Code-Switching for Multilingual LLMs

Haneul Yoo[1] **Juhyun Oh**[2] **Genta Indra Winata**[3]
Ruochen Zhang[4] **Thamar Solorio**[5] **Alice Oh**[2]
[1]New York University [2]KAIST [3]Capital One [4]Microsoft [5]MBZUAI
haneul.yoo@nyu.edu 411juhyun@kaist.ac.kr genta.winata@capitalone.com
ruochenzhang@microsoft.com thamar.solorio@mbzuai.ac.ae alice.oh@kaist.edu

Abstract

This tutorial provides a structured overview of code-switching for multilingual large language models (LLMs), connecting linguistic theory, NLP methodology, and human-centered perspectives. Code-switching is a natural communicative practice among multilingual speakers and is increasingly central to LLM interactions, where users bring mixed-language practices to prompts and models themselves may produce mixed-language outputs. The tutorial will cover theoretical foundations from sociolinguistics, code-switching behavior in human-AI interaction, datasets and evaluation benchmarks, synthetic data generation, multilingual LLM training and inference, mechanistic interpretability of unintended language switching, speech and multimodal code-switching, and societal impacts. By synthesizing work across NLP, linguistics, HCI, speech processing, and multilingual communities, the tutorial aims to give participants a coherent view of the field and identify open challenges for building multilingual LLMs that better reflect how multilingual users actually communicate.

1 Introduction

Code-switching, the practice of alternating between two or more languages within a single conversation or utterance, is a natural and widespread phenomenon among multilingual speakers worldwide (Myers-Scotton, 1993; Poplack, 1980). Far from reflecting linguistic deficiency, it demonstrates sophisticated communicative competence, serving pragmatic, social, and cultural functions that monolingual communication cannot always capture. Therefore, it has been extensively studied in (socio)linguistics, with theoretical frameworks developed to explain when, why, and how speakers switch between languages (Kamwangamalu, 1989; Dewaele and Li, 2014; Halpin and Melzi, 2021).

While multilingual large language models (LLMs) have advanced rapidly, realizing their promise for the world's multilingual users requires more than expanding the languages they support. It also requires engaging with how multilingual people actually communicate. Chief among these communicative practices is code-switching, which multilingual users naturally bring to LLM interactions (Oh et al., 2026; Bawa et al., 2020; Choi et al., 2023), and which LLMs also produce in their outputs (Li et al., 2025). As such, code-switching has emerged as a critical phenomenon at the intersection of natural language processing (NLP), linguistics, and human-computer interaction (HCI) (Doğruöz et al., 2021; Winata et al., 2023a; Dedema et al., 2026), spurring growing research from dataset construction and model evaluation, to training strategies, mechanistic interpretability, speech processing, and societal impact.

Despite this rapidly expanding landscape, research on code-switching for multilingual LLMs remains scattered across communities. This tutorial bridges this gap by providing a structured overview of theoretical foundations, empirical findings, open questions, and future directions.

2 Tutorial Outline

This tutorial will focus on multilingual LLMs through the lens of code-switching. For each session, we will present a systematic overview of relevant topics, followed by an in-depth exploration of recent advancements. The tutorial will conclude with an audience-engaged discussion of open challenges and emerging research opportunities.

2.1 Background and Preliminaries [30 min]

Linguistic Theories on Code-switching. Decades of research in sociolinguistics have produced rich theoretical frameworks for understanding code-switching. Structural accounts—ranging from the Matrix Language Frame (Joshi, 1982; Myers-Scotton, 1993) to the Equivalence

Proceedings of the 5th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 15th International Joint Conference on Natural Language Processing: Tutorial Abstracts, pages 1–10
November 6-10, 2026 ©2026 Association for Computational Linguistics

Constraint (Poplack, 1980) and Functional Head Constraint (Belazi et al., 1994)—attempt to characterize the syntactic conditions under which switching occurs, while constraint-free theories argue that switching is governed solely by the grammars of the contributing languages (MacSwan, 2005). More recent approaches, such as *translanguaging*, challenge strict language boundaries, viewing multilingual practice as a unified, fluid repertoire rather than alternation between discrete systems (Wei, 2018). Sociolinguists have also identified varied motivations for code-switching—from marking social identity and in-group solidarity to navigating power asymmetries and achieving pragmatic precision (Muysken, 2020; Poplack, 1980; Heller, 1992). We will survey these frameworks and examine how, and to what extent, they have been operationalized in NLP research on multilingual LLMs.

Code-switching in Human-AI Interaction. While code-switching in human-human interaction is well-documented, its manifestation in human-AI interaction remains far less understood (Dedema et al., 2026). Existing studies on LLM usage patterns suggest that multilingual users do switch languages when interacting with AI systems (Oh et al., 2026). Yet, the motivations and dynamics differ in important ways. For example, users interact with a text-based, non-human interlocutor whose language capabilities are often unpredictable, which may shape when and how they choose to code-switch (Choi et al., 2023; Bawa et al., 2020). We will review what is currently known about these patterns and highlight the key open questions around user expectations and behavior in multilingual human-AI interaction.

