German and two of the Swiss German sentences – Bern, Lucerne – from SwissDial). While there are some grammatical and lexical differences between the varieties, the most salient differences are on the level of pronunciation/spelling. Note that the xSID sentences were independently translated from English – while the sentence pair in this table is parallel on a word level, this does not apply to all sentences. For more examples, see Table 15 in Appendix F.2.

consistent performance differences between them. ITALIC contains recordings from different Italian regions (Koudounas et al., 2023), but no analysis of performance differences across regions. On the engineering side, Wang et al. (2023) and He and Garner (2023) explore the benefits of multi-task learning for spoken intent detection. In the context of cascaded systems, Faruqui and Hakkani-Tür (2022) argue for a better integration of ASR systems and text-based classification models.

## 3 Experiments

### 3.1 Setups

We compare three setups. In the **text-only** setup, the models are trained on the German text data and evaluated on the German and dialectal text data. In the **cascaded** setup, the same models are evaluated on ASR transcriptions of the German and dialectal speech data.<sup>3</sup> Lastly, in the **speech-only** setup, the models are trained on the German audio data and evaluated on the German and dialectal audio data.

<sup>3</sup>We additionally train some cascaded models on ASR transcriptions, but find that they show the same transfer trends as the cascaded models trained on gold-standard text data (§E).











Dataset	Mod	Lang	Train	Dev	Test
<i>Intent classification</i>					
Speech-MASSIVE		deu	2.5k	459	361
MASSIVE (deu)		deu	2.5k	459	361
MASSIVE (bar)		bar	–	–	361
xSID		deu, bar	–	–	2×412
xSID-audio ( <i>new</i> )		deu, bar	–	–	2×412
<i>Topic classification</i>					
SwissDial (deu)		deu	1.5k	194	396
SwissDial (gsw)	 	gsw (8 dials)	–	–	8×396

Table 2: **Datasets used in this study.** The numbers of instances reflect the subsets we used in our experiments (for instance, in order to have comparable intent label sets across datasets), e.g., the full version of xSID-audio contains 300 development and 500 test instances per variety. Several datasets are parallel (*p.*) to each other. Modality:  = text,  = speech.

### 3.2 Data

We use data in Standard German (deu) and two related dialects: Bavarian (bar) and Swiss German (gsw). Neither Bavarian nor Swiss German are standardized. Both are primarily spoken, but also occasionally written, especially in informal situations. The most salient differences between the standard and non-standard varieties are on the pronunciation level (which is reflected when they are written), although there are also some lexical and grammatical differences (Table 1).

To the best of our knowledge, no classification datasets with parallel spoken and written data in a standard and non-standard variety are available for any other language/dialect.

**Main experiment: intent classification** Our main set of experiments concerns intent classification. We use (and introduce) data in German and a closely related dialect, Bavarian. The data consist of queries and commands for virtual assistants annotated with user intents.

We use the German splits of the text-based MASSIVE dataset (FitzGerald et al., 2023) and its audio counterpart Speech-MASSIVE (Lee et al., 2024) for training models. For evaluation, we use the subset of the MASSIVE test set that has been translated into Bavarian (Winkler et al., 2024) as well as the corresponding German (Speech-)MASSIVE instances. We furthermore use the German and Bavarian test splits of xSID (van der Goot et al., 2021; Winkler et al., 2024).<sup>4</sup> Because xSID is only

<sup>4</sup>xSID is an evaluation dataset, and does not provide (high-

released as a text-based dataset, we additionally create audio recordings (read speech). The audios are recorded in-house by a native speaker of both German and Bavarian who also translated the written Bavarian split. The German and Bavarian recordings are identical in content, speaker, and recording conditions, and only differ in their language variety. Appendix A contains the data statement for xSID-audio.

We map the instances in (Speech-)MASSIVE to the label set used by xSID (Appendix C.1), and exclude all instances that do not match this mapping. As shown in Table 2, this leaves 2.5k/459/361 train/dev/test instances for (Speech-)MASSIVE and 412 test instances for xSID, and 10 labels.<sup>5</sup> There is no overlap in speakers between the training and evaluation splits. Because all text is lower-cased in MASSIVE, we lowercase the xSID data and the ASR transcriptions as well.

**Supporting experiment: topic classification** To replicate our findings to the extent possible with available data, we further include experiments on topic classification, a task conceptually similar to intent classification. Except for our extension of xSID, the dataset below is the only classification dataset that contains parallel (or otherwise comparable) dialectal and standard-variety data where at least the dialectal split is available as both spoken and written data.

We use SwissDial (Dogan-Schönberger et al., 2021), which contains German sentences that are also parallel in up to eight Swiss German dialects. The German sentences are only available as text data, but the dialectal sentences are both written and spoken, allowing us to use the dataset for the text-only and cascaded setups. We exclude rare or incoherent labels (Appendix C.2), resulting in a label set of 10 topics. We split the dataset into 1.5k/194/396 train/dev/test sentences. The test set only contains sentences that are parallel across all eight dialects. For training and development, we only use the Standard German versions.

### 3.3 Models

Table 3 provides an overview of the models we use. We focus on encoder models, as fine-tuned

quality) training splits for German or Bavarian. For the Bavarian data, we use the split from Upper Bavaria (“de-ba”).

<sup>5</sup>The full xSID-audio 0.1 release contains recordings for the entire Bavarian and German development and test sets (300+500 sentences, 16 intents).

	Model	Encoder size (M)	ASR
<i>Text</i>	mBERT	86	
	mDeBERTa	85	
	XLM-R base	86	
	XLM-R large	303	
<i>Speech</i>	mHuBERT	94	
	XLS-R 300M	315	
	XLS-R 300M DE	315	CTC
	MMS 300M	315	
	Whisper tiny	8	LM
	Whisper base	20	LM
	Whisper small	87	LM
	Whisper medium	306	LM
	Whisper large-v3	635	LM
<i>Add. ASR</i>	XLS-R 1B DE		CTC
	MMS 1B-all		CTC
	Whisper large-v3-turbo		LM

Table 3: **Text/speech encoders and additional ASR models.** Encoder size in millions of non-embedding parameters. ASR decoding strategies: CTC=connectionist temporal classification, LM=language modelling. For licenses and links, see Appendix B.

pre-trained encoder models have been shown to outperform larger instruction-tuned decoders on classification tasks (Adelani et al., 2024; Ojo et al., 2025; Kasa et al., 2025; Saatrup Nielsen et al., 2025). Furthermore, fine-tuning the pre-trained models below gives us more control over the experimental setup, as we can ensure each model is exposed to exactly the same labelled intent/topic classification training data.

**Text and speech models of comparable size** We focus on two sets of pretrained multilingual transformer models with comparable backbone parameter sizes (encoder parameters without embedding layers): ~90 M backbone parameters (12 encoder layers, 758 hidden dimensions) and ~300 M parameters (24 encoder layers, 1024 hidden dimensions). For encoder-decoder models, we only fine-tune the encoder (plus a classification head).

As text models, we select mBERT (Devlin et al., 2019), XLM-R base and large (Conneau et al., 2020), and mDeBERTa-v3 (He et al., 2021a,b). The latter has shown strong performance in previous work on intent classification for dialectal German data (Artemova et al., 2024; Krückl et al., 2025). As speech models, we use the encoders of Whisper

small and medium (Radford et al., 2023). We also use mHubERT-147 (Zanon Boito et al., 2024) and MMS 300M (Pratap et al., 2024). We additionally use both the original version of XLS-R 300M (Babu et al., 2022) and a version fine-tuned for German ASR (McDowell, 2022b).

All of the models we use include German in their pre-training data. Only mBERT includes Bavarian; none of the models were trained on Swiss German per their documentation.<sup>6</sup>

**Additional model sizes** To further investigate the effect of model size, we also include the Whisper model sizes tiny (8 M backbone encoder parameters), base (20 M), and large-v3 (635 M). No additional text model sizes are available.

**ASR models** To transcribe the evaluation data in the cascaded setup, we use all speech models mentioned above in this section that are trained/fine-tuned for ASR. To evaluate the effect of the ASR performance on the cascaded systems in greater detail, we additionally use a version of XLS-R 1B fine-tuned for German ASR (McDowell, 2022a), MMS 1B fine-tuned for multilingual ASR, and Whisper large-v3-turbo, a version of Whisper large-v3 with a reduced number of decoder layers. All Whisper models are trained for multilingual ASR. We use all ASR models without further fine-tuning and set German as the output language, as none of the ASR models support German dialects.

The Whisper models consist of a stack of encoder layers (that we also use in the speech-only setup) followed by a stack of decoder layers that act as a language model. The XLS-R and MMS models are based on wav2vec2 (Baevski et al., 2020) and trained for connectionist temporal classification (CTC), mapping the model’s internal representations to individual characters based on the audio alone without further information on plausible character or word sequences in the output language.

We use the word/character error rate for evaluating ASR performance. We also manually evaluate a subset of the transcriptions based on whether they are in fluent German, retain the overall meaning of the utterance, and retain keywords related to the

<sup>6</sup>The training datasets of the speech models are not openly accessible, but the metadata indicate no inclusion of dialectal training data. The fact that the ASR-capable speech models all perform very poorly on Bavarian ASR (Table 5) indicates that it is unlikely these models were exposed to Bavarian data, despite performing quite strongly on Bavarian in the speech-only setup. The qualitative ASR examples in §F.2 further illustrate this.

intent of the sentence (details in §F.1).

**Seeds and hyperparameters** We choose hyperparameter values based on prior work on written and spoken intent classification (Rajaa et al., 2022; Koudounas et al., 2023; Arora et al., 2024; Lee et al., 2024; Schmidt et al., 2025). We use the training and development splits of the German (Speech-)MASSIVE and Swissdial versions for selecting the learning rate<sup>7</sup> and maximum number of training epochs. Details on the hyperparameter search and model training are in §D. We report the mean accuracy values over three random seeds.

## 4 Results

Table 4 shows the overall intent/topic classification results per setup. The detailed results for each ASR and text model combination for the cascaded setup are in Appendix G, Table 17.

As in prior work on zero-shot standard-to-dialect transfer (§2), the results on the German test data are consistently better than on the dialectal data, with performance gaps between 6.1 and 36.7 percentage points (pp.) in accuracy for intent classification and gaps between 7.9 and 12.0 pp. for topic classification. For the intent classification data, the performance also drops in the cross-dataset transfer – the (German) test set accuracies for xSID are lower than for MASSIVE, although the dialect gap is much narrower for xSID (6.1–29.9 for xSID compared to 30.6–36.7 for MASSIVE).

For the speech-only models, the models trained/fine-tuned for German or multilingual ASR outperform their non-fine-tuned counterparts (at least when matched by model size), both on the German and dialectal data. This is consistent with prior works finding ASR to be a useful pre-training task for downstream tasks related to understanding the content of an utterance (He and Garner, 2023; Schmidt et al., 2025).

**Different setups work best for the German and dialectal data.** The text-only models achieve the best performances on the German data, followed by the cascaded models and the ASR-tuned speech models. The non-ASR speech models perform worst. For the Bavarian data however, the ASR-tuned speech models (and the one text model pretrained on Bavarian data, mBERT), perform

<sup>7</sup>We observe that the text models are much more robust to the choice of learning rate than the speech models, with the speech models without ASR tuning being especially sensitive (Table 11 in Appendix D).

