

The Tonogenesis Continuum in Tibetan: A Computational Investigation

Siyu Liang and Zhaxi Zerong

University of Washington

liangsy, tashi0@uw.edu

Abstract

Tonogenesis—the historical process by which segmental contrasts evolve into lexical tone—has traditionally been studied through comparative reconstruction and acoustic phonetics. We introduce a computational approach that quantifies the functional role of pitch at different stages of this sound change by measuring how pitch manipulation affects automatic speech recognition (ASR) performance. Through analysis on the sensitivity to pitch-flattening from a set of closely related Tibetan languages, we find evidence of a tonogenesis continuum: atonal Amdo dialects tolerate pitch removal the most, while fully tonal Ü-Tsang varieties show severe degradation, and intermediate Kham dialects fall measurably between these extremes. These gradient effects demonstrate how ASR models implicitly learn the shifting functional load of pitch as languages transition from consonant-based to tone-based lexical contrasts. Our findings show that computational methods can capture fine-grained stages of sound change and suggest that traditional functional load metrics, based solely on minimal pairs, may overestimate pitch dependence in transitional systems where segmental and suprasegmental cues remain phonetically intertwined.

1 Introduction

Tonogenesis refers to the emergence of lexical tone from earlier consonantal or laryngeal contrasts such as voicing and aspiration (Hombert, 1977; Haudricourt, 1954; Sæbø et al., 2025). Over time, secondary pitch perturbations can become the primary cue for distinguishing words, turning a formerly atonal language into a tonal one. For instance, in Lhasa Tibetan, the Classical Tibetan syllables *ka* (command) and *ga* (joy) were historically distinguished by aspiration and voicing; today, both are pronounced with an aspirated onset [k^h], but a high vs. low tone preserves the lexical contrast. This

process unfolds along a continuum: within a language family or group, some varieties maintain complex onset clusters and show little reliance on pitch (e.g. varieties of Amdo Tibetan), whereas others develop robust tone contrasts (e.g., Lhasa Tibetan) (Sun, 2015). Fully tonal languages (at later stages of tonogenesis) exhibit a strong dependence on fundamental frequency (f_0) for lexical contrasts, which can pose significant challenges for automatic speech recognition (ASR) if pitch information is removed (Zhang and Kirby, 2020; Fu et al., 1998).

Recent work has established that pitch-flattening—systematically removing f_0 contours while preserving spectral information—provides a computational method for quantifying tonal dependence in ASR systems (Liang and Levow, 2025). Across typologically diverse languages, fully tonal systems suffer dramatic performance degradation when pitch is removed, while non-tonal languages show minimal impact. Building on this validated methodology, we extend the investigation to multiple Tibetan languages representing distinct stages of tonogenesis. We hypothesize that flattening f_0 contours will cause a larger performance drop for languages with established tonal contrasts than for those that rely more on consonantal cues. By quantifying the impact of pitch loss on ASR using word error rate and character error rate, we offer computational evidence that captures the acoustic-level interplay between pitch and segmental cues—information that text-based functional load measures, which rely solely on minimal pair counts, cannot access. Our findings demonstrate that dialects at different points in the tonogenesis continuum exhibit correspondingly distinct degrees of pitch dependence, providing new empirical support for understanding how languages transition from consonant-based to tone-based lexical systems.

2 Literature Review

2.1 Tone and tonogenesis

Tonal languages use pitch (fundamental frequency, f_0) as a primary cue for lexical contrasts—that is, to distinguish words that would otherwise be identical. For example, in Cantonese, the syllable [si] with a high-level tone means ‘poem,’ whereas the same segments with a low-falling tone mean ‘time’ (Yip, 2002; Matthews and Yip, 2013). In contrast, non-tonal languages like English employ pitch mainly for intonation and stress (Edwards and Beckman, 1988). Crucially, tonogenesis can arise when older segmental cues (e.g., voicing or coda consonants) induce pitch perturbations that eventually become the main phonemic signal for lexical differentiation (Haudricourt, 1954; Hombert et al., 1979). In Vietnamese, for instance, the loss of final stops led to newly phonologized tones (Thurgood, 2002), while other Southeast Asian languages display partial or incomplete shifts (Matisoff, 1973).