2.2 Evaluation and Datasets [30 min]

Code-switching Data. We will survey the current landscape of code-switching corpora and evaluation benchmarks, focusing on how dataset scarcity shapes model assessment and data creation practices. We will review broad multilingual NLP benchmarks (Yang and Chai, 2025; Abdaljalil et al., 2025; Mohamed et al., 2026; Farhansyah et al., 2026) and task-specific evaluations for translation (Huzaifah et al., 2024; Zhang et al., 2025b), dialogue summarization (Suresh et al., 2025), red-teaming (Upadhayay and Behzadan, 2024; Song et al., 2025; Yoo et al., 2025b), and output-language confusion (Oh et al., 2026). Despite recent progress, we will emphasize that code-switching remains a challenging setting for multilingual LLM evaluation, not a solved proxy task (Zhang et al., 2023).

Synthetic Data Generation. We will then examine the data bottleneck. Most available corpora are concentrated in a small set of English-involving language pairs, such as Hindi–English (Khanuja et al., 2020; Singh et al., 2018) and Mandarin Chinese–English (Li et al., 2022; Lovenia et al., 2022; Xiang et al., 2020; Li et al., 2012; Lyu et al., 2010), while naturally occurring non-English or multilingual code-switching resources remain scarce (Winata et al., 2023a). This scarcity has made synthetic generation central to scaling evaluation and training resources. We will compare rule- or theory-based methods (Kuwanto et al., 2026; Winata et al., 2026; Yan et al., 2025), toolkit-driven generation for specific language pairs (Jayanthi et al., 2021; Rizvi et al., 2021), prompting-based generation with frontier LLMs (Yong et al., 2023; Kim et al., 2025; Yang and Chai, 2025), and multi-agent frameworks (Xie et al., 2025).

Implications of LLM-based Synthetic Data. Finally, we will discuss the methodological implications of using LLMs to generate code-switching data. While LLM-based generation improves scalability, coverage, and controllability (Winata et al., 2026), it also raises concerns about distributional mismatch with human bilingual speech, contamination between generation and evaluation, and the risk of encoding model-preferred rather than community-authentic switching patterns (Yoo et al., 2025a). We will frame these trade-offs as a central challenge for future code-switching evaluation.

2.3 Multilingual Language Modeling [30 min]

Code-Switching in LLM Training. We will examine how code-switching can be leveraged throughout the full lifecycle of multilingual LLMs, from pre-training to inference. We will first introduce recent work that incorporates code-switching into model training, including pre-training (Wang et al., 2025b) and continual pre-training (Yoo et al., 2025a) that expose models to mixed-language distributions and improve cross-lingual generalization. We will then cover downstream adaptation methods, including supervised fine-tuning (Son et al., 2026; Li et al., 2024), instruction tuning (Lee et al., 2024), and reinforcement learning (Li et al., 2025), highlighting how code-switching signals can improve multilingual capabilities across tasks.

Code-Switching in LLM Inference. Beyond training, we will discuss the role of code-switching in inference time, including in-context learning (Yoo et al., 2026) and reasoning (Wang et al., 2025a), where code-switching can serve as a mechanism to elicit latent multilingual knowledge. We will also introduce recent findings on cultural knowledge activation, where code-switching can help models access culturally grounded knowledge that may not surface in monolingual settings (Kim et al., 2025). Taken together, we will emphasize code-switching as a mechanism that interacts with multilingual representation learning and reasoning.

2.4 Mechanistic Interpretability [30 min]

Code-Switching from Interpretability Lens. We will examine code-switching in multilingual LLMs through the lens of mechanistic interpretability, focusing on recent work that identifies and controls the internal causes of unintended switching. We will discuss recent findings that language confusion emerges in final layers where models transition from shared to language specific representation (Nie et al., 2025), and complementary sparse-autoencoder analyses suggesting that unexpected code-switching can result from overactivated target-language features (Deng et al., 2026). Finally, we will highlight how these mechanistic insights enable targeted interventions, including neuron editing, feature suppression, and partial circuit tuning, which can reduce language confusion while preserving or even improving multilingual performance (Xiao et al., 2026). Taken together, code-switching is governed by identifiable internal mechanisms rather than diffuse model behavior, making it a tractable target for interpretability-driven control.

2.5 Speech and Multimodality [30 min]

Code-switching in Speech Processing. We will briefly review research on code-switching in speech processing, from early bilingual recognition to multilingual speech foundation models and speech-aware LLMs. Early work treated mixed-language speech primarily as a boundary-detection and decoding problem, focusing on language-boundary identification and bilingual recognition (Shia et al., 2004; Lyu et al., 2006). Subsequent progress was enabled by shared Mandarin–English and multilingual code-switching corpora (Lyu et al., 2010; Li et al., 2012, 2022; Lovenia et al., 2022). We will then cover the recent shift toward multilingual LLM-assisted speech models, including language alignment losses (Liu et al., 2025a), speech-conditioned LLMs with mixture-of-experts routing (Zhang et al., 2025a), and parameter-efficient adaptation for intra-sentential switching (Zhao et al., 2025; Yang et al., 2025). We will connect these methods to newer benchmark-style resources, which broaden the evaluation beyond spontaneous bilingual speech to multilingual and domain-specific settings (Ugan et al., 2024; Zhou et al., 2025; Yan et al., 2025; Xie et al., 2025). We will also discuss speech-to-speech interaction models, where recent results show that language alignment remains fragile even for large speech LLMs (Liu et al., 2025b). In addition, we will survey textless or zero-resource approaches for spoken settings where one or more languages lack standardized writing systems (Huang et al., 2024; Chen et al., 2023; Zhang et al., 2025c).