Model	Intents						Topics		
	MASSIVE			xSID			Swissdial		
	deu	bar	$\Delta$	deu	bar	$\Delta$	deu	gsw	$\Delta$
Majority class baseline	29.1	29.1		9.5	9.5		21.0	21.0	
Random baseline	10.0	10.0		10.0	10.0		10.0	10.0	
<i>Text-only</i>	<i>93.1</i> <sub>1.6</sub>	<i>59.7</i> <sub>4.4</sub>		<i>81.8</i> <sub>4.6</sub>	<i>58.8</i> <sub>5.2</sub>		<i>54.0</i> <sub>3.3</sub>	<i>43.7</i> <sub>3.8</sub>	
mBERT	94.1 <sub>0.6</sub>	62.8 <sub>5.1</sub>	-31.3	83.5 <sub>0.0</sub>	65.0 <sub>1.3</sub>	-18.5	52.8 <sub>2.1</sub>	41.2 <sub>1.0</sub>	-11.6
mDeBERTa	93.5 <sub>0.8</sub>	56.8 <sub>1.0</sub>	-36.7	79.0 <sub>0.6</sub>	54.4 <sub>2.6</sub>	-24.6	51.4 <sub>1.2</sub>	43.6 <sub>2.6</sub>	-7.9
XLM-R base	92.9 <sub>2.5</sub>	62.3 <sub>3.1</sub>	-30.6	78.6 <sub>7.1</sub>	57.0 <sub>5.0</sub>	-21.7	53.3 <sub>3.2</sub>	41.3 <sub>4.0</sub>	-12.0
XLM-R large	91.8 <sub>1.7</sub>	56.9 <sub>4.4</sub>	-34.9	85.9 <sub>2.5</sub>	59.0 <sub>4.8</sub>	-26.9	58.3 <sub>1.3</sub>	48.7 <sub>1.0</sub>	-9.6
<i>Cascaded</i>	<i>88.3</i> <sub>4.2</sub>			<i>77.1</i> <sub>6.0</sub>	<i>52.5</i> <sub>9.8</sub>		<i>42.9</i> <sub>8.7</sub>		
<i>... averaged over ASR models</i>									
mBERT	89.8 <sub>3.6</sub>			80.2 <sub>3.6</sub>	60.9 <sub>7.0</sub>	-19.4	41.5 <sub>8.5</sub>		
mDeBERTa	87.8 <sub>5.1</sub>			72.3 <sub>5.1</sub>	43.9 <sub>8.1</sub>	-28.4	40.8 <sub>8.8</sub>		
XLM-R base	88.3 <sub>3.8</sub>			75.5 <sub>6.1</sub>	50.5 <sub>7.5</sub>	-25.0	43.1 <sub>8.4</sub>		
XLM-R large	87.3 <sub>3.9</sub>			80.5 <sub>4.9</sub>	54.8 <sub>8.1</sub>	-25.7	46.3 <sub>8.6</sub>		
<i>... averaged over text models</i>									
XLS-R 300M DE	88.0 <sub>2.1</sub>			75.8 <sub>4.4</sub>	48.1 <sub>7.1</sub>	-27.7	32.1 <sub>3.1</sub>		
XLS-R 1B DE	89.0 <sub>1.8</sub>			79.6 <sub>4.5</sub>	49.7 <sub>6.6</sub>	-29.9	33.1 <sub>3.5</sub>		
MMS 1B-all	89.2 <sub>2.1</sub>			75.7 <sub>4.3</sub>	48.8 <sub>4.7</sub>	-26.9	37.3 <sub>3.3</sub>		
Whisper tiny	79.8 <sub>2.8</sub>			69.5 <sub>5.9</sub>	41.9 <sub>8.4</sub>	-27.6	35.6 <sub>2.3</sub>		
Whisper base	84.1 <sub>1.9</sub>			72.8 <sub>5.6</sub>	46.5 <sub>9.2</sub>	-26.3	42.2 <sub>2.3</sub>		
Whisper small	89.4 <sub>1.7</sub>			79.3 <sub>4.6</sub>	57.4 <sub>8.7</sub>	-21.9	49.2 <sub>2.5</sub>		
Whisper medium	91.2 <sub>1.4</sub>			80.0 <sub>4.7</sub>	57.0 <sub>8.0</sub>	-23.0	51.7 <sub>2.9</sub>		
Whisper large-v3-turbo	91.8 <sub>1.5</sub>			80.8 <sub>4.9</sub>	61.8 <sub>6.7</sub>	-19.1	52.6 <sub>3.1</sub>		
Whisper large-v3	92.2 <sub>1.7</sub>			80.8 <sub>4.8</sub>	61.4 <sub>6.6</sub>	-19.4	52.7 <sub>3.1</sub>		
<i>Speech-only</i>	<i>83.1</i> <sub>7.8</sub>			<i>70.6</i> <sub>8.9</sub>	<i>60.8</i> <sub>7.9</sub>				
mHuBERT	75.1 <sub>1.5</sub>			59.7 <sub>0.7</sub>	53.6 <sub>0.5</sub>	-6.1			
XLS-R 300M	76.1 <sub>0.2</sub>			62.1 <sub>0.4</sub>	54.3 <sub>0.6</sub>	-7.8			
XLS-R 300M DE <sup>ASR</sup>	90.9 <sub>0.3</sub>			76.1 <sub>2.2</sub>	66.0 <sub>1.1</sub>	-10.0			
MMS 300M	77.9 <sub>0.4</sub>			65.9 <sub>1.6</sub>	58.2 <sub>1.2</sub>	-7.8			
Whisper tiny <sup>ASR</sup>	71.3 <sub>2.2</sub>			59.6 <sub>1.8</sub>	47.6 <sub>2.8</sub>	-12.1			
Whisper base <sup>ASR</sup>	84.1 <sub>0.7</sub>			70.5 <sub>0.9</sub>	63.3 <sub>1.5</sub>	-7.1			
Whisper small <sup>ASR</sup>	89.9 <sub>0.2</sub>			80.3 <sub>1.2</sub>	69.2 <sub>2.8</sub>	-11.1			
Whisper medium <sup>ASR</sup>	91.0 <sub>0.8</sub>			79.0 <sub>1.3</sub>	62.6 <sub>3.6</sub>	-16.3			
Whisper large-v3 <sup>ASR</sup>	91.4 <sub>0.5</sub>			82.7 <sub>1.1</sub>	72.5 <sub>0.1</sub>	-10.2			
<i>... ASR models</i>	<i>86.4</i> <sub>7.5</sub>			<i>74.7</i> <sub>8.1</sub>	<i>63.5</i> <sub>8.3</sub>				
<i>... non-ASR models</i>	<i>76.4</i> <sub>1.5</sub>			<i>62.6</i> <sub>2.9</sub>	<i>55.4</i> <sub>2.2</sub>				

Table 4: **Although the speech-only models generally show the worst results for the German data, they give the overall best results in the dialectal setting for xSID and Swissdial. The performance of the cascaded models varies strongly for the dialectal data.** Accuracy (in %) on the German (deu) and dialectal (Bavarian, bar; Swiss German, gsw) test sets, and the difference between them ( $\Delta$  bar–deu,  $\Delta$  gsw–deu). Darker shades of green: better performance; darker shades of red: greater performance differences (the colour scales for intents and topics are independent of one another). All results are averaged over three random seeds. The cascaded systems are additionally averaged over ASR/text models (de-aggregated results in Appendix G, Table 17). Standard deviations are in subscripts. For performance differences between the setups, see Table 6. Average performances per setup type are in *italics*. <sup>ASR</sup> = speech models tuned for ASR.

Model	WER						CER					
	Intents			Topics			Intents			Topics		
	MAS.	xSID	xSID	Swissdial			MAS.	xSID	xSID	Swissdial		
	(deu)	(deu)	(bar)	deu ref	gsw ref		(deu)	(deu)	(bar)	deu ref	gsw ref	
<i>Word/character error rates per ASR model used in cascaded setups (<math>\downarrow</math>)</i>												
X. 300M DE	26.4 <sub>25.3</sub>	31.5 <sub>24.7</sub>	86.3 <sub>20.7</sub>	77.9 <sub>22.3</sub>	83.0 <sub>15.5</sub>	82.7 <sub>16.1</sub>	5.9 <sub>7.2</sub>	11.9 <sub>13.7</sub>	51.0 <sub>19.8</sub>	31.3 <sub>17.3</sub>	38.4 <sub>14.8</sub>	29.3 <sub>12.6</sub>
X. 1B DE	23.3 <sub>26.2</sub>	28.0 <sub>24.5</sub>	84.1 <sub>20.8</sub>	76.6 <sub>21.8</sub>	82.9 <sub>16.5</sub>	82.9 <sub>16.8</sub>	5.3 <sub>7.0</sub>	10.7 <sub>13.3</sub>	49.9 <sub>20.1</sub>	31.0 <sub>18.0</sub>	38.4 <sub>15.0</sub>	30.0 <sub>13.1</sub>
MMS 1B-all	20.5 <sub>23.9</sub>	28.9 <sub>24.6</sub>	82.5 <sub>22.4</sub>	75.1 <sub>23.0</sub>	81.0 <sub>20.2</sub>	80.9 <sub>19.7</sub>	5.7 <sub>8.4</sub>	11.0 <sub>11.4</sub>	50.2 <sub>19.3</sub>	32.0 <sub>16.6</sub>	36.6 <sub>14.9</sub>	29.0 <sub>13.3</sub>
Whisper tiny	39.2 <sub>34.4</sub>	42.6 <sub>30.5</sub>	87.4 <sub>25.8</sub>	85.4 <sub>23.8</sub>	85.7 <sub>169.3</sub>	96.3 <sub>136.2</sub>	16.4 <sub>16.7</sub>	16.9 <sub>14.4</sub>	53.4 <sub>18.3</sub>	39.7 <sub>15.9</sub>	46.6 <sub>162.4</sub>	44.6 <sub>149.0</sub>
Whisper base	28.9 <sub>32.3</sub>	31.2 <sub>28.7</sub>	85.1 <sub>30.0</sub>	81.9 <sub>28.6</sub>	64.9 <sub>49.2</sub>	86.0 <sub>54.8</sub>	11.5 <sub>14.3</sub>	12.5 <sub>13.2</sub>	51.4 <sub>21.0</sub>	37.8 <sub>20.9</sub>	31.7 <sub>38.9</sub>	33.7 <sub>41.0</sub>
W. small	18.0 <sub>24.1</sub>	16.6 <sub>21.8</sub>	75.6 <sub>27.3</sub>	72.5 <sub>25.8</sub>	42.9 <sub>25.5</sub>	79.3 <sub>20.2</sub>	13.3 <sub>126.8</sub>	6.1 <sub>9.2</sub>	46.2 <sub>21.0</sub>	31.6 <sub>19.5</sub>	20.6 <sub>17.5</sub>	28.2 <sub>14.3</sub>
W. medium	14.4 <sub>22.7</sub>	11.0 <sub>16.5</sub>	71.6 <sub>28.4</sub>	70.1 <sub>29.5</sub>	29.9 <sub>23.5</sub>	76.8 <sub>19.7</sub>	5.7 <sub>11.0</sub>	4.1 <sub>6.5</sub>	46.5 <sub>46.4</sub>	31.9 <sub>49.4</sub>	15.0 <sub>14.3</sub>	27.1 <sub>10.6</sub>
W. lg.-v3-tur.	14.3 <sub>23.9</sub>	9.0 <sub>14.5</sub>	68.0 <sub>28.0</sub>	67.0 <sub>28.3</sub>	26.0 <sub>22.1</sub>	75.5 <sub>20.1</sub>	5.3 <sub>10.1</sub>	3.2 <sub>5.9</sub>	42.5 <sub>21.6</sub>	28.1 <sub>17.1</sub>	13.3 <sub>13.6</sub>	25.8 <sub>9.9</sub>
W. lg.-v3	11.8 <sub>19.2</sub>	8.2 <sub>13.8</sub>	67.1 <sub>28.4</sub>	68.3 <sub>28.4</sub>	24.1 <sub>21.8</sub>	75.0 <sub>19.7</sub>	4.5 <sub>8.0</sub>	3.0 <sub>5.5</sub>	42.2 <sub>21.9</sub>	28.8 <sub>17.6</sub>	12.7 <sub>13.5</sub>	26.1 <sub>10.2</sub>
<i>Correlations between WER/CER and <math>\Delta</math> cascaded-text</i>												
mBERT	-0.86	-0.75	-0.86	-0.89	-0.97	-0.65	-0.73	-0.77	-0.91	-0.81	-0.94	-0.50
mDeBERTa	-0.98	-0.85	-0.86	-0.93	-0.97	-0.67	-0.72	-0.87	-0.92	-0.88	-0.94	-0.51
XLM-R base	-0.74	-0.77	-0.85	-0.82	-0.96	-0.65	-0.79	-0.78	-0.85	-0.66	-0.93	-0.50
XLM-R large	-0.84	-0.87	-0.87	-0.84	-0.99	-0.73	-0.81	-0.87	-0.88	-0.67	-0.97	-0.58

Table 5: **The better the ASR performance, the better the cascaded systems do compared to the text-only systems. Top:** Word/character error rates of ASR models (in %), averaged over sentences for each test set (standard deviations in subscripts). Lower (darker) is better. **Bottom:** Correlations between error rates and performance differences between cascaded and text-only setups (Pearson’s  $r$ , all  $p$ -values  $< 0.001$  except for the right-most column, where  $0.001 < p < 0.01$ ). For the performance differences (cascaded minus text-only), see Table 6 (right).

best. The accuracy of the cascaded models varies a lot, depending on both the ASR model and the text model. Similarly, for the Swiss German data, the cascaded systems contain both the best and the worst results – depending on the ASR model included – even though on average their performance is similar to that of the text-only models. This trend is consistent across the eight different dialects in Swissdial (Table 16 in Appendix G).

In most setups, bigger models outperform smaller models, with the exception of XLM-R large (text-only and cascaded) for the MASSIVE test data, even though it outperforms its smaller counterpart in the cross-dataset evaluation on xSID.

## 5 Analysis

### 5.1 Cascaded vs. text-only: How well do the text models generalize to transcribed data?

**The quality of the ASR system is an important factor in the cascaded setup.** Table 5 shows the word and character error rates (WER, CER) of the different ASR models for the different test sets, as well as the correlations between the error rates and the performance differences between the cascaded and text-only setups.

For the dialectal test sets, we compare the ASR hypotheses to both the dialectal and the German references. We compare them to the dialectal references as they directly correspond to the audio recordings. However, the ASR models – especially the Whisper ones – tend to produce more normalized, German-like output (Dolev et al., 2024; Gorisch and Schmidt, 2024; Blaschke et al., 2025b). We therefore also compare their outputs to the German text data, although there can be lexical and syntactic differences between the German and dialectal versions. In the case of xSID, the German and Bavarian splits were independently translated from English (Winkler et al., 2024), leading to sentence pairs that are parallel on a sentence level but not necessarily on a word level (Winkler et al., 2024; Krückl et al., 2025).