2.2 Tibetan languages

More than fifty distinct varieties have been identified as part of the Tibetic branch of the Sino-Tibetan language family (Tournadre, 2014). Among these, the three most widely discussed groups are Central, Amdo, and Khams, illustrated in Figure 1 (Gesang and Gesang, 2002). The long-standing debate over whether these varieties constitute separate languages or dialects remains relevant (Haugen, 1966). From a strictly linguistic standpoint, many varieties are only marginally mutually intelligible (Driem, 2001), suggesting that they could be treated as separate languages. Nevertheless, socio-political and cultural factors often motivate their classification as “Tibetan dialects,” highlighting a shared literary heritage (based on Classical Tibetan) and a unified Tibetan ethnic identity (Tournadre, 2014; Gesang and Gesang, 2002). In this work, we use the term *Tibetan languages*, while acknowledging the nuanced nature of the debate.

Tibetan languages exhibit a diverse range of tonogenesis outcomes. Central (Ü-Tsang) varieties such as Lhasa Tibetan have developed robust tonal distinctions from historical voicing contrasts: modern Lhasa has two phonemic tones that yield four different pitch contours depending on syllable structure (DeLancey, 2017; Lim, 2018). Khams varieties occupy an *intermediate* position, relying partially on pitch while retaining residual segmen-

tal cues such as voicing contrasts or breathiness (Sun, 2015; Suzuki, 2011). For instance, Dege Khams preserves historically voiced consonants alongside four emerging tonal categories (Gesang and Gesang, 2002). At the opposite end of the spectrum, Amdo varieties remain primarily atonal, retaining a rich inventory of consonantal clusters in lieu of pitch-based contrasts (Gesang and Gesang, 2002; Sun, 1986).




Figure 1: Approximate distribution of major Tibetan dialect groups. Ü-Tsang (red), Khams (green), and Amdo (purple) show distinct historical trajectories in tonogenesis.

2.3 Tone and ASR

Although pitch variations occur in all spoken languages for prosodic or paralinguistic reasons, *tone languages* uniquely depend on fundamental frequency (f_0) as a primary lexical contrast. As a result, failing to model pitch cues in automatic speech recognition (ASR) can lead to frequent confusions among words that differ only in tone (Fu et al., 1998; Patel et al., 2010; Wang et al., 2013; Zhang and Kirby, 2020). In practice, researchers typically address tone modeling in two broad ways. Some approaches add time-aligned pitch or voice-quality features directly into the acoustic model (Lei et al., 2006; Li et al., 2011), allowing it to learn the correlations between pitch contours and word identities. Other approaches explicitly encode tonal information in the transcript by assigning diacritics or numeric labels to otherwise identical syllables or phonemes (Yuan et al., 2021; Coto-Solano, 2021). Regardless of the chosen method, extensive evidence shows that removing or flattening pitch severely degrades ASR performance in tonal languages, mirroring the disruption that human listeners experience with f_0 suppression (Fu et al., 1998; Niu et al., 2019). In other words, in fully (or

partly) tonal systems, capturing pitch movements becomes essential for correctly mapping acoustic signals to lexical items.

Recent work has demonstrated that systematically removing pitch through flattening provides a computational method for quantifying how heavily languages rely on f_0 for lexical disambiguation: fully tonal systems suffer severe ASR degradation when pitch is removed, while non-tonal languages show minimal impact (Liang and Levow, 2025). This methodology offers a way to assess the functional load of tone in ASR systems, complementing traditional linguistic measures based on minimal pair counts.

2.4 Functional Load of Tone

In phonetics and phonology, the *functional load* of a contrast refers to how crucial that contrast is for distinguishing words (Surendran and Levow, 2004). In a strongly tonal language such as Mandarin or Cantonese, this notion conceptually relates to the amount of minimal pairs that depend solely on tone or to gauging how much pitch contributes to speech intelligibility. High functional load corresponds to greater confusion when f_0 is removed or altered, which in turn leads to marked increases in ASR error rates.