Multimodal Code-Switching. We will expand from speech processing to multimodal settings where speech is jointly modeled with text or vision. For translation, early work established code-switching speech translation as a distinct task and compared cascaded and end-to-end approaches (Weller et al., 2022). More recent systems explicitly interleave speech and text representations, with CoSTA (Shankar et al., 2025) showing that cross-modal alignment helps in low-resource code-switching speech translation. For video, we will discuss audiovisual settings, such as code-switching lip reading, highlighting the need to align language choice with visual cues rather than audio alone (Zhang et al., 2024). We will position code-switching as a genuinely multimodal phenomenon relevant to meeting video, social media clips, assistive communication, and embodied agents.

2.6 Impacts and Future Directions [30 min]

Societal Impacts. We will finally discuss the impacts of code-switching research in LLMs on the global society and the major research questions that remain. We begin with the research in the human-computer interaction field, where early studies show that multilingual users prefer chatbots that code-mix in a reciprocal way (Bawa et al., 2020; Choi et al., 2023). More specific challenges in multilingual LLM interaction have been discussed in Dedema et al. (2026); Oh et al. (2026). Users who belong in linguistic or cultural minorities are more disadvantaged by the prevalence of LLMs, and

code-switching can alleviate this problem, as we can see in domain-specific applications such as English education as a foreign language (Park et al., 2025) and health care (Harrington et al., 2022). Along with these papers, research shows that code-switching can help reduce cultural and linguistic challenges of immigrants (Lee et al., 2025).

Future Directions. We will turn this discussion into an interactive program that engages the audience as multilingual users, researchers, and practitioners. Participants will be invited to reflect on their own experiences with code-switching in LLM interactions, identify examples from their linguistic communities, discuss why such switching occurs, and surface edge cases that remain difficult for current models. Building on these observations, the tutorial will conclude by synthesizing open challenges and future research directions in code-switching for multilingual LLMs, emphasizing collaboration among NLP, linguistics, HCI, speech processing, and multilingual communities.

3 Tutorial Information

Type & Length. 3-hour, cutting-edge tutorial.

Target Audience. The tutorial will be accessible to anyone with basic knowledge of machine learning (ML) and natural language processing (NLP), and will be of interest to researchers, students, and practitioners in both academia and industry.

Prerequisites. Following knowledge is assumed:
- ML: supervised learning, transformer models
- NLP: language modeling, data and evaluation

Past Tutorials. Related tutorials include:
- EMNLP 2019 tutorial on *Processing and Understanding Mixed Language Data* (Choudhury et al., 2019) offered a comprehensive overview of the linguistic and computational aspects of code-mixing, at a time when the field lacked a coherent entry point. However, because it was developed before the rise of multilingual LLMs, it does not address recent advances and open questions surrounding LLM-based evaluation, generation, and use of synthetic code-switching data at scale.
- ACL 2023 tutorial on *Everything you need to know about Multilingual LLMs: Towards fair, performant and reliable models for languages of the world* (Sitaram et al., 2023) reviewed data, training, evaluation, responsible AI, and deployment for multilingual LLMs at a broad level, with limited space for an in-depth treatment of code-switching. Our tutorial complements these prior gaps by revisiting code-switching as a central key for multilingual LLMs.

4 Reading List

- A survey of Code-Switching: Linguistic and Social Perspectives for Language Technologies (Doğruöz et al., 2021)
- The decades progress on code-switching research in NLP: A systematic survey on trends and challenges (Winata et al., 2023b)
- Multilingual Large Language Models Are Not (Yet) Code-Switchers (Zhang et al., 2023)
- SwitchLingua: The First Large-Scale Multilingual and Multi-Ethnic Code-Switching Dataset (Xie et al., 2025)
- Investigating and Scaling up Code-Switching for Multilingual Language Model Pre-Training (Wang et al., 2025b)
- Language Mixing in Reasoning Language Models: Patterns, Impact, and Internal Causes (Wang et al., 2025a)
- Mechanistic Understanding and Mitigation of Language Confusion in English-Centric Large Language Models (Nie et al., 2025)
- CoSTA: Code-Switched Speech Translation using Aligned Speech-Text Interleaving (Shankar et al., 2025)
- Toward a Multilingual Conversational Agent: Challenges and Expectations of Code-mixing Multilingual Users (Choi et al., 2023)

5 Ethics Statement

This tutorial will present current research on the use of code-switching for training, evaluation, interpretation, and societal impact of multilingual LLMs in the field. Although we aim to promote linguistic diversity by discussing issues related to multilingual models trained in 100+ languages, many languages worldwide are yet to be supported by these models.