For the German test sets, the closer the ASR hypotheses are to the gold text, the smaller the performance gap between the text-only and cascaded setups (Pearson’s  $r$  for error rate and cascaded/text-only gap between  $-0.72$  and  $-0.98$ ). However, even the cascaded systems with the lowest error rates do not reach the accuracies of the corresponding text-only versions (Table 6, right). For the dialectal test sets, ASR performance is also strongly

Model	$\Delta$ speech–cascaded			$\Delta$ speech–text			$\Delta$ cascaded–text			
	MAS. deu	xSID deu	xSID bar	MAS. deu	xSID deu	xSID bar	MAS. deu	xSID deu	xSID bar	Swissd. gsw
mHuBERT				-18.0 <sub>1.0</sub>	-22.0 <sub>3.6</sub>	-5.2 <sub>4.5</sub>				
XLS-R 300M				-17.0 <sub>1.0</sub>	-19.6 <sub>3.6</sub>	-4.5 <sub>4.5</sub>				
XLS-R 300M DE	+2.8 <sub>1.4</sub>	+0.3 <sub>3.9</sub>	+17.9 <sub>7.2</sub>	-2.2 <sub>1.0</sub>	-5.7 <sub>3.6</sub>	+7.2 <sub>4.5</sub>	-5.1 <sub>1.1</sub>	-6.0 <sub>1.8</sub>	-10.7 <sub>2.8</sub>	-11.6 <sub>2.4</sub>
XLS-R 1B DE							-4.0 <sub>0.7</sub>	-2.2 <sub>1.4</sub>	-9.1 <sub>2.1</sub>	-10.6 <sub>2.2</sub>
MMS 300M				-15.1 <sub>1.0</sub>	-15.8 <sub>3.6</sub>	-0.6 <sub>4.5</sub>				
MMS 1B-all							-3.9 <sub>0.7</sub>	-6.1 <sub>2.7</sub>	-10.0 <sub>0.6</sub>	-6.4 <sub>2.0</sub>
Whisper tiny	-8.5 <sub>2.4</sub>	-9.9 <sub>5.2</sub>	+5.6 <sub>8.6</sub>	-21.8 <sub>1.0</sub>	-22.1 <sub>3.6</sub>	-11.2 <sub>4.5</sub>	-13.3 <sub>2.5</sub>	-12.2 <sub>3.1</sub>	-16.9 <sub>4.9</sub>	-8.1 <sub>2.8</sub>
Whisper base	+0.0 <sub>1.3</sub>	-2.3 <sub>4.9</sub>	+16.9 <sub>9.7</sub>	-9.0 <sub>1.0</sub>	-11.3 <sub>3.6</sub>	+4.5 <sub>4.5</sub>	-9.0 <sub>1.2</sub>	-9.0 <sub>3.4</sub>	-12.4 <sub>5.8</sub>	-1.5 <sub>2.7</sub>
Whisper small	+0.6 <sub>1.5</sub>	+1.0 <sub>3.8</sub>	+11.8 <sub>8.8</sub>	-3.1 <sub>1.0</sub>	-1.5 <sub>3.6</sub>	+10.4 <sub>4.5</sub>	-3.7 <sub>0.6</sub>	-2.5 <sub>1.8</sub>	-1.4 <sub>4.6</sub>	+5.5 <sub>2.5</sub>
Whisper medium	-0.2 <sub>1.0</sub>	-1.0 <sub>3.8</sub>	+5.6 <sub>7.5</sub>	-2.0 <sub>1.0</sub>	-2.8 <sub>3.6</sub>	+3.8 <sub>4.5</sub>	-1.8 <sub>0.3</sub>	-1.8 <sub>1.0</sub>	-1.8 <sub>3.5</sub>	+8.0 <sub>2.0</sub>
Whisper large-v3-turbo							-1.2 <sub>0.3</sub>	-1.0 <sub>0.8</sub>	+2.9 <sub>2.3</sub>	+9.0 <sub>1.7</sub>
Whisper large-v3	-0.8 <sub>0.9</sub>	+1.9 <sub>4.0</sub>	+11.1 <sub>6.0</sub>	-1.7 <sub>1.0</sub>	+0.9 <sub>3.6</sub>	+13.7 <sub>4.5</sub>	-0.9 <sub>0.3</sub>	-1.0 <sub>0.7</sub>	+2.6 <sub>1.8</sub>	+8.9 <sub>1.7</sub>

Table 6: **Using text data works best for the standard-language test sets, while speech data tends to yield better results for the dialect.** Accuracy differences between setups, in percentage points. We use the results that are already averaged across random seeds for calculating differences.  $\Delta$  speech–cascaded: Each speech-only model is compared to all cascaded systems that involve the given speech model as ASR model.  $\Delta$  speech–text: Each speech-only model is compared to all text-only models.  $\Delta$  cascaded–text: Each cascaded system is compared to the text-only system involving the same text model. Standard deviations are in subscripts.

correlated with the results of the cascaded systems. However, the cascaded systems with the best ASR models outperform their text-only counterparts.

For the Swiss German data, the correlations between the error rate and the performance gap are very strong when comparing the ASR hypotheses to the German references ( $-0.94 < r < -0.99$ ): if the ASR output is normalized German (i.e., the WER/CER is low with respect to the German reference), the text models can process it well. For the Bavarian data where there might be more syntactic variation between the German and Bavarian text versions, the correlations are not quite as strong as for the Swiss German data, but they are high both when considering the German *and* the Bavarian references ( $-0.85 < r < -0.92$  and  $-0.66 < r < -0.93$ ).

For the dialectal test sets, using ASR presents on the one hand the chance of acting as dialect-to-standard normalization, and on the other hand the risk of the ASR model producing nonsensical outputs. However, the automatic ASR quality measures might give a somewhat pessimistic impression with respect to sentence classification tasks. In our manual evaluation of ASR hypotheses for 125 German and Bavarian sentences from the xSID test set by four of the ASR models, we find that keywords related to the intent are often still retained

in the transcriptions even when the sentence has a (somewhat) different meaning and/or is not in fluent and orthographically correct German (§F.1).

In the case of Swissdial, the best cascaded models perform on the Swiss German data on par with the German text-only evaluations, suggesting successful normalization. For the cascaded intent classification systems, gaps of up to 19.1 pp. between the German and Bavarian data remain even for systems with the best-performing ASR models (Table 4). For the cascaded systems with the worst ASR models, the performance is worse than processing the written dialect data.

## 5.2 Speech-only vs. cascaded: How to best deal with spoken inputs?

There are two approaches for handling spoken input data: training a speech-based system (speech-only) or training a text-based system and transcribing the spoken input data. **For the German test data, this choice barely makes a difference.** When comparing the speech-only and cascaded systems involving the same speech/ASR models, the performance gaps are close to zero (Table 6, left), with the exception of the smallest model (Whisper tiny). **For Bavarian however, the speech-only setup outperforms the cascaded setup** across the board (performance gains of 5.6 – 17.9 pp.).

### 5.3 Speech-only vs. text-only: Which modality is preferable?

The text-only models nearly always outperform the speech-only models for German (Table 6, middle). For the text and speech models of comparable sizes, the performance gaps are between  $-23.8$  and  $+1.6$  pp. ( $\sim 90$  M parameters) and between  $-23.8$  and  $-7.0$  pp. ( $\sim 300$  M parameters).<sup>8</sup> The gaps are smaller for ASR speech models (Table 4).

The results for the Bavarian data are different: **the speech models are generally more robust in standard-to-dialect transfer than the text models.** Most speech models outperform the text models that were not also pretrained on Bavarian data (i.e., all text models except mBERT). The performance gaps are between  $-3.3$  and  $+14.8$  pp. ( $\sim 90$  M parameters, ignoring mBERT) and between  $-4.7$  and  $+7.0$  pp. ( $\sim 300$  M parameters).<sup>9</sup> Except for the speech models whose performance gap is particularly large on the German data (accuracy differences of more than 15 pp.), the speech models outperform the text models on the Bavarian data.

## 6 Discussion

**Results for German vs. dialects** The text-only approach is generally best for the German data, followed by the cascaded setup and lastly the spoken setup. This is consistent with prior work finding that text-only models typically outperform speech-only models (Koudounas et al., 2023; Lee et al., 2024; Schmidt et al., 2025) and cascaded models (Lee et al., 2018; Bastianelli et al., 2020; Schmidt et al., 2025), and that – especially in cross-lingual transfer settings – cascaded models often outperform speech-only setups (Lee et al., 2024; Keller and Glavaš, 2025; Schmidt et al., 2025).

In standard-to-dialect transfer settings, however, the speech-only models perform best, and cascaded systems yield very heterogeneous results. This is more similar to the finding by Keller and Glavaš (2025) that speech-only models outperform cascaded systems in monolingual settings where the language is not supported by the ASR component of the cascaded model. Rajaa et al. (2022) find that speech-only models also outperform cascaded ones on non-native Indian English. This might be due to

potentially poor ASR performance on non-native accents (cf. Sanabria et al., 2023).

**Robustness and input representations** The speech models tend to be more robust towards dialectal variation than the text models in our experiments (§5.3). This might be due to the speech models processing continuous signals from the audio data as their input without relying on pre-defined tokens. While the dialects in our study also differ from German in lexical and morphosyntactic regards, the chief differences are on the phonological and phonetic level, and prior research has shown that multilingual speech models encode phonetically similar inputs similarly (Choi et al., 2024; Shim et al., 2025).

**Designing datasets for speech-related tasks** Creating written datasets for tasks like intent detection makes sense for high-resource languages where ASR systems work well, given that it is easier to collect written data and cascaded models perform on par with speech-only models (§5.2). However, for language varieties that are primarily spoken and not supported well by text models, collecting spoken data might be an especially promising avenue for creating realistic datasets (cf. Scharenborg et al., 2020; Chrupała, 2023).

## 7 Conclusion

Insights about processing standard languages do not necessarily generalize to even closely related dialects, as speech-only and cascaded systems can be more promising on dialect data than their results on standard-language data might indicate. We hope that our work on intent and topic classification inspires more research on expanding the burgeoning field of dialect NLP to focus on processing speech data beyond ASR.

### Limitations

**Language varieties** Our experiments involve a limited set of languages, dialects, and tasks, as these experiments require datasets that (i) have evaluation data that are parallel in a standard language and a related non-standard dialect, (ii) have parallel written and spoken evaluation data, (iii) have parallel written and spoken training data, and (iv) are labelled for a content classification task. To the best of our knowledge, we create the first dataset that fulfills all four conditions, and use all available datasets that fulfill conditions *i*, *ii*, and *iv*. We hope

<sup>8</sup>mHuB.-mBERT:  $-23.8$ , Whisper<sub>sm</sub>-XLM-R<sub>base</sub>:  $+1.6$ , XLS-R<sub>300M</sub>-XLM-R<sub>lg</sub>:  $-23.8$ , Whisper<sub>med</sub>-XLM-R<sub>lg</sub>:  $-7.0$ .

<sup>9</sup>mHuB.-XLM-R<sub>base</sub>:  $-3.3$  (if including mBERT: mHuB.-mBERT:  $-11.3$ ), Whisper<sub>small</sub>-mDeB.:  $+14.8$ , XLS-R<sub>300M</sub>-XLM-R<sub>lg</sub>:  $-4.7$ , XLS-R<sub>300M</sub>-XLM-R<sub>lg</sub>:  $+7.0$ .

our work encourages others to release more such datasets in additional languages.

**Cascaded setup** For the cascaded setup, we focus on models trained on gold-standard text data rather than on ASR transcriptions. This implementation matches that of prior work (e.g., Wang et al., 2023; Lee et al., 2024). We do this in order to keep computation feasible while being able to compare a large amount of ASR models, as we only need to train 12 models this way (four text models on three random seeds each), instead of training 108 models (nine ASR models  $\times$  four text models  $\times$  three seeds). To mitigate this limitation, we additionally train a subset of the possible cascaded configurations on automatically transcribed training data (three ASR models, two text models, three seeds). While the exact results differ from the systems trained on gold-standard text data (some of the cascaded systems trained on ASR data are stronger, but others weaker, than the cascaded systems trained on gold text), the transfer trends are the same when comparing the cascaded systems to the text-only and speech-only systems (§E).

**Models** We use encoder models because of their well-suitedness for classification tasks and to have more control over their task-specific training data (§3.3). While follow-up work on instruction-tuned text/audio LLMs would be interesting, we intentionally do not include them in our experiments. We cannot directly compare the performance of instruction-tuned text vs. audio LLMs – unlike our experiments here where we control the fine-tuning data, we would not know if the differences in the model results are due to differences in the internal representations or due to differences in possible classification-related instruction-tuning data.

