

# How do Multimodal Foundation Models Encode Text and Speech? An Analysis of Cross-Lingual and Cross-Modal Representations

Hyunji Lee

Danni Liu

Supriti Sinhamahapatra

Jan Niehues

Karlsruhe Institute of Technology, Germany

hyunji.lee@student.kit.edu, {firstname.lastname}@kit.edu

## Abstract

Multimodal foundation models aim to create a unified representation space that abstracts away from surface features like language syntax or modality differences. To investigate this, we study the internal representations of three recent models, analyzing the model activations from semantically equivalent sentences across languages in the text and speech modalities. Our findings reveal that: **1)** Cross-modal representations converge over model layers, except in the initial layers specialized at text and speech processing. **2)** Length adaptation is crucial for reducing the cross-modal gap between text and speech, although current approaches' effectiveness is primarily limited to high-resource languages. **3)** Speech exhibits larger cross-lingual differences than text. **4)** For models not explicitly trained for modality-agnostic representations, the modality gap is more prominent than the language gap.

## 1 Introduction

Recent progress in foundation models has sparked growing interest in expanding their text processing capabilities (NLLB Team et al., 2022; Chiang et al., 2023; Yang et al., 2024) to speech (Seamless Communication et al., 2023; Chu et al., 2024; Tang et al., 2024; Dubey et al., 2024). Despite the empirical successes, understandings of these models' internal representations remain limited, particularly on *language differences*, *modality gaps*, and the impact of *model architectures*. This work aims to fill this gap by studying how text and speech are represented in recent multimodal foundation models.

While the internal representations of multilingual models have been extensively studied, most prior works focus on single-modality analyses of text (Kudugunta et al., 2019; Sun et al., 2023) or speech (Belinkov and Glass, 2017; de Seyssel et al., 2022; Sicherman and Adi, 2023; Sun et al., 2023; Kheir et al., 2024). Moreover, as many multimodal




Figure 1: We use the similarity between model activations for the same sentences in different languages and modalities to measure language and modality gaps.

foundation models have dedicated subparts for languages or modalities, not all analysis techniques are directly applicable. For instance, similarity retrieval tasks (Conneau et al., 2020; Wang et al., 2023; Chen et al., 2023a) often require identical input feature dimensions, which is not always guaranteed for speech and text. Probing (Adi et al., 2017; Belinkov and Glass, 2017; de Seyssel et al., 2022) features with different dimensions leads to auxiliary classifiers of varying sizes and may skew the results. In this work, we use Singular Vector Canonical Correlation Analysis (SVCCA; Raghu et al., 2017) due to its invariance to affine transformations, which is suitable for comparing features from different architectures and dimensions that occurs frequently in speech-text representations.

Previous studies comparing speech-text representations do not involve the cross-lingual aspect, and use either task-specific (Dinh et al., 2022; Tsiamas et al., 2024) or proprietary (Wang et al., 2023) models. Recent studies on multilingual representations in large language models reveal different levels of language invariance depending on training data (Wendler et al., 2024) and model scales (Zeng et al., 2024). The study most related to ours is probably the concurrent work from Wu et al. (2024), who show that representations for semantically equivalent multilingual/modal inputs are similar in model intermediate layers. Our study differs in its focus on languages and modality gaps as well as its broad coverage of 30 languages at differ-

ent resource levels. To the best of our knowledge, we present the first cross-modal and cross-lingual analysis of representations over a wide variety of language in multimodal foundation models.

## 2 Methodology

An assumption in unified multimodal and multilingual models is that inputs are transformed into a semantic space independent of input forms. This abstraction from surface level motivates our method.

**Measuring similarity between semantically equivalent sentences:** As shown in Figure 1, we begin with semantically equivalent sentences in different languages and modalities. To compare their model activations at different layers, we extract these activations and employ SVCCA. Its invariance to affine transformations (Raghu et al., 2017) ensures comparability of activations across different modalities and languages, even when they originate from different model subparts. Given the extracted activations, we calculate the SVCCA scores between speech and text versions of the same sentence (intra-lingual cross-modal) and between translations (cross-lingual text/speech). Higher SVCCA scores indicate higher similarity. More explanations of SVCCA scores are in Appendix A.

**Model selection:** Model architectures introduce inductive biases in the learned representations. For speech and text representations, a critical factor is the significant length difference between speech utterances and texts. To explore different *architectures* and, in particular, *length adaptation* mechanisms, we analyze the following models:

- **Seamless** (Seamless Communication et al., 2023): encoder-decoder model with dedicated text and speech encoders, where the latter is followed by a *length adaptor* (Zhao et al., 2022) to *downsample* by a fixed factor. We analyze its encoder representations, as the decoder does not support parallel comparisons speech and text.
- **SONAR** (Duquenne et al., 2023): sentence embedding model with a multilingual text encoder and a set of monolingual speech encoders. It creates *fixed-size* embeddings by *pooling* over sequence lengths, and is explicitly trained to align multilingual and multimodal embeddings.
- **SALMONN** (Tang et al., 2024): decoder-only LLM (Vicuna; Chiang et al., 2023) adapted to ingest audio inputs. It *downsamples* encoded audio

representations<sup>1</sup> by *window-level Q-Former* (Li et al., 2023; Tang et al., 2024) by a fixed factor. We do not analyze the audio encoders’ internal representations as they are audio-only.