6 Tutorial Instructors

Haneul Yoo is an Assistant Professor/Faculty Fellow in the Department of Computer Science at the NYU Courant Institute School of Mathematics, Computing, and Data Science. She received her PhD in the School of Computing at KAIST. With a background in machine learning (ML) and natural language processing (NLP), she has worked on developing language models and methods to improve

generalization across languages and designing resources and evaluations for underrepresented languages and cultures. She has published papers in major NLP conferences, including ACL, EMNLP, and NAACL, and has been named a Rising Star at KAIST exploreCSR supported by Google. She has served as a Social Chair at NeurIPS 2026 and co-organized the Multilingual and Equitable Language Technologies (MELT) workshop at COLM 2025, the Women in Machine Learning (WiML) symposium at ICML 2026, and the Generative AI for the World (GenAI4World) workshop at COLM 2026. She delivered invited talks on code-switching and multilingual/multicultural LLMs at Princeton, UMD, Mila, and NYU, and served as a instructor for Natural Language Processing at NYU.

Juhyun Oh is a PhD student in the School of Computing at KAIST and is currently a visiting scholar at Mila. She previously studied Linguistics at Seoul National University, where she received both her BA and MA. Her research focuses on developing evaluation methodologies for language models that better reflect real-world user experiences, particularly across diverse linguistic and cultural contexts. Her work on improving support for underrepresented languages in NLP has been published at top NLP and ML conferences. She has served as a Guest Lecturer for the LLM Seminar at Georgia Tech, delivering a talk titled "Rethinking Evaluation for Language Models," and as a Teaching Assistant for the AI Tech Boostcamp at NAVER Connect Foundation. She has co-organized the Generative AI for the World (GenAI4World) workshop at COLM 2026.

Genta Indra Winata is a Senior Research Manager in LLM Core and Agentic AI team at Capital One. His research interests include language modeling, multilingual NLP, and conversational agents. He is actively involved in large-scale dataset collection and the development of multimodal and multilingual reward models and language models, with a particular focus on advancing research in low-resource languages. He has served as organizer of Sixth Workshop on Computational Approaches to Linguistic Code-Switching (CALCS) at EMNLP 2023, the Multilingual Representation Learning (MRL) at EMNLP 2023 and EMNLP 2026, and the South-East Asian Language Processing (SEALP) workshops collocated with AACL 2024 and COLING 2025, and Financial Technology and Natural Language Processing (FinNLP) at EMNLP 2025. He also delivered a tutorial on the Current Status of NLP in South East Asia with Insights from Multilingualism and Language Diversity at AACL 2023. He was a keynote speaker at SPELLL 2023, the Fields Matter Workshop at ACL 2024, and the Workshop of Language Models for Underserved Communities (LM4UC) at NAACL 2025.

Ruochen Zhang is a Senior Reseacher at Microsoft. She received her PhD from Brown University. Her research focuses on understanding the mechanisms underlying model weaknesses and brittleness, particularly through the lens of multilinguality. She has published at major NLP and ML venues, including EMNLP, ICLR, and NeurIPS. She has served as an Area Chair at LREC-COLING 2024 and as a panelist at the 7th Workshop on Computational Approaches to Linguistic Code-Switching (CALCS). She also co-organized the tutorial "Current Status of NLP in Southeast Asia: Insights from Multilingualism and Language Diversity" at AACL 2023. She is one of the initiators of the grassroots community for Southeast Asian researchers, SEACrowd, and has consistently co-organized Birds of a Feather events at major *CL conferences.

Thamar Solorio is Professor of NLP at MBZUAI, where she also serves as Vice Provost for Faculty Excellence and Advancement. She has been working on code-switching for over two decades, with seminal papers dating back to the pre-LLM era on topics such as prediction of code-switching points and enabling technology for the processing of code-switching data. She has served as General Chair of EMNLP 2024, Program Chair for NAACL 2019, and has served in numerous other organization roles including D&I co-chair, demonstration co-chair, and faculty advisor to the student research workshop for different *CL conferences. She has also led, and collaborated in the organization of multiple shared tasks on different topics, including code-switching. Thamar has delivered keynote talks at *CL and/or associated events such as NAACL 2021, CoNLL 2024, and WNUT on topics related to code-switching, and artificial social intelligence. Through her career she has prepared and delivered multiple courses at the undergraduate and graduate level.

Alice Oh is a KAIST ICT Endowed Professor in the School of Computing at KAIST. She received her PhD in Computer Science from MIT in 2008.

Her major research area is at the intersection of natural language processing (NLP) and computational social science (CSS), with a recent focus on multilingual and multicultural aspects of LLMs. She has served as Tutorial Chair for NeurIPS 2019, and DEI Chair for ICLR 2019 and COLM 2024, Program Chair for ICLR 2021 and NeurIPS 2022, General Chair for ACM FAccT 2022, NeurIPS 2023, and Festival of Learning 2026 which co-locates AIED, EDM, and ACM L@S. She presented on the topics of code-switching and multilingual/multicultural LLMs as an invited speaker at Stanford, Cornell Tech, CMU, and UCLA in 2025. She also gave invited keynotes at various international conferences including ACML, ICWSM, and RSS.

References

Samir Abdaljalil, Erchin Serpedin, Khalid Qaraqe, and Hasan Kurban. 2025. Evaluating multilingual and code-switched alignment in LLMs via synthetic natural language inference. *arXiv preprint arXiv:2508.14735*.

Anshul Bawa, Pranav Khadpe, Pratik Joshi, Kalika Bali, and Monojit Choudhury. 2020. Do multilingual users prefer chat-bots that code-mix? let's nudge and find out! *Proc. ACM Hum.-Comput. Interact.*, 4(CSCW1).