We use models that we deem representative of typical text- or speech-based encoders. However, not all models use the same kinds of input representations. E.g., a recently released speech model, OpusLM (Tian et al., 2025), uses audio tokens rather than continuous representations, and pixel-based models (Rust et al., 2023; Muñoz-Ortiz et al., 2025) use continuous input representations for text. We also do not use any multimodal models that are trained to minimize representation differences between parallel text and speech input (e.g., Duquette et al., 2023).

The pre-training datasets of the speech models are unfortunately not openly accessible, but the metadata indicate no dialectal training data.

The poor performance of all ASR-capable speech models on Bavarian ASR (Table 5 and §F.2) indicates that it is unlikely these models were exposed to much Bavarian data, despite performing quite strongly on Bavarian in the speech-only setup.

## Ethical Considerations

All recordings/annotations were carried out by a speaker/annotator who was employed and paid according to national standards and who agreed to share their recordings/annotations for research use. We share the audio recordings for research on processing spoken language data, but do not permit their use in the context of speech synthesis or voice cloning.

## Acknowledgements

We thank Simone Ciciliano, Alex Fraser, and Andreas Säuberli for helpful discussions. We also thank Siyao Peng, Ryan Soh-Eun Shim, and Shijia Zhou for their feedback on a draft of this paper. This research is supported by European Research Council (ERC) Consolidator Grant DI-ALECT 101043235.

## References

- David Ifeoluwa Adelani, Hannah Liu, Xiaoyu Shen, Nikita Vassilyev, Jesujoba O. Alabi, Yanke Mao, Haonan Gao, and En-Shiun Annie Lee. 2024. [SIB-200: A simple, inclusive, and big evaluation dataset for topic classification in 200+ languages and dialects](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 226–245, St. Julian’s, Malta. Association for Computational Linguistics.
- Noëmi Aeppli, Çağrı Çöltekin, Rob Van Der Goot, Tommi Jauhiainen, Mourhaf Kazzaz, Nikola Ljubešić, Kai North, Barbara Plank, Yves Scherrer, and Marcos Zampieri. 2023. [Findings of the VarDial evaluation campaign 2023](#). In *Tenth Workshop on NLP for Similar Languages, Varieties and Dialects (VarDial 2023)*, pages 251–261, Dubrovnik, Croatia. Association for Computational Linguistics.
- William Agnew, Julia Barnett, Annie Chu, Rachel Hong, Michael Feffer, Robin Netzorg, Harry H. Jiang, Ezra Awumey, and Sauvik Das. 2024. [Sound check: Auditing audio datasets](#). *Preprint*, arXiv:2410.13114.
- Orevaoghene Ahia, Anuoluwapo Aremu, Diana Abagyan, Hila Gonen, David Ifeoluwa Adelani, Daud Abolade, Noah A. Smith, and Yulia Tsvetkov. 2024. [Voices unheard: NLP resources and models for Yorùbá regional dialects](#). In *Proceedings of the 2024*

- Conference on Empirical Methods in Natural Language Processing*, pages 4392–4409, Miami, Florida, USA. Association for Computational Linguistics.
- Reihaneh Amooie, Wietse De Vries, Yun Hao, Jelske Dijkstra, Matt Coler, and Martijn Wieling. 2025. [Enhancing standard and dialectal Frisian ASR: Multilingual fine-tuning and language identification for improved low-resource performance](#). In *2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
- Siddhant Arora, Hayato Futami, Jee-weon Jung, Yifan Peng, Roshan Sharma, Yosuke Kashiwagi, Emiru Tsunoo, Karen Livescu, and Shinji Watanabe. 2024. [UniverSLU: Universal spoken language understanding for diverse tasks with natural language instructions](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 2754–2774, Mexico City, Mexico. Association for Computational Linguistics.
- Ekaterina Artemova, Verena Blaschke, and Barbara Plank. 2024. [Exploring the robustness of task-oriented dialogue systems for colloquial German varieties](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 445–468, St. Julian’s, Malta. Association for Computational Linguistics.
- Arun Babu, Changhan Wang, Andros Tjandra, Kushal Lakhota, Qiantong Xu, Naman Goyal, Kritika Singh, Patrick von Platen, Yatharth Saraf, Juan Pino, Alexei Baevski, Alexis Conneau, and Michael Auli. 2022. [XLS-R: Self-supervised cross-lingual speech representation learning at scale](#). In *Interspeech 2022*, pages 2278–2282.
- Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. 2020. [wav2vec 2.0: A framework for self-supervised learning of speech representations](#). In *Advances in Neural Information Processing Systems*, pages 12449–12460.
- Emanuele Bastianelli, Andrea Vanzo, Pawel Swietojanski, and Verena Rieser. 2020. [SLURP: A spoken language understanding resource package](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7252–7262, Online. Association for Computational Linguistics.
- Nicole Beringer, Florian Schiel, and Peter Regel-Brietzmann. 1998. [German regional variants – a problem for automatic speech recognition?](#) In *5th International Conference on Spoken Language Processing (ICSLP 1998)*, page paper 0201.
- Verena Blaschke, Felicia Körner, and Barbara Plank. 2025a. [Add noise, tasks, or layers? MaiNLP at the VarDial 2025 shared task on Norwegian dialectal slot and intent detection](#). In *Proceedings of the 12th Workshop on NLP for Similar Languages, Varieties and Dialects*, pages 182–199, Abu Dhabi, UAE. Association for Computational Linguistics.
- Verena Blaschke, Christoph Purschke, Hinrich Schuetze, and Barbara Plank. 2024. [What do dialect speakers want? A survey of attitudes towards language technology for German dialects](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 823–841, Bangkok, Thailand. Association for Computational Linguistics.
- Verena Blaschke, Hinrich Schütze, and Barbara Plank. 2023. [Does manipulating tokenization aid cross-lingual transfer? a study on POS tagging for non-standardized languages](#). In *Tenth Workshop on NLP for Similar Languages, Varieties and Dialects (VarDial 2023)*, pages 40–54, Dubrovnik, Croatia. Association for Computational Linguistics.
- Verena Blaschke, Miriam Winkler, Constantin Förster, Gabriele Wenger-Glemser, and Barbara Plank. 2025b. [A multi-dialectal dataset for German dialect ASR and dialect-to-standard speech translation](#). In *Interspeech 2025*, pages 913–917.
- Kwanghee Choi, Ankita Pasad, Tomohiko Nakamura, Satoru Fukayama, Karen Livescu, and Shinji Watanabe. 2024. [Self-supervised speech representations are more phonetic than semantic](#). In *Interspeech 2024*, pages 4578–4582.
- Grzegorz Chrupała. 2023. [Putting natural in Natural Language Processing](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 7820–7827, Toronto, Canada. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Alice Coucke, Alaa Saade, Adrien Ball, Théodore Bluche, Alexandre Caulier, David Leroy, Clément Doumouro, Thibault Gisselbrecht, Francesco Caltagirone, Thibaut Lavril, Maël Primet, and Joseph Dureau. 2018. [Snips voice platform: An embedded spoken language understanding system for private-by-design voice interfaces](#). *Preprint*, arXiv:1805.10190.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

- Amirbek Djanibekov, Hawau Olamide Toyin, Raghad Alshalan, Abdullah Alatir, and Hanan Aldarmaki. 2025. [Dialectal coverage and generalization in Arabic speech recognition](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 29490–29502, Vienna, Austria. Association for Computational Linguistics.
- Pelin Dogan-Schönberger, Julian Mäder, and Thomas Hofmann. 2021. [SwissDial: Parallel multidialectal corpus of spoken Swiss German](#). *Computing Research Repository*, arXiv:2103.11401.
- Eyal Dolev, Clemens Lutz, and Noëmi Aepli. 2024. [Does Whisper understand Swiss German? An automatic, qualitative, and human evaluation](#). In *Proceedings of the Eleventh Workshop on NLP for Similar Languages, Varieties, and Dialects (VarDial 2024)*, pages 28–40, Mexico City, Mexico. Association for Computational Linguistics.
- Paul-Ambroise Duquenne, Holger Schwenk, and Benoît Sagot. 2023. [SONAR: Sentence-level multimodal and language-agnostic representations](#). *Preprint*, arXiv:2308.11466.
- Fahim Faisal, Sharlina Keshava, Md Mahfuz Ibn Alam, and Antonios Anastasopoulos. 2021. [SD-QA: Spoken dialectal question answering for the real world](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3296–3315, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Manaal Faruqui and Dilek Hakkani-Tür. 2022. [Revisiting the boundary between ASR and NLU in the age of conversational dialog systems](#). *Computational Linguistics*, 48(1):221–232.
- Jack FitzGerald, Christopher Hench, Charith Peris, Scott Mackie, Kay Rottmann, Ana Sanchez, Aaron Nash, Liam Urbach, Vishesh Kakarala, Richa Singh, Swetha Ranganath, Laurie Crist, Misha Britan, Wouter Leeuwis, Gokhan Tur, and Prem Nataraajan. 2023. [MASSIVE: A 1M-example multilingual natural language understanding dataset with 51 typologically-diverse languages](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4277–4302, Toronto, Canada. Association for Computational Linguistics.
- Zhipeng Gao, Akihiro Tamura, and Tsuneo Kato. 2025. [Wenzhou dialect speech to Mandarin text conversion](#). In *Proceedings of the Eighth Workshop on Technologies for Machine Translation of Low-Resource Languages (LoResMT 2025)*, pages 36–43, Albuquerque, New Mexico, U.S.A. Association for Computational Linguistics.
- Daniela Gerz, Pei-Hao Su, Razvan Kusztos, Avishek Mondal, Michał Lis, Eshan Singhal, Nikola Mrkšić, Tsung-Hsien Wen, and Ivan Vulić. 2021. [Multilingual and cross-lingual intent detection from spoken data](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7468–7475, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Jan Gorisch and Thomas Schmidt. 2024. [Evaluating workflows for creating orthographic transcripts for oral corpora by transcribing from scratch or correcting ASR-output](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 6564–6574, Torino, Italia. ELRA and ICCL.
- Mutian He and Philip Garner. 2023. [The interpreter understands your meaning: End-to-end spoken language understanding aided by speech translation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 4408–4423, Singapore. Association for Computational Linguistics.
- Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2021a. [DeBERTaV3: Improving DeBERTa using ELECTRA-style pre-training with gradient-disentangled embedding sharing](#). *Preprint*, arXiv:2111.09543.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021b. [DeBERTa: Decoding-enhanced BERT with disentangled attention](#). In *International Conference on Learning Representations*.
- Aditya Joshi, Raj Dabre, Diptesh Kanojia, Zhuang Li, Haolan Zhan, Gholamreza Haffari, and Doris Dipold. 2025. [Natural language processing for dialects of a language: A survey](#). *ACM Computing Surveys*, 57(6).
- Vani Kanjirangat, Tanja Samardzic, Ljiljana Dolamic, and Fabio Rinaldi. 2025. [Tokenization and representation biases in multilingual models on dialectal NLP tasks](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 23992–24010, Suzhou, China. Association for Computational Linguistics.
- Siva Rajesh Kasa, Karan Gupta, Sumegh Roychowdhury, Ashutosh Kumar, Yaswanth Biruduraju, Santhosh Kumar Kasa, Pattisapu Nikhil Priyatam, Arindam Bhattacharya, Shailendra Agarwal, and Vijay Huddar. 2025. [Generative or discriminative? Revisiting text classification in the era of transformers](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 9615–9637, Suzhou, China. Association for Computational Linguistics.
- Lennart Keller and Goran Glavaš. 2025. [SpeechTaxi: On multilingual semantic speech classification](#). In *2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
- Alkis Koudounas, Moreno La Quatra, Lorenzo Vaiani, Luca Colomba, Giuseppe Attanasio, Eliana Pastor, Luca Cagliero, and Elena Baralis. 2023. [ITALIC: An](#)