Detailed model descriptions and our hidden representation extraction procedures are in Appendix B.

**Data and language:** We use the FLEURS dataset (Conneau et al., 2022), which contains  $n$ -way parallel speech dataset with their transcripts from the FLoRes-101 dataset (Goyal et al., 2022). We use its test split and analyze 30 languages from diverse resource levels, language families, and scripts as detailed in Appendix C. Due to differences in supported languages among the models, six of the 30 languages are not shared between SONAR and the others. To maximize comparability, we select these six languages to have the same resource level.

**Baseline similarity:** We calculate SVCCA scores between random vectors of the same sizes as the analyzed representations as baselines. This represents the state of no similarity at all.

## 3 Results

We analyze cross-modal (§3.1) and cross-lingual (§3.2) representations, and compare the impact of modality and language differences (§3.3).

### 3.1 Cross-Modal Analysis

Figure 2 shows the average speech-text similarity of language grouped by language resource levels.

**Progression through layers:** Generally, cross-modal similarity increases with the number of layers, as expected. This suggests a growing abstraction of semantic meaning independent of the input modalities. However, all three models consistently exhibit a dip in cross-modal similarity at the initial layers. We believe this is related to the different functionalities of the earlier layers in audio and text processing models. While the early layers of text processing models primarily capture syntactic information (Belinkov et al., 2017; Peters et al., 2018), audio encoders tend to focus on acoustic features like speaker identity in the lower layers (Chung et al., 2019; Chen et al., 2022). After this initial specialized processing, representations for both modalities exhibit more similarity based on their semantics. Moreover, cross-modal similarity for SALMONN flattens and drops slightly at later

<sup>1</sup>encoded by the encoder of Whisper (Radford et al., 2023) and the BEATS encoder (Chen et al., 2023b) (both frozen)




Figure 2: Average cross-modal similarity over all languages over model layers. X-axis markers “in”: input word embeddings or audio features, “len”: after length adaptor in Seamless, “pool”: after pooling in SONAR, “enc”: after the frozen audio encoder, before length adaptation by window-level Q-Former in SALMONN.




Figure 3: Average cross-lingual similarities between all language pairs in speech/text modality over model layers.

layers. This is likely due to its decoder-only architecture which generates text outputs only and makes the model diverge from the shared representations that are learned for both modalities.

**Impact of language resource level:** As shown in the lower part of Figure 2, the overall trend of increasing similarity over the layers remains consistent across all resource levels. However, lower-resource languages consistently exhibit lower similarity scores, suggesting that they are less effectively mapped into a shared representation space than their higher-resource counterparts.

**Impact of length adaptor:** In Figure 2, by comparing the yellow crosses with their preceding data points, we can assess the impact of the length adaptors. SONAR’s pooling mechanism, coupled with its dedicated losses, are the most effective in minimizing the modality gap.<sup>2</sup> For Seamless and SALMONN, while their length adaptor and window-level Q-Former exhibit a slight positive impact in reducing the modality gap, this effect appears limited to high- and medium-resource languages. In low-resource languages, as evidenced by the flat slope towards the yellow crosses in Figure 2, these length adaptation mechanisms do not seem to be as effective. This limitation may be attributed to weaker representations for speech in lower-resource languages, hindering the learning of effective shrinking mechanisms.

<sup>2</sup>This complete elimination of the length difference also limits the model’s expressiveness and therefore performance on downstream tasks, as shown in Duquenne et al. (2023).

### 3.2 Cross-Lingual Analysis

After assessing the representational differences across modalities (§3.1), we hold modality constant and examine cross-lingual differences.

**Higher overall cross-lingual similarity in text than speech:** Figure 3 shows the average cross-lingual similarities for speech and text across model layers. Overall, with the exception of the initial layers in SALMONN, the higher cross-lingual similarities observed for text suggest that the models more effectively create a unified cross-lingual space for the text modality compared to speech. This is likely due to the greater variability in speech, as the same utterance can be expressed in various ways, different in vocal characteristics, speaking pace, and recording conditions. In contrast, text typically adheres to a single, standardized form of writing. This greater variability can pose more challenges in abstracting towards semantic representations independent of input languages.

**Initial drop in SALMONN text cross-lingual similarity due to fragmented tokenization:** We suppose that the initial drop in cross-lingual similarity in the text modality within Figure 3 is related to insufficient tokenizer coverage for diverse languages. As Vicuna’s vocabulary size is limited to 32k (inherited from LLaMA (Touvron et al., 2023a)), many languages with diverse scripts are inadequately supported. This results in texts being tokenized at the character or byte level, which are shared across many languages, inflating initial similarity in the input embeddings but settling soon