However, this measure can be inadequate for tonogenesis languages, which are still shifting from consonant-based distinctions (e.g. voicing or glottal stop codas) to primarily pitch-based contrasts (Hombert et al., 1979; Ratliff, 2015). In many transitional systems, tone can remain partially predictable from segmental cues, coarticulation, or laryngeal features (Hu, 2012; Sun, 2015). For example, pitch contours might be strongly influenced by the presence of voiced versus aspirated onsets (Hombert, 1977), or by residual glottalization that masks or enhances emerging pitch distinctions (Kirby et al., 2022). These extra cues mean that a simple tally of tone-based minimal pairs overestimates how much f_0 alone is responsible for contrasts. Even after f_0 becomes phonemic, features like breathiness or preaspiration sometimes continue to overlap with pitch distinctions, complicating the idea of an independent “tone functional load.” Consequently, although functional load remains a productive concept for mature tone languages, a single pitch-based metric does not fully capture the nuanced interplay between pitch and segmental cues in incipient or partial tone systems. In these contexts, more holistic approaches are

needed to account for how tone interacts with residual voicing, glottal stops, and other acoustic features—both for phonological description and for building robust ASR models.

3 Methods

We employ a pitch-flattening methodology to investigate the tonogenesis continuum in Tibetan varieties. This method allows us to quantify how heavily each dialect relies on pitch for lexical contrasts by measuring ASR performance degradation when f_0 information is systematically removed.

3.1 Data

We examine six Tibetan varieties drawn from TIBMD@MUC (Zhao et al., 2020), described in Table 1. Recordings featured both formal and colloquial read speech, recorded at 16 kHz, transcribed in Tibetan script. In each variety, we sampled approximately two hours of audio for training. We attempted to limit training data to two speakers per variety; in Changdu and Dege, however, three speakers were used due to data constraints. From each variety, we also reserved around 30 minutes of speech to form a test set. Data from Yushu were excluded due to insufficient data.

Since Tibetan orthography relies on stacked letters and diacritics, we converted the script to Wylie transliterations via pyewts (Wylie, 1959). We then segmented each Wylie transcription at the character level (treating spaces as distinct symbols) to form a language-specific vocabulary. Audio files were resampled at 16 kHz, and we fine-tuned separate XLS-R 300m models (Babu et al., 2021), a multilingual self-supervised speech model pretrained on 128 languages, using standard hyperparameters for low-resource ASR.

3.2 Model Training and Evaluation

We fine-tuned individual XLS-R 300m models for each of the six Tibetan varieties. During pre-processing, we normalized text and generated a character-based vocabulary for CTC training. Following model convergence, we evaluated each model on both the original test set and a pitch-flattened version of the same data.

Pitch flattening was performed using Praat’s Pitch-Synchronous OverLap and Add (PSOLA) algorithm, which replaces each utterance’s natural f_0 contour with its mean pitch. This manipulation effectively removes lexical tone cues while preserving the spectral envelope and temporal structure

Group	Duration (hrs)	Spkr	Utterances
Amdo			
Xiahe	4.12	2	3549
Aba	8.16	2	6546
Qinghai	13.65	19	12859
Kham			
Changdu	2.79	7	2558
Dege	2.31	3	1245
Yushu	0.77	3	631
Ü-Tsang			
Lhasa	37.38	48	30349
Shigatse	15.15	4	10729
Total	84.33	88	68466

Table 1: Statistics of the TIBMD@MUC database, illustrating eight Tibetan languages grouped into Amdo, Kham, and Ü-Tsang.

of the speech signal. By comparing character error rate (CER) and word error rate (WER) before and after flattening, we obtain a measure of how much each variety’s ASR system relies on pitch for lexical disambiguation.

Because Ü-Tsang varieties (e.g., Lhasa) are further along the tonogenesis trajectory, we hypothesized they would suffer greater performance drops upon pitch removal, reflecting a heavier functional load for f_0 than in the more atonal Amdo dialects. Kham varieties, described linguistically as occupying an intermediate position, were expected to show moderate sensitivity to pitch manipulation.