- Italian intent classification dataset. In *Interspeech 2023*, pages 2153–2157.
- Xaver Maria Krückl, Verena Blaschke, and Barbara Plank. 2025. [Improving dialectal slot and intent detection with auxiliary tasks: A multi-dialectal Bavarian case study](#). In *Proceedings of the 12th Workshop on NLP for Similar Languages, Varieties and Dialects*, pages 128–146, Abu Dhabi, UAE. Association for Computational Linguistics.
- Beomseok Lee, Ioan Calapodescu, Marco Gaido, Matteo Negri, and Laurent Besacier. 2024. [SpeechMASSIVE: A multilingual speech dataset for SLU and beyond](#). In *Interspeech 2024*, pages 817–821.
- Chia-Hsuan Lee, Szu-Lin Wu, Chi-Liang Liu, and Hung-yi Lee. 2018. [Spoken SQuAD: A study of mitigating the impact of speech recognition errors on listening comprehension](#). *Proc. Interspeech 2018*, pages 3459–3463.
- Liam Loneragan, Mengjie Qian, Neasa Ní Chiaráin, Christer Gobl, and Ailbhe Ní Chasaide. 2023. [Towards dialect-inclusive recognition in a low-resource language: Are balanced corpora the answer?](#) In *Interspeech 2023*, pages 5082–5086.
- Sanad Malaysha, Mo El-Haj, Saad Ezzini, Mohammed Khalilia, Mustafa Jarrar, Sultan Almujaewel, Ismail Berrada, and Houda Bouamor. 2024. [AraFinNLP 2024: The first Arabic financial NLP shared task](#). In *Proceedings of the Second Arabic Natural Language Processing Conference*, pages 393–402, Bangkok, Thailand. Association for Computational Linguistics.
- Andrew McDowell. 2022a. [Wav2vec2-xls-r-1b fine-tuned on CommonVoice8 - DE](#). Hugging Face, <https://huggingface.co/AndrewMcDowell/wav2vec2-xls-r-1b-german>.
- Andrew McDowell. 2022b. [Wav2vec2-xls-r-300m fine-tuned on CommonVoice7 - DE](#). Hugging Face, <https://huggingface.co/AndrewMcDowell/wav2vec2-xls-r-300m-german-de>.
- 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.
- Jessica Ojo, Odunayo Ogundepo, Akintunde Oladipo, Kelechi Ogueji, Jimmy Lin, Pontus Stenetorp, and David Ifeoluwa Adelani. 2025. [AfroBench: How good are large language models on African languages?](#) In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 19048–19095, Vienna, Austria. Association for Computational Linguistics.
- Orestis Papakyriakopoulos, Anna Seo Gyeong Choi, William Thong, Dora Zhao, Jerone Andrews, Rebecca Bourke, Alice Xiang, and Allison Koenecke. 2023. [Augmented datasheets for speech datasets and ethical decision-making](#). In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, pages 881–904, New York, NY, USA. Association for Computing Machinery.
- Emmy Phung, Harsh Deshpande, Ahmad Emami, and Kanishk Singh. 2024. [AR-NLU: A framework for enhancing natural language understanding model robustness against ASR errors](#). In *Interspeech 2024*, pages 1325–1329.
- Vineel Pratap, Andros Tjandra, Bowen Shi, Paden Tomasello, Arun Babu, Sayani Kundu, Ali Elkahky, Zhaoheng Ni, Apoorv Vyas, Maryam Fazel-Zarandi, Alexei Baevski, Yossi Adi, Xiaohui Zhang, Wei-Ning Hsu, Alexis Conneau, and Michael Auli. 2024. [Scaling speech technology to 1,000+ languages](#). *JMLR*, 25(97):1–52.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2023. [Robust speech recognition via large-scale weak supervision](#). In *ICML’23: Proceedings of the 40th International Conference on Machine Learning*. JMLR.org.
- Shangeth Rajaa, Swaraj Dalmia, and Kumarmanas Nethil. 2022. [Skit-S2I: An Indian accented speech to intent dataset](#). *Preprint*, arXiv:2212.13015.
- Phillip Rust, Jonas F Lotz, Emanuele Bugliarello, Elizabeth Salesky, Miryam de Lhoneux, and Desmond Elliott. 2023. [Language modelling with pixels](#). In *The Eleventh International Conference on Learning Representations*.
- Dan Saatrup Nielsen, Kenneth Enevoldsen, and Peter Schneider-Kamp. 2025. [Encoder vs decoder: Comparative analysis of encoder and decoder language models on multilingual NLU tasks](#). In *Proceedings of the Joint 25th Nordic Conference on Computational Linguistics and 11th Baltic Conference on Human Language Technologies (NoDaLiDa/Baltic-HLT 2025)*, pages 561–572, Tallinn, Estonia. University of Tartu Library.
- Ramon Sanabria, Nikolay Bogoychev, Nina Markl, Andrea Carmantini, Ondrej Klejch, and Peter Bell. 2023. [The Edinburgh International Accents of English Corpus: Towards the democratization of English ASR](#). In *2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
- Odette Scharenborg, Laurent Besacier, Alan Black, Mark Hasegawa-Johnson, Florian Metze, Graham Neubig, Sebastian Stüker, Pierre Godard, Markus Müller, Lucas Ondel, Shruti Palaskar, Philip Arthur, Francesco Ciannella, Mingxing Du, Elin Larsen, Danny Merx, Rachid Riad, Liming Wang, and Emmanuel Dupoux. 2020. [Speech technology for unwritten languages](#). *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28:964–975.
- Yves Scherrer, Rob van der Goot, and Petter Mæhlum. 2025. [Findings of the VarDial evaluation campaign 2025: The NorSID shared task on Norwegian slot, intent and dialect identification](#). In *Proceedings of*

- the 12th Workshop on NLP for Similar Languages, Varieties and Dialects*, pages 1–8, Abu Dhabi, UAE. Association for Computational Linguistics.
- Fabian David Schmidt, Ivan Vulić, Goran Glavaš, and David Ifeoluwa Adelani. 2025. [Fleurs-SLU: A massively multilingual benchmark for spoken language understanding](#). *Preprint*, arXiv:2501.06117.
- Sebastian Schuster, Sonal Gupta, Rushin Shah, and Mike Lewis. 2019. [Cross-lingual transfer learning for multilingual task oriented dialog](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3795–3805, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ryan Soh-Eun Shim, Domenico De Cristofaro, Chengzhi Martin Hu, Alessandro Vietti, and Barbara Plank. 2025. [Languages in multilingual speech foundation models align both phonetically and semantically](#). *Preprint*, arXiv:2505.19606.
- Clément Sicard, Victor Gillioz, and Kajetan Pyszkowski. 2023. [Spaiche: Extending state-of-the-art ASR models to Swiss German dialects](#). In *Proceedings of the 8th edition of the Swiss Text Analytics Conference*, pages 76–83, Neuchatel, Switzerland. Association for Computational Linguistics.
- Aarohi Srivastava and David Chiang. 2023. [Fine-tuning BERT with character-level noise for zero-shot transfer to dialects and closely-related languages](#). In *Tenth Workshop on NLP for Similar Languages, Varieties and Dialects (VarDial 2023)*, pages 152–162, Dubrovnik, Croatia. Association for Computational Linguistics.
- Aarohi Srivastava and David Chiang. 2025. [We’re calling an intervention: Exploring fundamental hurdles in adapting language models to nonstandard text](#). In *Proceedings of the Tenth Workshop on Noisy and User-generated Text*, pages 45–56, Albuquerque, New Mexico, USA. Association for Computational Linguistics.
- Zhiyuan Tang, Dong Wang, Yanguang Xu, Jianwei Sun, Xiaoning Lei, Shuaijiang Zhao, Cheng Wen, Xingjun Tan, Chuandong Xie, Shuran Zhou, Rui Yan, Chenjia Lv, Yang Han Han, Wei Zou, and Xiangang Li. 2021. [KeSpeech: An open source speech dataset of Mandarin and its eight subdialects](#). In *NeurIPS 2021 Datasets and Benchmarks Track*.
- Louis ten Bosch. 2000. [ASR, dialects, and acoustic/phonological distances](#). In *6th International Conference on Spoken Language Processing (ICSLP 2000)*, pages vol. 3, 1009–1012.
- Jinchuan Tian, William Chen, Yifan Peng, Jiatong Shi, Siddhant Arora, Shikhar Bharadwaj, Takashi Maekaku, Yusuke Shinohara, Keita Goto, Xiang Yue, Huck Yang, and Shinji Watanabe. 2025. [OpusLM: A family of open unified speech language models](#). In *Interspeech 2025*, pages 3259–3263. ISCA.
- Melissa Torgbi, Andrew Clayman, Jordan J. Speight, and Harish Tayyar Madabushi. 2025. [Adapting Whisper for regional dialects: Enhancing public services for vulnerable populations in the United Kingdom](#). In *Proceedings of the 12th Workshop on NLP for Similar Languages, Varieties and Dialects*, pages 29–38, Abu Dhabi, UAE. Association for Computational Linguistics.
- Socrates Vakirtzian, Chara Tsoukala, Stavros Bompolas, Katerina Mouzou, Vivian Stamou, Georgios Paraskevopoulos, Antonios Dimakis, Stella Markantonatou, Angela Ralli, and Antonios Anastasopoulos. 2024. [Speech recognition for Greek Dialects: A challenging benchmark](#). In *Interspeech 2024*, pages 3974–3978.
- Rob van der Goot, Ibrahim Sharaf, Aizhan Imankulova, Ahmet Üstün, Marija Stepanović, Alan Ramponi, Siti Oryza Khairunnisa, Mamoru Komachi, and Barbara Plank. 2021. [From masked language modeling to translation: Non-English auxiliary tasks improve zero-shot spoken language understanding](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2479–2497, Online. Association for Computational Linguistics.
- Minghan Wang, Yinglu Li, Jiaxin Guo, Xiaosong Qiao, Zongyao Li, Hengchao Shang, Daimeng Wei, Shimin Tao, Min Zhang, and Hao Yang. 2023. [WhiSLU: End-to-end spoken language understanding with Whisper](#). In *Interspeech 2023*, pages 770–774.
- Miriam Winkler, Virginija Juozapaityte, Rob van der Goot, and Barbara Plank. 2024. [Slot and intent detection resources for Bavarian and Lithuanian: Assessing translations vs natural queries to digital assistants](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 14898–14915, Torino, Italia. ELRA and ICCL.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, and 3 others. 2020. [Huggingface’s transformers: State-of-the-art natural language processing](#). *Preprint*, arXiv:1910.03771.
- Aditya Yadavalli, Ganesh Mirishkar, and Anil Kumar Vuppala. 2022. [Multi-task end-to-end model for Telugu dialect and speech recognition](#). In *Interspeech 2022*, pages 1387–1391.
- Marcos Zampieri, Preslav Nakov, and Yves Scherrer. 2020. [Natural language processing for similar languages, varieties, and dialects: A survey](#). *Natural Language Engineering*, 26(6):595–612.
- Marcely Zanon Boito, Vivek Iyer, Nikolaos Lagos, Laurent Besacier, and Ioan Calapodescu. 2024.

## A Data Statement

### A.1 Summary and motivation

- Dataset name and version: xSID-audio, version 0.1
- Dataset curators: Verena Blaschke, Miriam Winkler, Barbara Plank
- Dataset and data statement citation: Please cite this paper.
- Data statement version: 1.0

### Executive summary and curation rationale

This dataset extends the slot and intent classification dataset xSID by providing audio recordings for 300 development and 500 test sentences in German and Bavarian. Its purpose is to allow research on intent classification for spoken dialectal data.

### A.2 Documentation for source datasets

We use the written German and Bavarian test datasets from xSID 0.5 (van der Goot et al., 2021; Aepli et al., 2023; Winkler et al., 2024; CC BY-SA 4.0). xSID itself is derived from Snips (Coucke et al., 2018; CC0 1.0 Universal) and a dataset by Schuster et al. (2019; CC-BY-SA). For details on the German translation, see van der Goot et al. (2021); for details on the Bavarian translation, Winkler et al. (2024).

### A.3 Composition

Our dataset consists of spoken sentences with intent labels. All of the sentences are from the development and test splits of xSID.

**Date** The original sentences are likely from ca. 2018 (Coucke et al., 2018; Schuster et al., 2019). The German translation (van der Goot et al., 2021) is from ca. 2020 and the Bavarian translation (Winkler et al., 2024) is from 2023. Our recordings are from 2024.

**Modality** Spoken (read speech).

**Genre** Questions and commands for a virtual personal assistant.

Split		N	# Int.	Duration (s)
<i>xSID-audio 0.1</i>				
de	dev	300	15*	3.33±1.05
	test	500	15*	3.47±1.00
de-ba	dev	300	15*	3.04±0.98
	test	500	15*	3.56±0.99
<i>Subset used for experiments in this paper</i>				
de (“deu”)	test	412	10	3.48±1.05
de-ba (“bar”)	test	412	10	3.52±1.02

Table 7: **Dataset statistics for xSID-audio.** Intents: \*The development and test splits of xSID(-audio) contain 16 intent labels in total, but only 15 per split (14 are shared across splits). Duration: mean duration ± standard deviation.

**Language varieties and speaker** We record audios in Standard German spoken in Germany (“de”) and in the dialectal variant of Bavarian as spoken in rural Upper Bavaria (“de-ba”). The audios were recorded by a native speaker of both German and Bavarian in her 20s. The speaker is the same person who translated the text version of xSID de-ba.

**Dataset statistics** The dataset consists of one German and one Bavarian audio recording for each of the 300 development and 500 test sentences. On average, each audio clip in xSID-audio 0.1 has a duration of 3.4 s for German and 3.3 s for Bavarian. Detailed statistics are in Table 7.

**Labels** Each sentence is annotated with one of 16 intent labels, although the development set and test set each only use a subset of 15 labels. The labels are identical to the ones in xSID. The label distribution for the subset of instances used in the experiments in this paper is in Table 10 in §C.1. For annotation details, see van der Goot et al. (2021).

**Data quality** We directly worked with xSID’s German and Bavarian sentences. The German sentences contain a few intentional grammatical errors to reflect errors in the original English data (van der Goot et al., 2021). The Bavarian dataset does not contain such errors (Winkler et al., 2024). The recordings directly reflect the written versions.