Hedi M. Belazi, Edward J. Rubin, and Almeida Jacqueline Toribio. 1994. Code switching and X-bar theory: The functional head constraint. *Linguistic Inquiry*, 25(2):221–237.

Peng-Jen Chen, Kevin Tran, Yilin Yang, Jingfei Du, Justine Kao, Yu-An Chung, Paden Tomasello, Paul-Ambroise Duquenne, Holger Schwenk, Hongyu Gong, Hirofumi Inaguma, Sravya Popuri, Changhan Wang, Juan Pino, Wei-Ning Hsu, and Ann Lee. 2023. Speech-to-speech translation for a real-world unwritten language. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 4969–4983, Toronto, Canada. Association for Computational Linguistics.

Yunjae J. Choi, Minha Lee, and Sangsu Lee. 2023. Toward a multilingual conversational agent: Challenges and expectations of code-mixing multilingual users. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, New York, NY, USA. Association for Computing Machinery.

Monojit Choudhury, Anirudh Srinivasan, and Sandipan Dandapat. 2019. Processing and understanding mixed language data. 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): Tutorial Abstracts*, Hong Kong, China. Association for Computational Linguistics.

Meredith Dedema, Chiao Xuen Goh, and Pengyi Zhang. 2026. When generative AI mixes languages: Multilingual users' code-switching behavior in human-LLM interaction. In *Proceedings of the Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems*, CHI EA '26, New York, NY, USA. Association for Computing Machinery.

Boyi Deng, Yu Wan, Baosong Yang, Fei Huang, Wenjie Wang, and Fuli Feng. 2026. SASFT: Sparse autoencoder-guided supervised finetuning to mitigate unexpected code-switching in LLMs. In *The Fourteenth International Conference on Learning Representations*.

Jean-Marc Dewaele and Wei Li. 2014. Intra- and inter-individual variation in self-reported code-switching patterns of adult multilinguals. *International Journal of Multilingualism*, 11(2):225–246.

A. Seza Doğruöz, Sunayana Sitaram, Barbara E. Bullock, and Almeida Jacqueline Toribio. 2021. A survey of code-switching: Linguistic and social perspectives for language technologies. 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 1654–1666, Online. Association for Computational Linguistics.

Mohammad Rifqi Farhansyah, Hanif Muhammad Zhafran, Farid Adilazuarda, Shamsuddeen Hassan Muhammad, Maryam Ibrahim Mukhtar, Nedjma Ousidhoum, Genta Indra Winata, Ayu Purwarianti, and Alham Fikri Aji. 2026. PingPong: A natural benchmark for multi-turn code-switching dialogues. *arXiv preprint arXiv:2601.17277*.

Emily Halpin and Gigliana Melzi. 2021. Code-switching in the narratives of dual-language latino preschoolers. *International Journal of Bilingual Education and Bilingualism*, 24(9):1271–1287.

Christina N. Harrington, Radhika Garg, Amanda Woodward, and Dimitri Williams. 2022. "it's kind of like code-switching": Black older adults' experiences with a voice assistant for health information seeking. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI '22, New York, NY, USA. Association for Computing Machinery.

Monica Heller. 1992. The politics of codeswitching and language choice. *Journal of Multilingual and Multicultural Development*, 13(1-2):123–142.

Kuan-Po Huang, Chih-Kai Yang, Yu-Kuan Fu, Ewan Dunbar, and Hung-Yi Lee. 2024. Zero resource code-switched speech benchmark using speech utterance pairs for multiple spoken languages. In *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 10006–10010.

Muhammad Huzaifah, Weihua Zheng, Nattapol Chanpaisit, and Kui Wu. 2024. Evaluating code-switching translation with large language models. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 6381–6394, Torino, Italia. ELRA and ICCL.

Sai Muralidhar Jayanthi, Kavya Nerella, Khyathi Raghavi Chandu, and Alan W Black. 2021. CodemixedNLP: An extensible and open NLP toolkit for code-mixing. In *Proceedings of the Fifth Workshop on Computational Approaches to Linguistic Code-Switching*, pages 113–118, Online. Association for Computational Linguistics.

Aravind K. Joshi. 1982. Processing of sentences with intra-sentential code-switching. In *Coling 1982: Proceedings of the Ninth International Conference on Computational Linguistics*.

Nkonko Mudipanu Kamwangamalu. 1989. Code-mixing and modernization. *World Englishes*, 8(3):321–332.

Simran Khanuja, Sandipan Dandapat, Sunayana Sitaram, and Monojit Choudhury. 2020. A new dataset for natural language inference from code-mixed conversations. In *Proceedings of the 4th Workshop on Computational Approaches to Code Switching*, pages 9–16, Marseille, France. European Language Resources Association.

Seoyeon Kim, Huiseo Kim, Chanjun Park, Jinyoung Yeo, and Dongha Lee. 2025. Can code-switched texts activate a knowledge switch in LLMs? a case study on English-Korean code-switching. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 22303–22327, Suzhou, China. Association for Computational Linguistics.