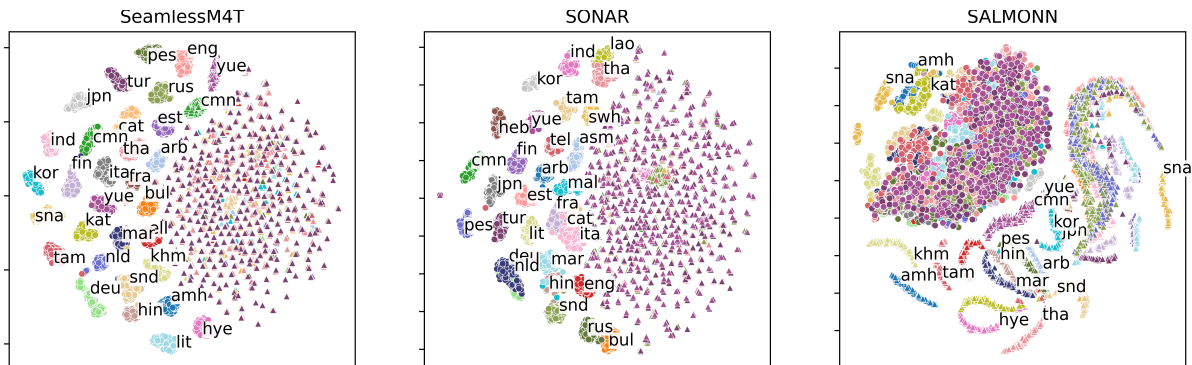


Figure 4: To visually verify how the models progressively process language and modality gaps, we use 2D visualization with t-SNE (van der Maaten and Hinton, 2008) for speech and text at a middle layer (14<sup>th</sup>, 14<sup>th</sup>, 18<sup>th</sup> from left to right). For Seamless and SONAR, texts are organized by semantics while speech remains clustered by language or language family. For SALMONN, languages with diverse scripts remain distinct in text representations.

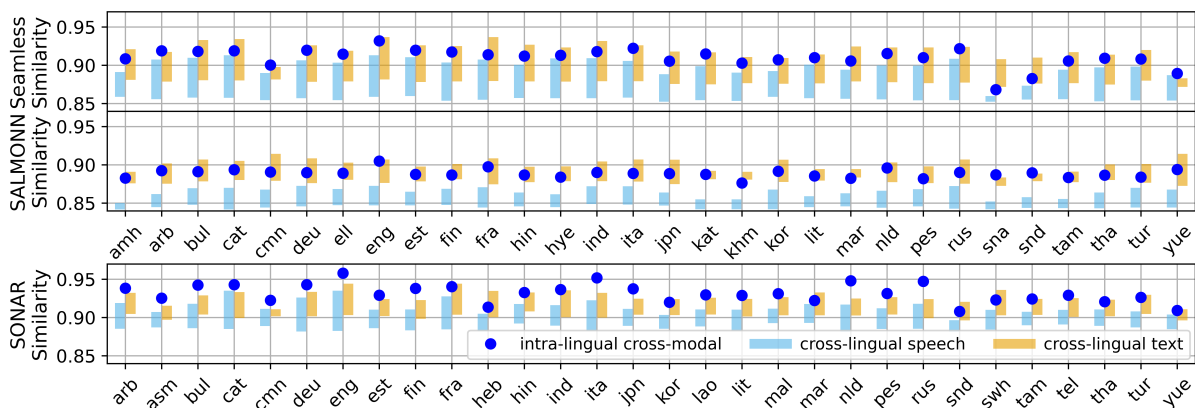


Figure 5: Given representations of text sentences at the last layer in one language, similarity to the same sentences in speech (“intra-lingual cross-modal”), their translations in text (“cross-lingual text”), and their translations in speech (“cross-lingual speech”). Latter two shown as range over all 29 language pairs. Language codes in Table 1.

in the subsequent layers. To visually verify this hypothesis, we use t-SNE plots. The visualizations for SALMONN in Figure 4 further support the tokenizer deficiency, as text representations that remain distinct are predominantly languages with diverse scripts, such as Khmer (khm), Armenian (hye), and Japanese (jpn). To further support this finding, we quantified the relation between fragmented tokenization and similarity scores after the word embedding layer. We calculated the Pearson correlation coefficient between the proportion of shared tokens between parallel sentences in two languages (averaged over all sentences in FLoRes devtest set) and their pairwise similarity scores on SALMONN. Over the 435 language pairs, we found a positive correlation coefficient of 0.228 ( $p$ -value=1.48e-06).<sup>3</sup>

<sup>3</sup>This relation may be even stronger after disentangling the effects of resource level. Fragmented tokenization often occurs in lower-resource languages, which in general have lower similarity scores (as shown in Figure 2).

**Language gap reduced earlier in text than speech:** In the other plots in Figure 4 for Seamless and SONAR, the language gap appears to be reduced earlier in text than in speech. While text data points are primarily organized by semantics in a middle layer, speech data points are still clustered by language or language family, as evidenced by clusters by language or language families like Sindhi (snd) and Hindi (hin). This aligns with our previous findings on higher overall cross-lingual similarity in text compared to speech.

### 3.3 Comparing Language and Modality Gaps

Our previous analyses have held language or modality constant and compared cross-modal and cross-lingual differences. A logical next step is to study the relative influence of the modality gap and the language gap. As shown in Figure 5, for Seamless and SALMONN, intra-lingual cross-modal similarity (between text and speech of the same sentence) is mostly always lower than the highest



cross-lingual text similarity (between a sentence and its text translation). This means that for texts in a given language, there exists a text translation in another, presumably related, language that is representationally more similar than the same sentences in speech. It also implies that for these models, the modality gap is larger than the language gap. This is somewhat counter-intuitive since the intra-lingual cross-modal setup involves the same language, but can be explained by previous findings on modality differences such as length mismatch.