4 Results

Table 2 presents ASR outcomes for six Tibetan languages under original vs. pitch-flattened conditions, revealing clear differences in how each variety responds to the removal of f_0 . Amdo dialects (Xiahe, Aba), generally described as *atonal* or pre-tonogenetic (Sun, 1986; Gesang and Gesang, 2002), exhibit only moderate increases in error rates after pitch flattening, suggesting that segmental contrasts (e.g., voiced vs. voiceless onsets) continue to carry the primary lexical burden. By contrast, the Ü-Tsang group (Lhasa, Shigatse), which has been characterized as fully tonal (DeLancey, 2017), incurs notably larger performance drops. Shigatse in particular shows a sharp spike in WER (+13.9%), indicating that obscuring pitch removes a principal mechanism of word-level contrast, much as in other mature tone systems (Lim, 2018; Kirby, 2014).

Falling between these two poles are the Kham varieties (Changdu, Dege), which illustrate a partially developed reliance on pitch. Changdu undergoes a moderate rise in CER and WER when pitch is flattened, implying that f_0 is emerging as a contrastive feature yet does not fully supersede residual segmental or phonation-based cues (Sun, 2015). Dege, despite having a higher baseline error rate overall, exhibits comparatively small deltas under pitch removal, suggesting that its lexical contrasts may still be anchored in non-pitch features, or that limited training data has led to underrepresentation of the language’s tonal cues. Together, these outcomes corroborate the view that Kham varieties have not completed the shift toward fully developed tone systems (Suzuki, 2011; Kirby et al., 2022).

Overall, this pattern of results supports the broader idea of a *tonogenesis continuum* in Tibetan: atonal Amdo dialects appear minimally impacted by pitch manipulation, Ü-Tsang’s well-established tones produce sizable recognition deficits, and Kham falls in between, reflecting an incomplete transition away from historical voicing contrasts toward f_0 -based distinctions. These findings underscore the importance of capturing both segmental and suprasegmental factors in Tibetan ASR. They also suggest that, as a dialect shifts deeper into tonogenesis, the stakes of losing pitch cues rise commensurately, culminating in fully tonal systems whose lexical contrasts become acutely vulnerable to f_0 flattening.

5 Discussion

Our experiments offer insight into how pitch-flattening impacts ASR across Tibetan dialects at different stages of tonogenesis. Table 2 reveals a clear gradient pattern: the atonal Amdo varieties (Xiahe, Aba) experience only small increases in CER and WER upon pitch-flattening, pointing to a reliance on residual consonantal cues (e.g. complex onset clusters, voicing) rather than f_0 (Gesang and Gesang, 2002; Sun, 1986). At the other end of the tonogenesis continuum, the fully tonal Ü-Tsang dialects (Lhasa, Shigatse) incur larger recognition penalties, indicating that pitch has become a primary mechanism for lexical contrast (DeLancey, 2017; Lim, 2018).

The Kham varieties (Changdu, Dege) occupy an intermediate position consistent with descriptions of partial tonogenesis in which pitch has begun to assume a contrastive load but still coexists with

Language	Tone Status	Original		Flattened		Δ	
		CER	WER	CER	WER	CER	WER
Amdo							
Xiahe	Non-tonal	11.4	32.0	13.9	37.8	2.5	5.8
Aba	Non-tonal	18.2	52.5	20.2	56.3	2.0	3.8
Kham							
Changdu	Intermediate	24.7	52.3	30.3	61.3	5.6	9.0
Dege	Intermediate	47.5	90.2	49.2	91.7	1.7	1.5
Ü-Tsang							
Lhasa	Tonal	17.7	48.6	23.7	59.3	6.0	10.7
Shigatse	Tonal	17.5	49.0	25.0	62.9	7.5	13.9

Table 2: Character error rate (CER) and word error rate (WER) in % for six Tibetan dialects under original vs. pitch-flattened conditions, along with their respective Δ (Flattened–Original). Dialects are ordered along the tonogenesis continuum from atonal (Amdo) through intermediate (Kham) to fully tonal (Ü-Tsang).

other cues (Sun, 2015; Suzuki, 2011). Changdu exhibits moderate performance drops (+5.6% CER, +9.0% WER), suggesting that both pitch and residual segmental features remain important. Dege shows a smaller delta despite a higher overall error rate, possibly indicating that the dialect’s emergent tonal distinctions are underpinned by breathiness or voicing contrasts, or that limited training data dampened the model’s reliance on f_0 cues. Notably, Shigatse shows the largest WER increase (+13.9%), likely reflecting its well-documented four-way tonal system where pitch carries substantial lexical load with fewer redundant segmental cues than Lhasa (Lim, 2018). Taken together, these patterns strengthen the conclusion that Tibetan dialects vary in pitch dependence according to their historical stage of tonogenesis, mirroring prior research highlighting the gradual shift from archaic voicing contrasts to full-blown pitch-based systems (Haudricourt, 1954; Hombert et al., 1979; Ratliff, 2015).