Garry Kuwanto, Chaitanya Agarwal, Genta Indra Winata, and Derry Tanti Wijaya. 2026. Linguistics theory meets LLM: Code-switched text generation via equivalence constrained large language models. In *Proceedings of the 1st Workshop on Computational Developmental Linguistics (CDL)*, pages 1–14, Grand Hyatt Manchester San Diego, 1 Market Pl, San Diego, CA 92101. Association for Computational Linguistics.

Jaeseong Lee, YeonJoon Jung, and Seung-won Hwang. 2024. COMMIT: Code-mixing English-centric large language model for multilingual instruction tuning. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 3130–3137, Mexico City, Mexico. Association for Computational Linguistics.

Sunok Lee, Dasom Choi, Lucy Truong, Nitin Sawhney, and Henna Paakki. 2025. Into the unknown: Leveraging conversational AI in supporting young migrants' journeys towards cultural adaptation. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, CHI '25, New York, NY, USA. Association for Computing Machinery.

Chengfei Li, Shuhao Deng, Yaoping Wang, Guangjing Wang, Yaguang Gong, Changbin Chen, and Jinfeng Bai. 2022. TALCS: An open-source Mandarin-English code-switching corpus and a speech recognition baseline. In *Proc. Interspeech 2022*, pages 1741–1745.

Yihao Li, Jiayi Xin, Miranda Muqing Miao, Qi Long, and Lyle Ungar. 2025. The impact of language mixing on bilingual LLM reasoning. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 32531–32548, Suzhou, China. Association for Computational Linguistics.

Ying Li, Yue Yu, and Pascale Fung. 2012. A Mandarin-English code-switching corpus. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 2515–2519, Istanbul, Turkey. European Language Resources Association (ELRA).

Zhuoran Li, Chunming Hu, Junfan Chen, Zhijun Chen, Xiaohui Guo, and Richong Zhang. 2024. Improving zero-shot cross-lingual transfer via progressive code-switching. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24*, pages 6388–6396. International Joint Conferences on Artificial Intelligence Organization. Main Track.

Hexin Liu, Xiangyu Zhang, Haoyang Zhang, Leibny Paola Garcia-Perera, Andy W. H. Khong, Eng Siong Chng, and Shinji Watanabe. 2025a. Aligning speech to languages to enhance code-switching speech recognition. *IEEE Transactions on Audio, Speech and Language Processing*, 33:4712–4725.

Heyang Liu, Yuhao Wang, Ziyang Cheng, Ronghua Wu, Qunshan Gu, Yanfeng Wang, and Yu Wang. 2025b. Cs3-bench: Evaluating and enhancing speech-to-speech LLMs for Mandarin–English code-switching. *arXiv preprint arXiv:2510.07881*.

Holy Lovenia, Samuel Cahyawijaya, Genta Indra Winata, Peng Xu, Yan Xu, Zihan Liu, Rita Frieske, Tiezheng Yu, Wenliang Dai, Elham J. Barezi, Qifeng Chen, Xiaojuan Ma, Bertram E. Shi, and Pascale Fung. 2022. ASCEND: A spontaneous Chinese-English dataset for code-switching in multi-turn conversation. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 7259–7268, Marseille, France. European Language Resources Association.

Dau-cheng Lyu, Ren-yuan Lyu, Yuang-chin Chiang, and Chun-nan Hsu. 2006. Speech recognition on code-switching among the chinese dialects. In *2006 IEEE International Conference on Acoustics Speech and Signal Processing Proceedings*, volume 1, pages I–I.

Dau-Cheng Lyu, Tien-Ping Tan, Eng Siong Chng, and Haizhou Li. 2010. SEAME: a Mandarin-English code-switching speech corpus in South-east Asia. In *Proc. Interspeech 2010*, pages 1986–1989.

Jeff MacSwan. 2005. Codeswitching and generative grammar: A critique of the MLF model and some remarks on "modified minimalism". *Bilingualism: language and cognition*, 8(1):1–22.

Amr Mohamed, Yang Zhang, Michalis Vazirgiannis, and Guokan Shang. 2026. Lost in the mix: Evaluating LLM understanding of code-switched text. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 44922–44938, San Diego, California, United States. Association for Computational Linguistics.

Pieter Muysken. 2020. Code-switching and grammatical theory. In *The bilingualism reader*, pages 280–297. Routledge.

Carol Myers-Scotton. 1993. *Duelling Languages: Grammatical Structure in Codeswitching*. Oxford University Press.

Ercong Nie, Helmut Schmid, and Hinrich Schuetze. 2025. Mechanistic understanding and mitigation of language confusion in English-centric large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 690–706, Suzhou, China. Association for Computational Linguistics.

Juhyun Oh, Haneul Yoo, Faiz Ghifari Haznitrama, and Alice Oh. 2026. OLA: Output language alignment in code-switched LLM interactions. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 46603–46624, San Diego, California, United States. Association for Computational Linguistics.

Junyeong Park, Jieun Han, Yeon Su Park, Youngbin Lee, Suin Kim, Juho Kim, Alice Oh, and So-Yeon Ahn. 2025. Understanding EFL learners' code-switching and teachers' pedagogical approaches in LLM-supported speaking practice. *arXiv preprint arXiv:2512.23136*.

Shana Poplack. 1980. Sometimes I'll start a sentence in Spanish y termino en español: toward a typology of code-switching. *Linguistics*, 18(7-8):581–618.