### A.4 Audio recording and formatting

We record the samples with a Røde NT-USB microphone with pop filter in a quiet room. The microphone was set up at the desk and the speaker spoke from a distance of about 20 cm. The instances were

read one sentence at a time directly from the dataset text with breaks in between. After recording, we convert the audio files to the wav mulaw format, resulting in stereo audios with a recording rate of 48 kHz and a bit rate of 768 kBit/s.

The audios are in a folder each for each language and dataset split with corresponding TSV files containing the sentence ID, intent, written sentence and audio filename.

### A.5 Distribution, use, and maintenance

**Distribution and usage terms** We release the dataset on Zenodo ([10.5281/zenodo.19554427](https://zenodo.org/record/105281/files/zenodo.19554427)) under the [CC BY-SA 4.0](https://creativecommons.org/licenses/by-sa/4.0/) license (in accordance with xSID’s license). We prohibit the use of xSID-audio for voice cloning or speech synthesis.

**Other uses** The dataset could also be used for slot filling (in combination with the slot annotations in xSID). It could further be used for ASR and for phonetic studies.

**Maintenance and updates** We have not planned any specific updates. If you find errors in the dataset, please email Verena Blaschke or Barbara Plank.

### A.6 Limitations, acknowledgements, and further information

**Limitations** In order to be parallel to the already existing test version of xSID, this dataset consists of read speech. It is therefore not representative of spontaneous speech.

This dataset reflects the dialect of a single speaker.

**Disclosures and acknowledgements** There are no conflicts of interest. This research is supported by European Research Council (ERC) Consolidator Grant DIALECT 101043235.

**About this data statement** A data statement is a characterization of a dataset that provides context to allow developers and users to better understand how experimental results might generalize, how software might be appropriately deployed, and what biases might be reflected in systems built on the software.

This data statement was written based on the template for the Data Statements Version 3 Schema. The template was prepared by Angelina McMillan-Major and Emily M. Bender and can be found at <http://techpolicylab.uw.edu/data-statements>.

We adapted the structure and questions of the template and integrated parts of the audio-specific data statements by [Papakyriakopoulos et al. \(2023\)](#) and [Agnew et al. \(2024\)](#).

## B Resource Links and Licenses

Table 8 provides license details and resource links for the models and datasets we use. We use all resources in accordance with their intended use.

The MASSIVE dataset ([FitzGerald et al., 2023](#)) is a multilingual, text-only extension of SLURP ([Bastianelli et al., 2020](#)). Speech-MASSIVE ([Lee et al., 2024](#)) adds audio recordings for several of the languages in MASSIVE. The Bavarian translation of a subset of MASSIVE that we use was introduced in [Winkler et al. \(2024\)](#), where it is called “MASSIVE de:ba”. For details on the source datasets that xSID is built on, see [A.2](#).

## C Data Preprocessing

### C.1 Mapping the intent labels

We map the (Speech-)MASSIVE labels to the ones used by xSID: If a sentence in MASSIVE was manually re-annotated with an xSID label by [Winkler et al. \(2024\)](#), we use that label. Otherwise, if it can be mapped according to Table 9, we update the label accordingly. Otherwise, we exclude the sentence. Sentences with labels that occur less than ten times in MASSIVE’s training data are excluded as well. The resulting label distribution is shown in Table 10.

### C.2 Excluded topic labels

When using SwissDial, we remove sentences with labels that are not content-based topics: Random (these appear to be randomly chosen sentences from Wikipedia) and Special (sentences showcasing lexical variation across dialects). We also exclude sentences with labels that occur very rarely: Earth/Space and Story.

## D Hyperparameters and Training

We use HuggingFace’s *transformers* library ([Wolf et al., 2020](#), v4.51.3, Apache 2.0 license) to fine-tune our models. We use GPUs of type NVIDIA H200 for fine-tuning the Whisper large-v3 models and NVIDIA A100-SXM4-80GB for the remaining models.

We use the training and development datasets of the German splits of (Speech-)MASSIVE and

Resource	License	Hugging Face ID / Dataset link
<i>Models</i>		
mBERT	Apache 2.0	google-bert/bert-base-multilingual-cased
mDeBERTa	MIT	microsoft/mdeberta-v3-base
XLM-R base	MIT	FacebookAI/xlm-roberta-base
XLM-R large	MIT	FacebookAI/xlm-roberta-large
mHuBERT	CC BY-NC-SA 3.0	utter-project/mHuBERT-147
XLS-R 300M	Apache 2.0	facebook/wav2vec2-xls-r-300m
XLS-R 300M DE	Apache 2.0	AndrewMcDowell/wav2vec2-xls-r-300m-german-de
XLS-R 1B DE	Apache 2.0	AndrewMcDowell/wav2vec2-xls-r-1B-german
MMS 300M	CC BY-NC 4.0	facebook/mms-300m
MMS 1B-all	CC BY-NC 4.0	facebook/mms-1b-all
Whisper tiny	Apache 2.0	openai/whisper-tiny
Whisper base	Apache 2.0	openai/whisper-base
Whisper small	Apache 2.0	openai/whisper-small
Whisper medium	Apache 2.0	openai/whisper-medium
Whisper large-v3	Apache 2.0	openai/whisper-large-v3
Whisper large-v3-turbo	MIT	openai/whisper-large-v3-turbo
<i>Datasets</i>		
MASSIVE	Apache 2.0	github.com/alexa/massive
Speech-MASSIVE	CC BY-NC-SA 4.0	github.com/hlt-nt/Speech-MASSIVE
MASSIVE de:ba	not specified	github.com/mainlp/NaLiBaSID
xSID	CC BY-SA 4.0	github.com/mainlp/xsid
xSID-audio	CC BY-SA 4.0	doi.org/10.5281/zenodo.19554428
SwissDial	CC BY-NC 4.0	mtc.ethz.ch/publications/open-source/swiss-dial.html

Table 8: **Links and licenses for the resources we use/release.**

MASSIVE label	→	xSID label
play_music		PlayMusic
recommendation_movies		SearchScreeningEvent
(if slots contain place_name		and/or business_type)
alarm_remove		alarm/cancel_alarm
alarm_set		alarm/set_alarm
alarm_query		alarm/show_alarms
calendar_remove		reminder/cancel_reminder
(if English text includes “remind”		and/or “notif”)
calendar_set		reminder/set_reminder
(if English text includes “remind”		and/or “notif”)
calendar_query		reminder/show_reminders
(if English text includes “remind”)		
weather_query		weather/find

Table 9: **Intent label mappings.** See §C.1 for additional mapping details.

Intent	(Speech-)MASSIVE			xSID
	Train	Dev	Test	Test
AddToPlaylist	24	4	8	34
PlayMusic	923	168	105	39
BookRestaurant	23	3	1	43
SearchScreeningEvent	43	9	9	37
alarm/cancel_alarm	71	13	21	33
alarm/set_alarm	173	30	40	29
alarm/show_alarms	132	21	31	19
reminder/set_reminder	350	58	41	37
reminder/show_reminders	156	27	32	19
Total	2468	459	361	412

Table 10: Label distribution in the intent classification datasets after mapping the labels.

Swissdial for selecting the learning rate and maximum number of training epochs. We train each model configuration on three random seeds, both for hyperparameter tuning and for the main experiments. We always choose the model epoch with the highest development accuracy. The hyperparameter values are based on prior work on written and spoken intent classification (Rajaa et al., 2022; Koudounas et al., 2023; Arora et al., 2024; Lee et al., 2024; Schmidt et al., 2025). We otherwise use the default values of the *TrainingArguments* class.

- Maximum number of training epochs (written input): 10
- Maximum number of training epochs (spoken input): 30 (used in hyperparameter search). Maximum number picked based on hyperparameter experiments: 15 (Whisper small, Whisper large-v3), 17 (Whisper medium), 19 (Whisper tiny, Whisper base), 23 (XLS-R 300M DE), 29 (MMS 300M, mHuBERT), 30 (XLS-R 300M)
- Learning rate: {1e-5, 5e-5, 1e-4}. Learning rates chosen:
  - Intents: 1e-5 (XLM-R base, mDeBERTa, Whisper medium, Whisper large-v3), 5e-5 (mBERT, Whisper small, XLS-R 300M, MMS 300M, mHuBERT), 1e-4 (XLM-R large, XLS-R 300M DE, Whisper tiny, Whisper base)
  - Topics: 1e-5 (XLM-R large), 5e-5 (mBERT, XLM-R base), 1e-4 (mDeBERTa)
- Batch size: 32
- Warm-up ratio: 0.1
- Weight decay: 0.01

For the intent classification task, we observe that the text models are much more robust to the learning rate choice than the speech models (Table 11).

## E Cascaded Setups Trained on ASR Data

When testing on automatically transcribed data, it can make a difference whether the model was fine-tuned on gold text data or on transcriptions produced by the same ASR model (Lee et al., 2018; Phung et al., 2024).

Intents			
Model	1e-5	5e-5	1e-4
mBERT	95.6 <sub>0.2</sub>	<b>96.0</b> <sub>0.3</sub>	95.6 <sub>0.3</sub>
mDeBERTa	<b>97.3</b> <sub>0.1</sub>	97.2 <sub>0.6</sub>	96.9 <sub>0.4</sub>
XLM-R base	<b>97.5</b> <sub>0.3</sub>	97.2 <sub>0.3</sub>	96.8 <sub>0.3</sub>
XLM-R large	96.7 <sub>0.4</sub>	96.4 <sub>1.8</sub>	<b>96.9</b> <sub>0.2</sub>
mHuBERT	67.1 <sub>1.6</sub>	<b>85.1</b> <sub>1.4</sub>	49.0 <sub>21.5</sub>
XLS-R 300M	42.4 <sub>0.7</sub>	<b>86.3</b> <sub>0.6</sub>	81.7 <sub>2.6</sub>
XLS-R 300M DE <sup>ASR</sup>	88.0 <sub>0.9</sub>	95.1 <sub>0.7</sub>	<b>95.4</b> <sub>0.4</sub>
MMS 300M	66.7 <sub>0.8</sub>	<b>83.0</b> <sub>0.7</sub>	74.4 <sub>9.5</sub>
Whisper tiny <sup>ASR</sup>	71.2 <sub>0.3</sub>	74.1 <sub>0.8</sub>	<b>77.1</b> <sub>0.5</sub>
Whisper base <sup>ASR</sup>	74.8 <sub>2.5</sub>	85.4 <sub>4.3</sub>	<b>89.4</b> <sub>1.9</sub>
Whisper small <sup>ASR</sup>	91.1 <sub>2.1</sub>	<b>94.0</b> <sub>0.3</sub>	92.3 <sub>1.1</sub>
Whisper medium <sup>ASR</sup>	<b>95.3</b> <sub>0.8</sub>	94.4 <sub>0.8</sub>	93.5 <sub>1.1</sub>
Whisper large-v3 <sup>ASR</sup>	<b>94.1</b> <sub>0.9</sub>	93.1 <sub>1.0</sub>	92.2 <sub>0.3</sub>

Topics			
Model	1e-5	5e-5	1e-4
mBERT	55.0 <sub>2.4</sub>	<b>57.7</b> <sub>1.8</sub>	55.5 <sub>1.1</sub>
mDeBERTa	44.2 <sub>1.7</sub>	54.3 <sub>3.1</sub>	<b>57.4</b> <sub>2.5</sub>
XLM-R base	56.0 <sub>1.3</sub>	<b>58.1</b> <sub>0.3</sub>	55.2 <sub>0.9</sub>
XLM-R large	<b>60.0</b> <sub>0.8</sub>	34.7 <sub>23.5</sub>	21.1 <sub>0.0</sub>

Table 11: **Effect of the learning rate on development accuracies (in %)**. For the intent classification task, the text models are more robust to learning rate differences than the speech models, and the ASR speech models (<sup>ASR</sup>) are more robust than the non-ASR ones. The best development score per model and task is bolded. Scores are averaged over three random seeds, with standard deviations in subscripts.

To supplement our experiments on cascaded setups, we additionally train six cascaded systems (three ASR models  $\times$  two text models) on training data that was automatically transcribed by the same ASR models as the evaluation data. We select this subset since training on all combinations of ASR and text models is not feasible. We use the two text models with the best and worst performance on the Bavarian data (mBERT and mDeBERTa, respectively). For the ASR models, we select the overall best and worst ASR model (Whisper large-v3 and tiny, respectively), and additionally XLS-R 300M DE so as to also include a CTC-based ASR model. We only carry out these supplemental experiments on the intent classification data, since no spoken training data is available for the topic classification dataset.