An interesting fact of Tibetan languages is the use of a shared script dating back to Classical Tibetan. This orthography encodes older consonant clusters and voicing distinctions, thus leaving tonal contrasts unmarked for modern varieties as a result of diachronic development (Wylie, 1959). Consequently, the ASR models must learn any pitch-based distinctions directly from acoustic input. For dialects like Amdo, which retain many of the historical segmental cues, pitch contributes little to lexical disambiguation and can be removed with minimal impact on recognition. In fully tonal dialects like Lhasa, however, sound change has reduced or eliminated these consonantal cues, render-

ing f_0 indispensable for word identification (Lim, 2018).

From a typological standpoint, these results shed fresh light on how functional load interacts with partial or incipient tone. The notion of *functional load* proposes that the more minimal pairs rely on a given contrast, the more catastrophic it is to remove that contrast from the signal (Surendran and Levow, 2004). However, in a transitional system such as Kham, pitch may be correlated with, or redundant to, residual voicing or breathiness cues (Sun, 2015; Brunelle and Kirby, 2016). As a result, simply counting tone-based minimal pairs can overstate how often pitch alone distinguishes lexical items (Hombert et al., 1979; Ratliff, 2015). Our work shows that flattening f_0 does degrade recognition in Kham, but not universally to the extent observed in fully tonal dialects. Thus, while functional load captures *some* of the story, it cannot fully explain how segmental and suprasegmental cues get dynamically reweighted throughout tonogenesis. A purely text-based estimate of tone’s load would miss the subtle interplay between pitch and other laryngeal features in transitional dialects (Hu, 2012; Kirby et al., 2022).

Overall, these findings highlight the practical challenge for Tibetan ASR of capturing both the legacy of archaic consonant distinctions and the emergent or fully developed reliance on pitch. As Zhao et al. (2020); Qin et al. (2022) have shown, new acoustic features such as voice quality and breathiness may improve robustness for dialects in flux. A next step is a more granular error analysis, focusing on whether certain segments or syllables are more prone to confusion when f_0 is removed.




Figure 2: A schematic illustrating how Tibetan script maps to Wylie transliteration (Wylie, 1959), then to acoustic features such as pitch and voicing, which the ASR model weighs during decoding.

Such studies could clarify how partial tone systems, where voicing and pitch contrasts remain intertwined, may ultimately evolve into the strongly pitch-dependent systems exemplified by Lhasa. In addition, these insights could motivate new approaches to multi-dialect Tibetan ASR, perhaps by selectively conditioning on pitch or phonation cues only where they carry crucial lexical information.

In sum, our experiments confirm that tone systems operate along a broad continuum, from atonal to partly tonal to fully tonal. By mapping each Tibetan language’s performance profile under pitch-flattening, we obtain evidence for how the process of tonogenesis reshapes a language’s reliance on f_0 . The results not only corroborate long-standing theories of tone development (Matisoff, 1999; Hombert et al., 1979; Haudricourt, 1954), but also point to new computational challenges. For emerging tone systems, we must move beyond static estimates of functional load to account for how segmental and suprasegmental cues blend to encode lexical meaning.

6 Conclusion

We have shown that pitch-flattening reveals systematic variation in ASR performance across Tibetan varieties that directly corresponds to their position along the tonogenesis continuum. Amdo dialects (e.g., Xiahe, Aba) incurred minimal penalty from pitch suppression, reflecting their atonal nature and continued reliance on consonantal contrasts. Ü-Tsang dialects (Lhasa, Shigatse) suffered sharp error increases that underscore the functional load of pitch for lexical distinctions in fully tonal systems. Kham varieties (Changdu, Dege) fell between these extremes, showing partial reliance on f_0 alongside residual voicing or phonation cues.