Mohd Sanad Zaki Rizvi, Anirudh Srinivasan, Tanuja Ganu, Monojit Choudhury, and Sunayana Sitaram. 2021. GCM: A toolkit for generating synthetic code-mixed text. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 205–211, Online. Association for Computational Linguistics.

Bhavani Shankar, Preethi Jyothi, and Pushpak Bhattacharyya. 2025. CoSTA: Code-switched speech translation using aligned speech-text interleaving. In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 9194–9208, Abu Dhabi, UAE. Association for Computational Linguistics.

Chi-Jiun Shia, Yu-Hsien Chiu, Jia-Hsin Hsieh, and Chung-Hsien Wu. 2004. Language boundary detection and identification of mixed-language speech based on map estimation. In *2004 IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 1, pages I–381.

Kushagra Singh, Indira Sen, and Ponnurangam Kumaraguru. 2018. A Twitter corpus for Hindi-English code mixed POS tagging. In *Proceedings of the Sixth International Workshop on Natural Language Processing for Social Media*, pages 12–17, Melbourne, Australia. Association for Computational Linguistics.

Sunayana Sitaram, Monojit Choudhury, Barun Patra, Vishrav Chaudhary, Kabir Ahuja, and Kalika Bali. 2023. Everything you need to know about multilingual LLMs: Towards fair, performant and reliable models for languages of the world. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 6: Tutorial Abstracts)*, pages 21–26, Toronto, Canada. Association for Computational Linguistics.

Guijin Son, Donghun Yang, Hitesh Laxmichand Patel, Amit Agarwal, Hyunwoo Ko, Chanuk Lim, Srikant Panda, Minhyuk Kim, Nikunj Drolia, Dasol Choi, Kyong-Ha Lee, and Youngjae Yu. 2026. Pushing on multilingual reasoning models with language-mixed chain-of-thought. In *The Fourteenth International Conference on Learning Representations*.

Jiayang Song, Yuheng Huang, Zhehua Zhou, and Lei Ma. 2025. Multilingual blending: Large language model safety alignment evaluation with language mixture. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 3433–3449, Albuquerque, New Mexico. Association for Computational Linguistics.

Sathya Krishnan Suresh, Tanmay Surana, Lim Zhi Hao, and Eng Siong Chng. 2025. CS-sum: A benchmark for code-switching dialogue summarization and the limits of large language models. In *Proceedings of The 5th New Frontiers in Summarization Workshop*, pages 31–47, Hybrid. Association for Computational Linguistics.

Enes Yavuz Ugan, Ngoc-Quan Pham, and Alexander Waibel. 2024. DECM: Evaluating bilingual ASR performance on a code-switching/mixing benchmark. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 4468–4475, Torino, Italia. ELRA and ICCL.

Bibek Upadhayay and Vahid Behzadan. 2024. Sandwich attack: Multi-language mixture adaptive attack on LLMs. In *Proceedings of the 4th Workshop on Trustworthy Natural Language Processing (TrustNLP 2024)*, pages 208–226, Mexico City, Mexico. Association for Computational Linguistics.

Mingyang Wang, Lukas Lange, Heike Adel, Yunpu Ma, Jannik Strötgen, and Hinrich Schuetze. 2025a. Lan-

guage mixing in reasoning language models: Patterns, impact, and internal causes. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 2637–2665, Suzhou, China. Association for Computational Linguistics.

Zhijun Wang, Jiahuan Li, Hao Zhou, Rongxiang Weng, Jingang Wang, Xin Huang, Xue Han, Junlan Feng, Chao Deng, and Shujian Huang. 2025b. Investigating and scaling up code-switching for multilingual language model pre-training. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 11032–11046, Vienna, Austria. Association for Computational Linguistics.

Li Wei. 2018. Translanguaging as a practical theory of language. *Applied Linguistics*, 39(1):9–30.

Orion Weller, Matthias Sperber, Telmo Pires, Hendra Setiawan, Christian Gollan, Dominic Telaar, and Matthias Paulik. 2022. End-to-end speech translation for code switched speech. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1435–1448, Dublin, Ireland. Association for Computational Linguistics.

Genta Winata, Alham Fikri Aji, Zheng Xin Yong, and Thamar Solorio. 2023a. The decades progress on code-switching research in NLP: A systematic survey on trends and challenges. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 2936–2978, Toronto, Canada. Association for Computational Linguistics.

Genta Indra Winata, Alham Fikri Aji, Zheng-Xin Yong, and Thamar Solorio. 2023b. The decades progress on code-switching research in nlp: A systematic survey on trends and challenges. *Findings of the Association for Computational Linguistics: ACL 2023*, pages 2936–2978.

Genta Indra Winata, David Anugraha, Patrick Amadeus Irawan, Anirban Das, Haneul Yoo, Paresh Dashore, Shreyas Kulkarni, Ruochen Zhang, Haruki Sakajo, Frederikus Hudi, Anaelia Ovalle, Syrielle Montariol, Felix Gaschi, Michael Anugraha, Rutuj Ravindra Puranik, Zawad Hayat Ahmed, Adril Putra Merin, and Emmanuele Chersoni. 2026. Can large language models understand, reason about, and generate code-switched text? *arXiv preprint arXiv:2601.07153*.