The picture is slightly different with SONAR, which was explicitly trained to bridge modality and language gaps. For most languages, intra-lingual cross-modal similarity surpasses other types of similarity, suggesting that SONAR more effectively reduces the modality gap than the language gap. The different observations from SONAR highlight that, unless explicitly optimized for reduction, both modality and language gaps persist in multimodal foundation models, with the modality gap often being more pronounced than language gaps.

## 4 Conclusion

To study how multimodal foundation models process text and speech across diverse languages, we analyzed their internal representations based on the similarity of semantically equivalent sentences. Our findings highlight that while these models present a *unified architecture* for handling various modalities and languages, they do not inherently create *fully unified representations* by semantic meaning. Representational gaps, some of which already observed in task-specific models, including speech-text length mismatches (e.g., Gaido et al. 2021; Zhao et al. 2022), weak representation for low-resource languages, and tokenizer bottleneck (e.g., Zhang et al. 2022, Salesky et al. 2023), still persist in current multimodal foundation models.

Besides the findings presented earlier, our study offers several practical recommendations. The first is to incorporate representation analyses into the development cycle of models, especially on models designed to reduce modality gaps. Another recommendation is model choices for speech-text downstream tasks. Practitioners working with low-resource or zero-shot use cases may consider initializing their models with foundation models explicitly trained for closing modality and language gaps.

## Limitations

**Data and Modality Coverage** First, our study is limited by its reliance on multiway aligned text and speech data, which is scarce. Specifically, our findings are based on the FLEURS dataset (Conneau et al., 2022), which is created from Wikipedia texts. This may limit the generalizability of our findings to other domains, such as informal or spoken texts. Additionally, this study focuses on two modalities of speech and text. Exploring other modalities like images would be very interesting. However, as our research question focuses on speech-text foundation models, we consider analyzing other modalities out of the scope of the current work.

**Model Coverage** In this study, we analyzed three multimodal foundation models widely-adopted at the beginning of this project. Since then, many new multimodal models supporting text and speech have emerged, such as Qwen-Audio (Chu et al., 2024) and Llama 3 (Dubey et al., 2024). Extending our analyses to more of these models would be a valuable addition. Nonetheless, we believe our coverage of encoder-decoder, sentence embedding, and decoder-only architectures, including the decoder-only SALMONN model, provides a sufficiently diverse representation of model types.

**Language Coverage** Due to differences in supported languages among the analyzed models, six of the 30 languages are not shared between SONAR and the two models. A fully overlapping set of languages would have provided a cleaner experimental setup. However, since our conclusions are based on comparisons of similarity scores within the same model between modalities and languages, rather than across different models, we believe that the differing sets of languages do not compromise the validity of our findings.

**Analysis Type** Our findings are only drawn from intrinsic analyses based on feature vectors, i.e., SVCCA scores on activations. Additional results by other explainability methods will complement the current findings, e.g., Logit lens (nostalgebraist, 2020) or performance on downstream tasks.

**Prompt Variation** We do not vary prompts in the experiments on SALMONN, meanwhile NLLB and SONAR do not support prompting. Prompting multimodal large language models itself is activate research field, and to the best of our knowledge, there is no established prompt for bridging speech-

text modality gaps. Given the analysis-oriented nature of this work, we did not focus on prompt optimization. However, it would be interesting as the field of MLLM prompting advances.

## Acknowledgments

Part of this work was performed on the HoreKa supercomputer funded by the Ministry of Science, Research and the Arts Baden-Württemberg and by the Federal Ministry of Education and Research. Part of this work was supported by funding from the pilot program Core-Informatics of the Helmholtz Association (HGF). Part of this work received support from the European Union’s Horizon research and innovation programme under grant agreement No 101135798, project Meetween (My Personal AI Mediator for Virtual MEETtings BetWEEN People).