How big the performance difference is between

ASR model	F $\uparrow$	M $\uparrow$	K $\uparrow$	A $\uparrow$	Acc $\uparrow$	WER $\downarrow$
<i>German audio</i>						
X. 300M DE	46.4	46.4	76.8	32.0	76.9 <sub>4.6</sub>	29.8 <sub>26.9</sub>
MMS 1B-all	53.6	45.6	81.6	36.8	75.9 <sub>5.3</sub>	30.8 <sub>30.0</sub>
Whisper small	88.8	68.0	90.4	64.8	79.9 <sub>4.5</sub>	19.3 <sub>28.8</sub>
W. large-v3	98.4	84.0	97.6	84.0	80.3 <sub>5.3</sub>	8.6 <sub>14.4</sub>
<i>Bavarian audio</i>						
X. 300M DE	2.4	3.2	28.8	2.4	45.8 <sub>8.7</sub>	90.5 <sub>21.7</sub>
MMS 1B-all	4.0	5.6	33.6	3.2	47.1 <sub>7.3</sub>	89.8 <sub>29.3</sub>
Whisper small	14.4	8.0	45.6	4.8	53.2 <sub>11.1</sub>	85.7 <sub>30.7</sub>
W. large-v3	29.6	20.0	56.8	14.4	59.9 <sub>7.3</sub>	81.2 <sub>24.5</sub>

Table 12: **Manual ASR evaluation of a subset of xSID.** All values in %. Manual evaluation: F=fluency, M=meaning, K=keywords, A=all combined (F and M and K are all yes). The “Acc” column contains the intent classification accuracy of the cascaded models on the same subset, averaged over text models and random seeds.

ASR model	Train	MASSIVE		xSID			
		deu	$\Delta$	deu	$\Delta$	bar	$\Delta$
<i>mBERT as text model</i>							
XLS-R	gold	89.8 <sub>1.3</sub>	0.5	79.6 <sub>2.3</sub>	-2.3	57.5 <sub>2.5</sub>	-5.7
	ASR	90.3 <sub>1.0</sub>		77.3 <sub>0.7</sub>		51.8 <sub>1.4</sub>	
W. tiny	gold	82.4 <sub>1.5</sub>	4.2	72.8 <sub>1.7</sub>	1.5	50.6 <sub>2.9</sub>	4.8
	ASR	86.6 <sub>1.4</sub>		74.3 <sub>1.6</sub>		55.3 <sub>1.9</sub>	
W. large	gold	93.3 <sub>1.0</sub>	0.6	83.3 <sub>0.4</sub>	1.6	68.8 <sub>2.6</sub>	-1.4
	ASR	93.9 <sub>0.0</sub>		84.9 <sub>0.8</sub>		67.4 <sub>1.8</sub>	
<i>mDeBERTa as text model</i>							
XLS-R	gold	86.9 <sub>1.0</sub>	-0.9	71.4 <sub>0.4</sub>	-1.9	41.1 <sub>1.2</sub>	6.6
	ASR	86.0 <sub>3.7</sub>		69.5 <sub>5.3</sub>		47.7 <sub>2.0</sub>	
W. tiny	gold	76.6 <sub>0.9</sub>	-0.6	62.3 <sub>1.7</sub>	-1.1	30.3 <sub>1.1</sub>	9.4
	ASR	76.1 <sub>3.2</sub>		61.2 <sub>4.6</sub>		39.6 <sub>2.1</sub>	
W. large	gold	92.7 <sub>1.0</sub>	-5.9	77.0 <sub>1.0</sub>	-9.9	54.4 <sub>2.9</sub>	-1.9
	ASR	86.8 <sub>4.7</sub>		67.2 <sub>7.4</sub>		52.5 <sub>2.1</sub>	

Table 13: **Performance differences between cascaded models that were trained on gold-standard vs. automatically transcribed data.** Each row shows the test-set performance of one cascaded setup (accuracy, in %, averaged over three seeds with standard deviations in subscripts). The  $\Delta$  columns show the differences in percentage points: trained on data transcribed by the same ASR model as the evaluation data minus trained on gold text data. ASR models: XLS-R 300M DE, Whisper tiny, Whisper large-v3.

Text model	Speech/ASR mod.	$\Delta$ speech-casc.			$\Delta$ cascaded-text		
		MAS. deu	xSID deu	xSID bar	MAS. deu	xSID deu	xSID bar
<i>Train on gold text data, evaluate on ASR data</i>							
mBERT	XLS-R	1.0	-3.6	8.5	-4.2	-3.9	-7.4
	W. tiny	-11.1	-13.2	-3.0	-11.7	-10.7	-14.4
	W. large	-1.8	-0.6	3.7	-0.8	-0.2	3.8
mDeB.	XLS-R	4.0	4.6	24.9	-6.6	-7.5	-13.3
	W. tiny	-5.4	-2.7	17.3	-16.9	-16.7	-24.1
	W. large	-1.3	5.7	18.1	-0.8	-1.9	0.0
<i>Train/eval on data transcribed by the same ASR model</i>							
mBERT	XLS-R	0.6	-1.3	14.2	-3.8	-6.1	-13.2
	W. tiny	-15.3	-14.6	-7.8	-7.5	-9.2	-9.6
	W. large	-2.5	-2.2	5.1	-0.2	1.4	2.4
mDeB.	XLS-R	4.9	6.6	18.3	-7.6	-9.5	-6.6
	W. tiny	-4.8	-1.6	7.9	-17.5	-17.7	-14.7
	W. large	4.6	15.5	20.0	-6.7	-11.8	-1.9

Table 14: **The overall trends are similar for cascaded systems that were fine-tuned on gold-standard text data (top) vs. automatically transcribed data (bottom).** Accuracy differences between setups, in percentage points. Text models: mBERT and mDeBERTa, speech/ASR models: XLS-R 300M DE, Whisper tiny, Whisper large-v3.

the cascaded models trained on gold-standard vs. automatically transcribed data – and which version actually yields the better results – varies across the ASR and text model combinations (Table 13). While the ASR errors in the training data help some models better generalize to the evaluation data, this is not the case for all models: the ASR transcriptions might on the one hand introduce infelicitous noise into the training data, and on the other hand, the types of ASR errors can be different across data splits (see §F.2).

Comparing the performance differences between cascaded and text-/speech-only models results in somewhat different deltas, but the trends are the same regardless of whether the cascaded models are trained on gold or transcribed text (Table 14). As in the overall analysis (§5, Table 6), the speech-only setups generally outperform the cascaded ones on the Bavarian test data, but not necessarily on the German data. Most cascaded-setups are outperformed by the text-only models on both German and Bavarian.

Thus, while training cascaded models on automatically transcribed data results in somewhat different predictions than when training them on gold text data, the overall patterns stay the same.

## F Details on ASR Results

### F.1 Manual Evaluation of ASR Results

We annotate the ASR hypotheses of four models (XLS-R 300M DE, MMS 1B-all, Whisper small, and Whisper large-v3) for the first 75 and last 50 instances of the xSID test set in German and Bavarian based on the following criteria:

- **Fluency:** Is the ASR hypothesis in fluent, grammatical, and orthographically correct German? Colloquial and regionally marked language is okay, major spelling errors are not (unless they concern named entities).
- **Meaning:** Does the hypothesis have the same meaning as the reference? Spelling errors are okay as long as it would still be clear out of context what the sentence means, but not if they concern key entities (i.e., we keep slot filling in mind).
- **Keywords:** Does the hypothesis contain obvious clues for the intent label? This can be the case even if the more nuanced meaning of the sentence did not get preserved. The keywords signalling the intent should be spelled correctly (so as to match the spelling in the training data) and should not be merged with other words in the hypothesis due to wrong word segmentation, unless they occur at the beginning of the merged word.  
This measure is from the annotator’s perspective, but the pertinent keywords (keyword subtokens) / combinations thereof learned by the models might be different. This measure is inspired by [Choi et al. \(2024\)](#), whose work suggests that keyword spotting is often sufficient for intent classification.

We annotate each instance with *yes* or *no* for each of the criteria. If in doubt, we choose *no*. We ignore capitalization (since the model input is lowercased, §3.2). We also ignore punctuation. For each measure, we calculate the proportion of instances that were annotated as *yes*. The annotator is a native speaker of German who is familiar with Bavarian and with the dataset.

**Results** Table 12 shows the results. All of these measures are highly correlated with each other and nearly always show the same model ranking. However, the *keyword* measure almost always has the highest score, and it is always much higher than the

*meaning* score (with differences of 13.6–37.6 pp.). This shows that even if the meaning of the transcribed sentence does not match the original meaning, the sentence often still contains keywords related to the intent. This is even the case for Bavarian sentences whose transcriptions neither match the original meaning nor are in fluent German.

### F.2 ASR Sample Outputs

Table 15 shows ASR sample transcriptions. While the German transcriptions are usually understandable, the transcriptions for Bavarian audio contain many more mistakes.

## G De-aggregated Results

Table 16 shows the topic classification per Swiss German dialect in Swissdial. Table 17 shows the detailed results for the cascaded systems (by text and ASR model). Both tables complement Table 4 in §4.

---

“Show (me) all reminders”

DEU	Zeige alle Erinnerungen
W. lg	Zeige alle Erinnerungen.
W. tiny	Zeige <b>aller</b> Erinnerungen. (“show all.GEN reminders”)
XLS-R	zeige alle erinnerungen
BAR	Zoag ma olle Erinnerungen
W. lg	<b>Zeug</b> mir alle Erinnerungen. (“beget me all reminders”)
W. tiny	<b>Talk mal</b> alle <b>ja in der Wungen</b> . (“[nonce] once all yes in the [nonce]”)
XLS-R	<b>zorgmal olle</b> erinnerungen (“[nonce] [‘old’ in DEU / ‘all’ in BAR] reminders”)

---

“Add a reminder for 4pm today”

DEU	Eine Erinnerung für heute um 4 Uhr nachmittags hinzufügen
W. lg	Eine Erinnerung für heute um 4 Uhr nachmittags hinzufügen.
W. tiny	Eine Erinnerung für heute um vier Uhr nachmittags hinzufügen.
XLS-R	eine erinnerung für heute um vier uhr nachmittags hinzufügen
BAR	Stei ma a Erinnerung fua heid um viere Nammiddog ei
W. lg	Stell <b>mal</b> Erinnerung für <b>Heidumphiri</b> <b>Namidogai</b> . (“set reminder for [nonce] [nonce]”)
W. tiny	<b>Stämmerer</b> Erinnerung für <b>Heidung</b> <b>4 Nm</b> . (“[nonce] reminder for [nonce] 4 [nonce]”)
XLS-R	<b>steimeerinnerung</b> für <b>heidung</b> <b>virenamidock ei</b> (“[nonce]reminder for [nonce] [nonce] [nonce in DEU / ‘in’ in BAR]”)

---

“Show all alarms”

DEU	Zeige alle Wecker
W. lg	Zeige alle Wecker.
W. tiny	Zeige alle Wecker.
XLS-R	zeige alle <b>weger</b> (“show all [nonce]”)
BAR	Zoag alle Wegga
W. lg	<b>Zwerg-Olewecke</b> (“dwarf-[nonce]”)
W. tiny	<b>Zuhack und lege</b> . (“[nonce] and lay”)
XLS-R	<b>zorg ollyvicker</b> (“[nonce] [nonce]”)

---

Table 15: **Sample gold sentences and ASR outputs** for German (DEU) and Bavarian (BAR) audios by Whisper large-v3, Whisper-tiny, and XLS-R 300M DE. Wrong transcriptions are in **red**; transcriptions that are wrong in German but accurate in Bavarian are in **blue**.