Rong Xiang, Mingyu Wan, Qi Su, Chu-Ren Huang, and Qin Lu. 2020. Sina Mandarin alphabetical words:a web-driven code-mixing lexical resource. In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 833–842, Suzhou, China. Association for Computational Linguistics.

Yuxin Xiao, Zhen Huang, Wenxiao Wang, Binbin Lin, Xiaofei He, Xu Shen, and Jieping Ye. 2026. How do language models speak languages? a case study on unintended code-switching.

Peng Xie, Xingyuan Liu, Yequan Bie, Tsz Wai Chan, Yangqiu Song, Yang Wang, Hao Chen, and Kani Chen. 2025. SwitchLingua: The first large-scale multilingual and multi-ethnic code-switching dataset. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

Brian Yan, Injy Hamed, Shuichiro Shimizu, Vasista Sai Lodagala, William Chen, Olga Iakovenko, Bashar Talafha, Amir Hussein, Alexander Polok, Kalvin Chang, Dominik Klement, Sara Althubaiti, Puyuan Peng, Matthew Wiesner, Thamar Solorio, Ahmed Ali, Sanjeev Khudanpur, and Shinji Watanabe. 2025. CS-FLEURS: A massively multilingual and code-switched speech dataset. In *Interspeech 2025*, pages 743–747.

Hongli Yang, Yizhou Peng, Hao Huang, and Sheng Li. 2025. Adapting Whisper for parameter-efficient code-switching speech recognition via soft prompt tuning. In *Interspeech 2025*, pages 5203–5207.

Yilun Yang and Yekun Chai. 2025. CodeMixBench: Evaluating code-mixing capabilities of LLMs across 18 languages. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 2139–2169, Suzhou, China. Association for Computational Linguistics.

Zheng-Xin Yong, Ruochen Zhang, Jessica Zosa Forde, Skyler Wang, Arjun Subramonian, Holy Lovenia, Samuel Cahyawijaya, Genta Indra Winata, Lintang Sutawika, Jan Christian Blaise Cruz, Yin Lin Tan, Long Phan, Rowena Garcia, Thamar Solorio, and Alham Fikri Aji. 2023. Prompting multilingual large language models to generate code-mixed texts: The case of south East Asian languages. In *Proceedings of the 6th Workshop on Computational Approaches to Linguistic Code-Switching*, pages 43–63, Singapore. Association for Computational Linguistics.

Haneul Yoo, Jiho Jin, Kyunghyun Cho, and Alice Oh. 2026. Gradual code-switching as inference-time cross-lingual representational alignment for LLMs. In *Third Conference on Language Modeling*.

Haneul Yoo, Cheonbok Park, Sangdoo Yun, Alice Oh, and Hwaran Lee. 2025a. Code-switching curriculum learning for multilingual transfer in LLMs. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 7816–7836, Vienna, Austria. Association for Computational Linguistics.

Haneul Yoo, Yongjin Yang, and Hwaran Lee. 2025b. Code-switching red-teaming: LLM evaluation for safety and multilingual understanding. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13392–13413, Vienna, Austria. Association for Computational Linguistics.

Fengrun Zhang, Wang Geng, Hukai Huang, Yahui Shan, Cheng Yi, and He Qu. 2025a. Boosting code-switching ASR with mixture of experts enhanced

speech-conditioned LLM. In *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5.

Ruochen Zhang, Samuel Cahyawijaya, Jan Christian Blaise Cruz, Genta Winata, and Alham Fikri Aji. 2023. Multilingual large language models are not (yet) code-switchers. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12567–12582, Singapore. Association for Computational Linguistics.

Wenbo Zhang, Aditya Majumdar, and Amulya Yadav. 2025b. CHAI for LLMs: Improving code-mixed translation in large language models through reinforcement learning with AI feedback. *arXiv preprint arXiv:2411.09073*.

Xueyi Zhang, Chengwei Zhang, Mingrui Lao, Peng Zhao, Jun Tang, Yanming Guo, Siqi Cai, Xianghu Yue, and Haizhou Li. 2024. Language without borders: A dataset and benchmark for code-switching lip reading. In *Advances in Neural Information Processing Systems*, volume 37, pages 30727–30739. Curran Associates, Inc.

Yuhao Zhang, Xiangnan Ma, Kaiqi Kou, Peizhuo Liu, Weiqiao Shan, Benyou Wang, Tong Xiao, Yuxin Huang, Zhengtao Yu, and JingBo Zhu. 2025c. Leveraging unit language guidance to advance speech modeling in textless speech-to-speech translation. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 1448–1460, Vienna, Austria. Association for Computational Linguistics.

Jiahui Zhao, Hao Shi, Chenrui Cui, Tianrui Wang, Hexin Liu, Zhaoheng Ni, Lingxuan Ye, and Longbiao Wang. 2025. Adapting Whisper for code-switching through encoding refining and language-aware decoding. In *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5.

Jiaming Zhou, Yujie Guo, Shiwan Zhao, Haoqin Sun, Hui Wang, Jiabei He, Aobo Kong, Shiyao Wang, Xi Yang, Yequan Wang, Yonghua Lin, and Yong Qin. 2025. CS-Dialogue: A 104-hour dataset of spontaneous Mandarin–English code-switching dialogues for speech recognition. *arXiv preprint arXiv:2502.18913*.