## References

- Yossi Adi, Einat Kermany, Yonatan Belinkov, Ofer Lavi, and Yoav Goldberg. 2017. [Fine-grained analysis of sentence embeddings using auxiliary prediction tasks](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Yonatan Belinkov, Nadir Durrani, Fahim Dalvi, Hassan Sajjad, and James Glass. 2017. [What do neural machine translation models learn about morphology?](#) In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 861–872, Vancouver, Canada. Association for Computational Linguistics.
- Yonatan Belinkov and James R. Glass. 2017. [Analyzing hidden representations in end-to-end automatic speech recognition systems](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 2441–2451.
- Mingda Chen, Kevin Heffernan, Onur Çelebi, Alexandre Mourachko, and Holger Schwenk. 2023a. [xSIM++: An improved proxy to bitext mining performance for low-resource languages](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 101–109, Toronto, Canada. Association for Computational Linguistics.
- Sanyuan Chen, Chengyi Wang, Zhengyang Chen, Yu Wu, Shujie Liu, Zhuo Chen, Jinyu Li, Naoyuki Kanda, Takuya Yoshioka, Xiong Xiao, Jian Wu, Long Zhou, Shuo Ren, Yanmin Qian, Yao Qian, Jian Wu, Michael Zeng, Xiangzhan Yu, and Furu Wei. 2022. [Wavlm: Large-scale self-supervised pre-training for full stack speech processing](#). *IEEE J. Sel. Top. Signal Process.*, 16(6):1505–1518.
- Sanyuan Chen, Yu Wu, Chengyi Wang, Shujie Liu, Daniel Tompkins, Zhuo Chen, Wanxiang Che, Xiangzhan Yu, and Furu Wei. 2023b. [Beats: Audio pre-training with acoustic tokenizers](#). In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 5178–5193. PMLR.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. [Vicuna: An open-source chatbot impressing gpt-4 with 90%\\* chatgpt quality](#).
- Yunfei Chu, Jin Xu, Qian Yang, Haojie Wei, Xipin Wei, Zhifang Guo, Yichong Leng, Yuanjun Lv, Jinzheng He, Junyang Lin, Chang Zhou, and Jingren Zhou. 2024. [Qwen2-audio technical report](#). *CoRR*, abs/2407.10759.
- Yu-An Chung, Wei-Ning Hsu, Hao Tang, and James R. Glass. 2019. [An unsupervised autoregressive model for speech representation learning](#). In *20th Annual Conference of the International Speech Communication Association, Interspeech 2019, Graz, Austria, September 15-19, 2019*, pages 146–150. ISCA.
- Yu-An Chung, Yu Zhang, Wei Han, Chung-Cheng Chiu, James Qin, Ruoming Pang, and Yonghui Wu. 2021. [w2v-bert: Combining contrastive learning and masked language modeling for self-supervised speech pre-training](#). In *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 244–250.
- Alexis Conneau, Min Ma, Simran Khanuja, Yu Zhang, Vera Axelrod, Siddharth Dalmia, Jason Riesa, Clara Rivera, and Ankur Bapna. 2022. [FLEURS: few-shot learning evaluation of universal representations of speech](#). In *IEEE Spoken Language Technology Workshop, SLT 2022, Doha, Qatar, January 9-12, 2023*, pages 798–805. IEEE.
- Alexis Conneau, Shijie Wu, Haoran Li, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Emerging cross-lingual structure in pretrained language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6022–6034, Online. Association for Computational Linguistics.
- Maureen de Seyssel, Marvin Lavechin, Yossi Adi, Emmanuel Dupoux, and Guillaume Wisniewski. 2022. [Probing phoneme, language and speaker information in unsupervised speech representations](#). In *23rd Annual Conference of the International Speech Communication Association, Interspeech 2022, Incheon, Korea, September 18-22, 2022*, pages 1402–1406. ISCA.

- Tu Anh Dinh, Danni Liu, and Jan Niehues. 2022. [Tackling data scarcity in speech translation using zero-shot multilingual machine translation techniques](#). In *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6222–6226.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurélien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Rozière, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny Livshits, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Graeme Nail, Grégoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel M. Kloumann, Ishan Misra, Ivan Evtimov, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, and et al. 2024. [The llama 3 herd of models](#). *CoRR*, abs/2407.21783.
- Paul-Ambroise Duquenne, Holger Schwenk, and Benoît Sagot. 2023. [SONAR: sentence-level multimodal and language-agnostic representations](#). *CoRR*, abs/2308.11466.
- Marco Gaido, Mauro Cettolo, Matteo Negri, and Marco Turchi. 2021. [CTC-based compression for direct speech translation](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 690–696, Online. Association for Computational Linguistics.
- Naman Goyal, Cynthia Gao, Vishrav Chaudhary, Peng-Jen Chen, Guillaume Wenzek, Da Ju, Sanjana Krishnan, Marc’Aurelio Ranzato, Francisco Guzmán, and Angela Fan. 2022. [The Flores-101 evaluation benchmark for low-resource and multilingual machine translation](#). *Transactions of the Association for Computational Linguistics*, 10:522–538.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [Lora: Low-rank adaptation of large language models](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Yassine El Kheir, Ahmed Ali, and Shammur Absar Chowdhury. 2024. [Speech representation analysis based on inter- and intra-model similarities](#). In *IEEE International Conference on Acoustics, Speech, and Signal Processing, ICASSP 2024 - Workshops, Seoul, Republic of Korea, April 14-19, 2024*, pages 848–852. IEEE.
- Sneha Kudugunta, Ankur Bapna, Isaac Caswell, and Orhan Firat. 2019. [Investigating multilingual NMT representations at scale](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1565–1575, Hong Kong, China. Association for Computational Linguistics.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven C. H. Hoi. 2023. [BLIP-2: bootstrapping language-image pre-training with frozen image encoders and large language models](#). In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 19730–19742. PMLR.
- Danni Liu, Jan Niehues, James Cross, Francisco Guzmán, and Xian Li. 2021. [Improving zero-shot translation by disentangling positional information](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1259–1273, Online. Association for Computational Linguistics.
- NLLB Team, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Hefernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barrault, Gabriel Mejia Gonzalez, Prangthip Hansanti, John Hoffman, Semarley Jarrett, Kaushik Ram Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, and Jeff Wang. 2022. [No language left behind: Scaling human-centered machine translation](#). *arXiv preprint*.
- nostalgebraist. 2020. [Interpreting gpt: the logit lens](#). <https://www.lesswrong.com/posts/AcKRB8wDpdaN6v6ru/interpreting-gpt-the-logit-lens>.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of*