Model	deu	gsw	ag	be	bs	gr	lu	sg	vs	zh
<i>Text-only</i>	54.0 <sub>3.3</sub>	43.7 <sub>3.8</sub>	41.9 <sub>3.3</sub>	40.0 <sub>4.5</sub>	46.4 <sub>3.9</sub>	46.3 <sub>4.7</sub>	41.8 <sub>4.2</sub>	45.6 <sub>4.4</sub>	42.7 <sub>3.7</sub>	44.9 <sub>3.4</sub>
mBERT	52.8 <sub>2.1</sub>	41.2 <sub>1.0</sub>	39.7 <sub>0.6</sub>	37.5 <sub>2.1</sub>	43.9 <sub>0.8</sub>	43.3 <sub>1.2</sub>	39.6 <sub>2.4</sub>	42.2 <sub>1.9</sub>	41.0 <sub>0.8</sub>	42.6 <sub>1.1</sub>
mDeBERTa	51.4 <sub>1.2</sub>	43.6 <sub>2.6</sub>	40.7 <sub>1.4</sub>	41.1 <sub>1.5</sub>	48.0 <sub>4.1</sub>	45.5 <sub>2.8</sub>	40.9 <sub>2.7</sub>	45.5 <sub>2.9</sub>	41.2 <sub>2.4</sub>	45.5 <sub>3.5</sub>
XLNet base	53.3 <sub>3.2</sub>	41.3 <sub>4.0</sub>	40.7 <sub>4.2</sub>	35.7 <sub>4.1</sub>	43.4 <sub>4.2</sub>	44.1 <sub>5.5</sub>	39.2 <sub>3.6</sub>	43.4 <sub>4.4</sub>	41.3 <sub>4.7</sub>	42.7 <sub>2.8</sub>
XLNet large	58.3 <sub>1.3</sub>	48.7 <sub>1.0</sub>	46.4 <sub>0.6</sub>	45.6 <sub>1.9</sub>	50.3 <sub>1.7</sub>	52.4 <sub>1.5</sub>	47.5 <sub>1.7</sub>	51.4 <sub>1.0</sub>	47.1 <sub>2.3</sub>	48.9 <sub>0.6</sub>
<i>Cascaded</i>		42.9 <sub>8.7</sub>	42.3 <sub>9.4</sub>	43.1 <sub>9.9</sub>	43.7 <sub>7.8</sub>	45.6 <sub>7.5</sub>	42.4 <sub>9.3</sub>	43.8 <sub>7.9</sub>	39.9 <sub>9.3</sub>	42.7 <sub>9.7</sub>
<i>... averaged over ASR models</i>										
mBERT		41.5 <sub>8.5</sub>	40.2 <sub>9.3</sub>	41.9 <sub>9.7</sub>	42.2 <sub>7.6</sub>	44.4 <sub>7.2</sub>	40.5 <sub>8.9</sub>	42.9 <sub>7.7</sub>	38.9 <sub>9.3</sub>	41.2 <sub>9.2</sub>
mDeBERTa		40.8 <sub>8.8</sub>	39.7 <sub>9.0</sub>	41.3 <sub>10.2</sub>	42.0 <sub>7.6</sub>	43.4 <sub>7.9</sub>	39.9 <sub>9.6</sub>	41.4 <sub>8.0</sub>	38.3 <sub>9.2</sub>	40.7 <sub>9.7</sub>
XLNet base		43.1 <sub>8.4</sub>	42.9 <sub>9.1</sub>	42.9 <sub>10.0</sub>	44.1 <sub>7.9</sub>	45.4 <sub>7.2</sub>	42.6 <sub>8.3</sub>	43.4 <sub>7.4</sub>	40.3 <sub>9.3</sub>	43.3 <sub>9.8</sub>
XLNet large		46.3 <sub>8.6</sub>	46.3 <sub>9.4</sub>	46.3 <sub>9.5</sub>	46.4 <sub>7.9</sub>	49.3 <sub>6.7</sub>	46.7 <sub>9.5</sub>	47.4 <sub>7.6</sub>	42.3 <sub>9.5</sub>	45.4 <sub>10.0</sub>
<i>... averaged over text models</i>										
XLS-R 300M DE		32.1 <sub>3.1</sub>	31.0 <sub>3.1</sub>	30.6 <sub>3.3</sub>	34.3 <sub>3.4</sub>	36.5 <sub>3.2</sub>	32.6 <sub>4.5</sub>	33.5 <sub>3.8</sub>	27.9 <sub>3.2</sub>	30.6 <sub>3.1</sub>
XLS-R 1B DE		33.1 <sub>3.5</sub>	33.0 <sub>4.7</sub>	31.0 <sub>3.7</sub>	36.6 <sub>4.2</sub>	37.2 <sub>4.2</sub>	33.2 <sub>3.9</sub>	35.5 <sub>3.6</sub>	28.2 <sub>3.2</sub>	30.1 <sub>3.1</sub>
MMS 1B-all		37.3 <sub>3.3</sub>	36.9 <sub>3.9</sub>	35.1 <sub>3.8</sub>	38.6 <sub>2.5</sub>	42.3 <sub>4.6</sub>	34.8 <sub>3.6</sub>	41.4 <sub>3.8</sub>	33.2 <sub>2.9</sub>	35.9 <sub>3.2</sub>
Whisper tiny		35.6 <sub>2.3</sub>	32.6 <sub>2.7</sub>	35.6 <sub>3.1</sub>	36.0 <sub>2.4</sub>	39.7 <sub>3.8</sub>	34.5 <sub>2.7</sub>	36.6 <sub>2.7</sub>	34.3 <sub>2.3</sub>	35.2 <sub>2.2</sub>
Whisper base		42.2 <sub>2.3</sub>	41.6 <sub>3.2</sub>	44.9 <sub>2.7</sub>	42.0 <sub>2.6</sub>	44.0 <sub>2.6</sub>	40.2 <sub>3.4</sub>	41.8 <sub>2.5</sub>	39.9 <sub>3.0</sub>	43.4 <sub>3.1</sub>
Whisper small		49.2 <sub>2.5</sub>	48.2 <sub>4.2</sub>	51.2 <sub>2.4</sub>	49.2 <sub>2.8</sub>	50.7 <sub>2.5</sub>	48.6 <sub>2.8</sub>	49.3 <sub>3.2</sub>	45.3 <sub>2.2</sub>	51.1 <sub>2.0</sub>
Whisper medium		51.7 <sub>2.9</sub>	52.0 <sub>3.5</sub>	52.3 <sub>2.8</sub>	51.1 <sub>2.6</sub>	52.3 <sub>2.6</sub>	52.3 <sub>4.2</sub>	51.1 <sub>3.1</sub>	50.0 <sub>3.3</sub>	52.5 <sub>2.5</sub>
Whisper large-v3-turbo		52.6 <sub>3.1</sub>	52.5 <sub>3.9</sub>	53.7 <sub>2.9</sub>	52.5 <sub>2.9</sub>	53.9 <sub>2.7</sub>	52.9 <sub>4.0</sub>	52.4 <sub>3.3</sub>	50.0 <sub>2.9</sub>	52.8 <sub>3.5</sub>
Whisper large-v3		52.7 <sub>3.1</sub>	52.6 <sub>3.8</sub>	53.5 <sub>2.9</sub>	53.0 <sub>3.0</sub>	54.1 <sub>3.1</sub>	53.0 <sub>4.2</sub>	52.4 <sub>3.2</sub>	50.8 <sub>2.5</sub>	52.5 <sub>3.1</sub>

Table 16: **Topic classification accuracies for German, Swiss German (all dialects), and Swiss German by region.** Key: deu = German, gsw = Swiss German (all of the following dialects), ag = Aargau, be = Bern, bs = Basel, gr = Grisons, lu = Lucerne, sg = St. Gallen, vs = Valais, zh = Zurich. Accuracies in %, averaged over three random seeds (standard deviations in subscripts). Average performances per setup type are in *italics*.

Text model	ASR model	Intents				Topics
		MASSIVE	xSID			Swissdial
			deu	deu	bar	
						gsw
mBERT	XLS-R 300M DE	89.8 <sub>1.3</sub>	79.6 <sub>2.3</sub>	57.5 <sub>2.5</sub>	-22.1	30.6 <sub>0.8</sub>
	XLS-R 1B DE	90.7 <sub>0.4</sub>	82.8 <sub>1.0</sub>	57.8 <sub>0.6</sub>	-25.0	31.4 <sub>0.7</sub>
	MMS 1B-all	90.5 <sub>1.4</sub>	79.1 <sub>1.3</sub>	54.0 <sub>2.1</sub>	-25.1	35.3 <sub>0.5</sub>
	Whisper tiny	82.4 <sub>1.5</sub>	72.8 <sub>1.7</sub>	50.6 <sub>2.9</sub>	-22.2	34.6 <sub>0.9</sub>
	Whisper base	85.8 <sub>2.1</sub>	77.2 <sub>2.0</sub>	56.2 <sub>1.2</sub>	-21.0	41.2 <sub>1.2</sub>
	Whisper small	90.9 <sub>1.0</sub>	82.2 <sub>1.4</sub>	67.6 <sub>2.5</sub>	-14.6	47.8 <sub>1.0</sub>
	Whisper medium	92.2 <sub>0.7</sub>	81.9 <sub>0.9</sub>	65.5 <sub>2.3</sub>	-16.3	50.5 <sub>1.4</sub>
	Whisper large-v3-turbo	93.1 <sub>0.8</sub>	83.3 <sub>0.0</sub>	69.6 <sub>1.8</sub>	-13.7	51.2 <sub>1.3</sub>
	Whisper large-v3	93.3 <sub>1.0</sub>	83.3 <sub>0.4</sub>	68.8 <sub>2.6</sub>	-14.5	51.2 <sub>1.4</sub>
mDeBERTa	XLS-R 300M DE	86.9 <sub>1.0</sub>	71.4 <sub>0.4</sub>	41.1 <sub>1.2</sub>	-30.3	30.0 <sub>4.3</sub>
	XLS-R 1B DE	88.5 <sub>1.9</sub>	75.4 <sub>1.2</sub>	42.8 <sub>0.6</sub>	-32.6	30.5 <sub>4.1</sub>
	MMS 1B-all	88.8 <sub>1.2</sub>	71.2 <sub>0.4</sub>	44.5 <sub>2.4</sub>	-26.7	35.0 <sub>2.9</sub>
	Whisper tiny	76.6 <sub>0.9</sub>	62.3 <sub>1.7</sub>	30.3 <sub>1.1</sub>	-32.0	33.5 <sub>2.1</sub>
	Whisper base	82.7 <sub>0.9</sub>	65.9 <sub>2.3</sub>	33.7 <sub>0.7</sub>	-32.2	40.1 <sub>1.8</sub>
	Whisper small	89.8 <sub>0.8</sub>	74.5 <sub>0.5</sub>	47.1 <sub>2.8</sub>	-27.4	47.3 <sub>0.9</sub>
	Whisper medium	92.0 <sub>0.8</sub>	75.9 <sub>0.4</sub>	47.3 <sub>3.4</sub>	-28.6	49.6 <sub>1.5</sub>
	Whisper large-v3-turbo	92.5 <sub>0.7</sub>	76.9 <sub>1.1</sub>	54.0 <sub>3.3</sub>	-22.8	50.6 <sub>1.8</sub>
	Whisper large-v3	92.7 <sub>1.0</sub>	77.0 <sub>1.0</sub>	54.4 <sub>2.9</sub>	-22.7	50.7 <sub>1.7</sub>
XLM-R base	XLS-R 300M DE	88.4 <sub>1.7</sub>	73.5 <sub>4.3</sub>	44.1 <sub>2.6</sub>	-29.4	32.7 <sub>2.8</sub>
	XLS-R 1B DE	88.9 <sub>1.3</sub>	77.3 <sub>6.1</sub>	46.8 <sub>4.1</sub>	-30.5	33.4 <sub>3.2</sub>
	MMS 1B-all	89.8 <sub>2.2</sub>	75.4 <sub>5.3</sub>	47.3 <sub>3.8</sub>	-28.1	37.5 <sub>3.3</sub>
	Whisper tiny	80.3 <sub>2.5</sub>	69.2 <sub>7.4</sub>	41.4 <sub>2.7</sub>	-27.8	36.3 <sub>2.7</sub>
	Whisper base	84.5 <sub>2.0</sub>	72.7 <sub>6.0</sub>	44.7 <sub>2.4</sub>	-28.0	43.0 <sub>2.4</sub>
	Whisper small	89.7 <sub>1.2</sub>	77.9 <sub>6.7</sub>	54.1 <sub>4.1</sub>	-23.8	49.9 <sub>3.4</sub>
	Whisper medium	90.6 <sub>1.7</sub>	77.9 <sub>7.2</sub>	56.3 <sub>5.2</sub>	-21.6	51.3 <sub>3.8</sub>
	Whisper large-v3-turbo	91.1 <sub>1.5</sub>	77.9 <sub>7.2</sub>	59.9 <sub>3.8</sub>	-18.0	51.9 <sub>4.0</sub>
	Whisper large-v3	91.6 <sub>2.3</sub>	77.8 <sub>7.1</sub>	59.9 <sub>3.8</sub>	-18.0	52.1 <sub>3.6</sub>
XLM-R large	XLS-R 300M DE	87.0 <sub>3.1</sub>	78.5 <sub>3.7</sub>	49.7 <sub>4.9</sub>	-28.8	35.2 <sub>0.7</sub>
	XLS-R 1B DE	88.1 <sub>2.4</sub>	82.8 <sub>2.7</sub>	51.4 <sub>6.0</sub>	-31.4	37.2 <sub>0.6</sub>
	MMS 1B-all	87.7 <sub>3.0</sub>	77.0 <sub>4.9</sub>	49.4 <sub>4.7</sub>	-27.7	41.3 <sub>0.9</sub>
	Whisper tiny	80.0 <sub>2.8</sub>	73.9 <sub>2.9</sub>	45.6 <sub>5.6</sub>	-28.3	37.8 <sub>1.1</sub>
	Whisper base	83.4 <sub>1.7</sub>	75.3 <sub>4.2</sub>	51.1 <sub>6.2</sub>	-24.2	44.6 <sub>0.9</sub>
	Whisper small	87.3 <sub>1.3</sub>	82.4 <sub>2.2</sub>	60.8 <sub>5.9</sub>	-21.6	51.8 <sub>1.2</sub>
	Whisper medium	90.1 <sub>1.5</sub>	84.3 <sub>1.5</sub>	58.9 <sub>7.2</sub>	-25.4	55.2 <sub>1.0</sub>
	Whisper large-v3-turbo	90.6 <sub>1.7</sub>	85.2 <sub>1.9</sub>	63.5 <sub>5.6</sub>	-21.7	56.5 <sub>0.9</sub>
	Whisper large-v3	91.2 <sub>1.9</sub>	85.0 <sub>2.4</sub>	62.5 <sub>6.8</sub>	-22.5	56.8 <sub>0.9</sub>

Table 17: **Detailed results for the cascaded systems** (accuracy, in %). Results are averaged over three random seeds; standard deviations in subscripts. Models are trained on German text and evaluated on all available speech test sets. For a version of this table aggregated over text or ASR models, see Table 4.