From an ASR standpoint, these findings rein-

force that a one-size-fits-all acoustic modeling approach may be insufficient for languages undergoing or nearing completion of tonogenesis. Systems trained on orthographies that only encode historical clusters or voicing may inadvertently overlook emergent pitch contrasts, or conversely fail to leverage lingering segmental cues in transitional dialects. By showing that pitch removal disproportionately degrades recognition for tonal varieties, our results quantify how much current models already rely on f_0 , and suggest that this reliance varies predictably with a dialect’s stage of tonogenesis. These observations motivate future work on multi-dialect Tibetan ASR that explicitly accounts for the shifting balance between consonantal and pitch-based cues across the tonogenesis continuum.

Overall, our work highlights how pitch manipulation offers a straightforward window on the evolving functional load of tone. While techniques like measuring minimal pairs or examining segmental environments can provide a textual or phonological perspective on tone, computational experiments using pitch-flattening yield concrete evidence of how critical f_0 is for modern-day lexical discrimination. In doing so, we situate Tibetan within a cross-linguistic tonogenetic trajectory that further underscores the interplay of segmental and suprasegmental features. Future research could extend these experiments by incorporating more dialects in flux, examining detailed confusion patterns for specific segments, or exploring advanced acoustic correlates (e.g., breathiness) that may serve as additional or alternative cues in emergent tone systems.

7 Limitations

Several limitations affect the scope and generalizability of our findings. First, our study examines only six Tibetan varieties, which, while representing key points on the tonogenesis continuum, cannot capture the full diversity of Tibetan dialectology. The TIBMD@MUC corpus provides valuable data but is constrained by uneven speaker representation across varieties and relatively small test sets (approximately 30 minutes per dialect).

Second, due to the realities of linguistic fieldwork, our training and test datasets often contained recordings from the same speakers. This may inflate performance estimates and limit our ability to assess model robustness to speaker variability. Future work should validate these findings with genuinely held-out speakers to better gauge gener-

alization.

Third, the shared Tibetan orthography, while linguistically interesting, presents challenges for ASR evaluation. The script encodes historical consonant clusters rather than modern phonological contrasts, and transcription conventions may vary within and across datasets. This inconsistency particularly affects our ability to assess errors in phonologically complex categories.