- the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2023. [Robust speech recognition via large-scale weak supervision](#). In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 28492–28518. PMLR.
- Maithra Raghu, Justin Gilmer, Jason Yosinski, and Jascha Sohl-Dickstein. 2017. [SVCCA: singular vector canonical correlation analysis for deep learning dynamics and interpretability](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 6076–6085.
- Elizabeth Salesky, Neha Verma, Philipp Koehn, and Matt Post. 2023. [Multilingual pixel representations for translation and effective cross-lingual transfer](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13845–13861, Singapore. Association for Computational Linguistics.
- Seamless Communication, Loïc Barrault, Yu-An Chung, Mariano Coria Meglioli, David Dale, Ning Dong, Paul-Ambroise Duquenne, Hady Elsahar, Hongyu Gong, Kevin Heffernan, John Hoffman, Christopher Klaiber, Pengwei Li, Daniel Licht, Jean Maillard, Alice Rakotoarison, Kaushik Ram Sadagopan, Guillaume Wenzek, Ethan Ye, Bapi Akula, Peng-Jen Chen, Naji El Hachem, Brian Ellis, Gabriel Mejia Gonzalez, Justin Haaheim, Prangthip Hansanti, Russ Howes, Bernie Huang, Min-Jae Hwang, Hirofumi Inaguma, Somya Jain, Elahe Kalbassi, Amanda Kallet, Iliia Kulikov, Janice Lam, Daniel Li, Xutai Ma, Ruslan Mavlyutov, Benjamin Peloquin, Mohamed Ramadan, Abinesh Ramakrishnan, Anna Y. Sun, Kevin Tran, Tuan Tran, Igor Tufanov, Vish Vogeti, Carleigh Wood, Yilin Yang, Bokai Yu, Pierre Andrews, Can Balioglu, Marta R. Costa-jussà, Onur Celebi, Maha Elbayad, Cynthia Gao, Francisco Guzmán, Justine Kao, Ann Lee, Alexandre Mourachko, Juan Pino, Sravya Popuri, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, Paden Tomasello, Changhan Wang, Jeff Wang, and Skyler Wang. 2023. [SeamlessM4T-massively multilingual & multimodal machine translation](#). *CoRR*, abs/2308.11596.
- Amitay Sicherman and Yossi Adi. 2023. [Analysing discrete self supervised speech representation for spoken language modeling](#). In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5.
- Haoran Sun, Xiaohu Zhao, Yikun Lei, Shaolin Zhu, and Deyi Xiong. 2023. [Towards a deep understanding of multilingual end-to-end speech translation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 14332–14348, Singapore. Association for Computational Linguistics.
- Changli Tang, Wenyi Yu, Guangzhi Sun, Xianzhao Chen, Tian Tan, Wei Li, Lu Lu, Zejun Ma, and Chao Zhang. 2024. [SALMONN: towards generic hearing abilities for large language models](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023a. [Llama: Open and efficient foundation language models](#). *CoRR*, abs/2302.13971.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurélien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023b. [Llama 2: Open foundation and fine-tuned chat models](#). *CoRR*, abs/2307.09288.
- Ioannis Tsiamas, Gerard I. Gállego, José Fonollosa, and Marta Costa-jussà. 2024. [Pushing the limits of zero-shot end-to-end speech translation](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 14245–14267, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Laurens van der Maaten and Geoffrey Hinton. 2008. [Visualizing data using t-sne](#). *Journal of Machine Learning Research*, 9(86):2579–2605.
- Gary Wang, Kyle Kastner, Ankur Bapna, Zhehuai Chen, Andrew Rosenberg, Bhuvana Ramabhadran, and Yu Zhang. 2023. [Understanding shared speech-text representations](#). In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5.
- Chris Wendler, Veniamin Veselovsky, Giovanni Monea, and Robert West. 2024. [Do llamas work in English?](#)



on the latent language of multilingual transformers. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15366–15394, Bangkok, Thailand. Association for Computational Linguistics.

Zhaofeng Wu, Xinyan Velocity Yu, Dani Yogatama, Jiasen Lu, and Yoon Kim. 2024. [The semantic hub hypothesis: Language models share semantic representations across languages and modalities](#). *Preprint*, arXiv:2411.04986.

An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jianxin Yang, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Xuejing Liu, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, Zhifang Guo, and Zhihao Fan. 2024. [Qwen2 technical report](#). *CoRR*, abs/2407.10671.

Hongchuan Zeng, Senyu Han, Lu Chen, and Kai Yu. 2024. [Converging to a lingua franca: Evolution of linguistic regions and semantics alignment in multilingual large language models](#). *CoRR*, abs/2410.11718.

Shiyue Zhang, Vishrav Chaudhary, Naman Goyal, James Cross, Guillaume Wenzek, Mohit Bansal, and Francisco Guzman. 2022. [How robust is neural machine translation to language imbalance in multilingual tokenizer training?](#) In *Proceedings of the 15th biennial conference of the Association for Machine Translation in the Americas (Volume 1: Research Track)*, pages 97–116, Orlando, USA. Association for Machine Translation in the Americas.

Jinming Zhao, Hao Yang, Gholamreza Haffari, and Ehsan Shareghi. 2022. [M-adapter: Modality adaptation for end-to-end speech-to-text translation](#). In *23rd Annual Conference of the International Speech Communication Association, Interspeech 2022, Incheon, Korea, September 18-22, 2022*, pages 111–115. ISCA.

## A Details on SVCCA

We use the Singular Vector Canonical Correlation Analysis (SVCCA; [Raghu et al., 2017](#)) to evaluate the similarity of the extracted speech and text representations. Given two sets of representations  $X \in \mathbb{R}^{F_x \times M}$  and  $Y \in \mathbb{R}^{F_y \times N}$ , where  $M$  and  $N$  are the number of data points and  $F_x$  and  $F_y$  are the dimension of the features, SVCCA measures

their similarity. The inputs may differ in their feature dimension ( $F_x \neq F_y$ ), but the number of data points must be the same ( $M = N$ ). SVCCA first performs a singular value decomposition (SVD) on both  $X$  and  $Y$ , resulting in two sets of singular vectors and singular values. Then, Canonical Correlation Analysis (CCA) is applied on only the top  $m \leq M$  and top  $n \leq N$  singular vectors that explain  $k\%$  variance of  $X$  and  $Y$ . CCA will then find linear transformations that maximize the correlation between two vector sets, returning CCA correlation coefficients. The averaged value of all coefficients is the SVCCA similarity value  $\in [0, 1]$ , depending on how similar ( $= 1$ ) or different ( $= 0$ ) the two sets of representations are.

We use the implementation from [Raghu et al. \(2017\)](#)<sup>4</sup>. We take singular vectors that explain 90% variance of in the data. To stabilize the similarity computations, we use an epsilon of  $1e-10$ . To compare variable-length sequences, we follow prior works ([Kudugunta et al., 2019](#); [Liu et al., 2021](#); [Sun et al., 2023](#)) and meanpool over the sequence length dimension.

## B Details on Models and Hidden Representation Extraction

### B.1 SeamlessM4T

Seamless Massively Multilingual & Multimodal Machine Translation (SeamlessM4T; [Seamless Communication et al., 2023](#)) is an encoder-decoder model that supports speech-to-text, text-to-text, and speech-to-speech translation/transcription. It covers over 100 languages. Text inputs go through the text encoder and decoder, which are initialized with SeamlessM4T-NLLB ([Seamless Communication et al., 2023](#)), a multilingual text-to-text translation model supporting 200 languages. Speech inputs first pass through the mel filterbank feature extraction, where the outputs are given to the Conformer speech encoder, initialized with the speech representation learning model W2v-BERT 2.0 ([Chung et al., 2021](#)) and is followed by a length adaptor. The length adaptor of SeamlessM4T is a modified version of the M-Adaptor ([Zhao et al., 2022](#)), which downsamples the speech with a fixed factor. We focus on the encoder, as the subsequent parts do not support both text and speech in parallel.

We use `seamless-m4t-v2-large`<sup>5</sup> for the anal-

<sup>4</sup><https://github.com/google/svcca>

<sup>5</sup><https://huggingface.co/facebook/seamless-m4t-v2-large>

yses. Both speech and text encoders have 24 layers with a feature size of 1024. We analyze the speech and text representations after the layers {1, 2, 4, 6, 8, 10, 12, 14, 16, 18, 20, 22, 23, 24}, both speech and text input embeddings, and the speech representations after the length adaptor.

## B.2 SONAR

Sentence-level multimodal and language-agnostic representations (SONAR; Duquenne et al., 2023) is a multimodal and multilingual sentence embedding model for 200 languages. It has *one multilingual* text encoder initialized with NLLB (NLLB Team et al., 2022) and *multiple monolingual* speech encoders initialized with Wav2Vec2-BERT 2.0 (Chung et al., 2021). A multilingual text decoder initialized with NLLB is used in training for translation and autoencoding. The encoder outputs are used to produce sentence embeddings by pooling along the sequence length dimension (meanpooling for text and learned attention pooling is used for speech). Additionally, the mean squared error (MSE) loss is used on the encoder outputs, which encourages aligning sentences in the shared embedding space by reducing the differences between embeddings of the same semantic meaning but of different languages and modality.

We use the pre-trained SONAR models from fairseq<sup>6</sup>. Like Seamless, all encoders have 24 layers and a feature dimension of 1024. We extract representations from layers {1, 2, 4, 6, 8, 10, 12, 14, 16, 18, 20, 22, 23, 24}, the input embeddings, and the final speech and text embeddings after pooling, which are all of the same feature dimension of 1024.

## B.3 SALMONN

Speech Audio Language Music Open Neural Network (SALMONN; Tang et al., 2024) is based on Vicuna<sup>7</sup> (Chiang et al., 2023), a text-based LLM fine-tuned from Llama2 (Touvron et al., 2023b) to follow text instructions. Vicuna is finetuned with low-rank adaptation (LoRA) (Hu et al., 2022) to ingest inputs from audio features from the Whisper (Radford et al., 2023) encoder and the BEATs (Chen et al., 2023b) encoder. Window-level Q-Former (Li et al., 2023; Tang et al., 2024) is used to downsample the audio features with a window size of 0.33 second.

<sup>6</sup><https://github.com/facebookresearch/SONAR>

<sup>7</sup><https://huggingface.co/lmsys/vicuna-7b-v1.5>

Since SALMONN only accepts audio and text inputs simultaneously and the auditory and textual embeddings are given to Vicuna as one concatenated input, the extracted raw representations equal the concatenated speech and text representations. To analyze hidden speech and text representations separately, the raw representations are split into speech and text representations with the input length dimension.

We use the 7B version of SALMONN<sup>8</sup>. The decoder has 32 layers. We analyze layers {1, 2, 4, 6, 8, 10, 12, 14, 16, 18, 20, 22, 24, 26, 28, 30, 31, 32}, the speech encoder outputs before the Q-Former, and the textual and the auditory embeddings after the Q-Former. The speech encoder outputs before the Q-former has the feature size of 2048, while the text embeddings, speech encoder outputs after the Q-former and all decoder layers have a feature size of 4096. The different feature sizes cause no problem for SVCCA, as it can handle different input feature dimensions.

## C Details on Selected Languages

We use the FLEURS (Conneau et al., 2022) test split to extract hidden representations, and use the normalized transcriptions instead of raw transcriptions for the text representations. As FLEURS has multiple utterances for the same sentence with different speakers, we remove these duplicates randomly, so that the same sentence only appears once in the audio set. We choose 30 language from the 102 languages supported by FLEURS for analyses. While deciding on the languages, we maintain an even distribution of different language characteristics such as script, family and resource-level (high, medium and low as classified by Seamless Communication et al. (2023)). The statistics per language are in Table 1. We analyze the the same set of languages for SeamlessM4T and SALMONN. For SONAR, as it uses monolingual audio encoders and does not cover support six languages from this set, we replace them with other languages at the same resource-level. A fully overlapping set of languages would have provided a cleaner experimental setup. However, since our conclusions are based on comparisons of similarity scores within the same model between modalities and languages, rather than across different models, we believe that

<sup>8</sup><https://huggingface.co/tsinghua-ee/SALMONN-7B>

Code	Name	Script	Family	Resource level	Seamless & SALMONN	SONAR	# sentences original	Deduplicated
amh	Amharic	Ethiopic	Afro-Asiatic	low	✓		516	296
arb	Arabic	Arabic	Afro-Asiatic	high	✓	✓	428	283
asm	Assamese	Bengali	Indo-European	low		✓	984	349
bul	Bulgarian	Cyrillic	Indo-European	low	✓	✓	658	344
cat	Catalan	Latin	Indo-European	high	✓	✓	940	350
cmn	Chinese Mandarin	Hant	Sino-Tibetan	high	✓	✓	945	349
deu	German	Latin	Indo-European	high	✓	✓	862	347
ell	Greek	Greek	Indo-European	medium	✓		650	333
eng	English	Latin	Indo-European	high	✓	✓	647	350
est	Estonian	Latin	Uralic	medium	✓	✓	893	345
fin	Finnish	Latin	Uralic	high	✓	✓	918	348
fra	French	Latin	Indo-European	high	✓	✓	676	332
heb	Hebrew	Hebrew	Afro-Asiatic	low		✓	792	347
hin	Hindi	Devanagari	Indo-European	medium	✓	✓	418	265
hye	Armenian	Armenic	Indo-European	low	✓		932	350
ind	Indonesian	Latin	Austronesian	medium	✓	✓	687	328
ita	Italian	Latin	Indo-European	high	✓	✓	865	346
jpn	Japanese	Japanese	Japonic	high	✓	✓	650	321
kat	Georgian	Georgian	Kartvelian	low	✓		979	350
khm	Khmer	Khmer	Austroasiatic	low	✓		949	335
kor	Korean	Korean	Koreanic	medium	✓	✓	382	270
lao	Lao	Lao	Tai-Kadai	low		✓	405	260
lit	Lithuanian	Latin	Indo-European	low	✓	✓	986	349
mal	Malayalam	Malayalam	Dravidian	low		✓	985	344
mar	Marathi	Devanagari	Indo-European	low	✓	✓	1020	349
nld	Dutch	Latin	Indo-European	high	✓	✓	364	251
pes	Persian	Arabic	Indo-European	low	✓	✓	871	324
rus	Russian	Cyrillic	Indo-European	medium	✓	✓	775	344
sna	Shona	Latin	Atlantic-Congo	low	✓		925	348
snd	Sindhi	Arabic	Indo-European	low	✓	✓	980	350
swh	Swahili	Latin	Atlantic-Congo	low		✓	487	312
tam	Tamil	Tamil	Dravidian	medium	✓	✓	591	336
tel	Telugu	Telugu	Dravidian	medium		✓	472	302
tha	Thai	Thai	Tai-Kadai	medium	✓	✓	1020	349
tur	Turkish	Latin	Turkic	medium	✓	✓	743	329
yue	Cantonese	Hant	Sino-Tibetan	low	✓	✓	819	339

Table 1: List of analyzed languages. The column “# sentences original” lists the number of sentences of a language from FLEURS. The number of unique sentences in the test split for each language is given under the “Deduplicated” column.

the differing sets of languages do not compromise the validity of our findings.

For the cross-modal analysis (§3.1), each representation sets are reduced to the first 251 representations, as this is the smallest number of input data without duplicates. For the cross-lingual analysis (§3.2), each intersects were reduced to the first 194 intersecting representations for SeamlessM4T and SALMONN, and 192 for SONAR.