References

- Arun Babu, Changhan Wang, Andros Tjandra, Kushal Lakhotia, Qiantong Xu, Naman Goyal, Kritika Singh, Patrick von Platen, Yatharth Saraf, Juan Pino, Alexei Baevski, Alexis Conneau, and Michael Auli. 2021. [XLS-R: Self-supervised Cross-lingual Speech Representation Learning at Scale](#). *arXiv preprint*. ArXiv:2111.09296 [cs, eess].
- Marc Brunelle and James Kirby. 2016. [Tone and Phonation in Southeast Asian Languages: Tone and Phonation in Southeast Asian Languages](#). *Language and Linguistics Compass*, 10(4):191–207.
- Rolando Coto-Solano. 2021. [Explicit Tone Transcription Improves ASR Performance in Extremely Low-Resource Languages: A Case Study in Bribri](#). In *Proceedings of the First Workshop on Natural Language Processing for Indigenous Languages of the Americas*, pages 173–184, Online. Association for Computational Linguistics.
- Scott DeLancey. 2017. Lhasa Tibetan. In *The Sino-Tibetan Languages*.
- George van Driem. 2001. [Languages of the Himalayas](#). Brill.
- Jan Edwards and Mary E. Beckman. 1988. [Articulatory Timing and the Prosodic Interpretation of Syllable Duration](#). *Phonetica*, 45(2-4):156–174. Publisher: De Gruyter Mouton.
- Qian-Jie Fu, Fan-Gang Zeng, Robert V Shannon, and Sigfrid D Soli. 1998. Importance of tonal envelope cues in Chinese speech recognition. *The Journal of the Acoustical Society of America*, 104(1):505–510. Publisher: Acoustical Society of America.
- Jumian Gesang and Yangjing Gesang. 2002. Zangyu fangyan gailun. *Beijing: Minzu chubanshe*.
- André-Georges Haudricourt. 1954. De l’origine des tons en vietnamien. *Journal Asiatique*, 242:69–82.
- Einar Haugen. 1966. [Dialect, Language, Nation](#). *American Anthropologist*, 68(4):922–935.
- Jean-Marie Hombert. 1977. [Development of tones from vowel height?](#) *Journal of Phonetics*, 5(1):9–16.
- Jean-Marie Hombert, John J. Ohala, and William G. Ewan. 1979. [Phonetic Explanations for the Development of Tones](#). *Language*, 55(1):37–58. Publisher: Linguistic Society of America.
- Fang Hu. 2012. Tonogenesis in Lhasa Tibetan – Towards a gestural account. *Consonant Clusters and Structural Complexity*, 26:231.
- James Kirby, Pittayawat Pittayaporn, and Marc Brunelle. 2022. [Transphonologization of onset voicing: revisiting Northern and Eastern Kmhmu’](#). *Phonetica*, 79(6):591–629.
- James P. Kirby. 2014. [Incipient tonogenesis in Phnom Penh Khmer: Acoustic and perceptual studies](#). *Journal of Phonetics*, 43:69–85.
- Xin Lei, Manhung Siu, Mei-Yuh Hwang, Mari Ostendorf, and Tan Lee. 2006. [Improved tone modeling for Mandarin broadcast news speech recognition](#). In *Interspeech 2006*, pages paper 1752–Tue3A2O.4–0. ISCA.
- Shang-wen Li, Yow-bang Wang, Liang-che Sun, and Lin-shan Lee. 2011. [Improved tonal language speech recognition by integrating spectro-temporal evidence and pitch information with properly chosen tonal acoustic units](#). In *Interspeech 2011*, pages 2293–2296. ISCA.
- Siyu Liang and Gina-Anne Levow. 2025. [Tone in Perspective: A Computational Typological Analysis of Tone Function in ASR](#). In *Proceedings of the 7th Workshop on Research in Computational Linguistic Typology and Multilingual NLP*, pages 82–92, Vienna, Austria. Association for Computational Linguistics.
- Keh Sheng Lim. 2018. *The Tonal and Intonational Phonology of Lhasa Tibetan*. Ph.D. thesis.
- James A Matisoff. 1973. Tonogenesis in southeast Asia. *Consonant types and tone*, 1(1):71–96.
- James A Matisoff. 1999. Tibeto-Burman tonology in an areal context. In *Proceedings of the symposium “Crosslinguistic studies of tonal phenomena: Tonogenesis, Japanese Accentology, and Other Topics*, pages 3–31. Tokyo: Tokyo University of Foreign Studies, Institute for the Study of . . .
- Stephen Matthews and Virginia Yip. 2013. [Cantonese: A Comprehensive Grammar](#), 2 edition. Routledge, London.
- Yadong Niu, Fei Chen, and Jing Chen. 2019. [The effect of F0 contour on the intelligibility of Mandarin Chinese for hearing-impaired listeners](#). *The Journal of the Acoustical Society of America*, 146(2):EL85–EL91.
- Aniruddh D. Patel, Yi Xu, and Bei Wang. 2010. [The role of F0 variation in the intelligibility of Mandarin sentences](#). In *Speech Prosody 2010*, pages paper 890–0. ISCA.

- Siqing Qin, Longbiao Wang, Sheng Li, Jianwu Dang, and Lixin Pan. 2022. [Improving low-resource Tibetan end-to-end ASR by multilingual and multilevel unit modeling](#). *EURASIP Journal on Audio, Speech, and Music Processing*, 2022(1):2.
- Martha Ratliff. 2015. [Tonoexodus, Tonogenesis, and Tone Change](#). In Patrick Honeybone and Joseph Salmons, editors, *The Oxford Handbook of Historical Phonology*, page 0. Oxford University Press.
- Jackson (Ed) Sun. 2015. [Variegated tonal developments in Tibetan](#). pages 2.1M, 35–52 pages. Artwork Size: 2.1M, 35-52 pages Medium: PDF Publisher: Pacific Linguistics Version Number: 1.0.
- Jackson T-S Sun. 1986. Aspects of the Phonology of Amdo Tibetan. (*No Title*).
- Dinoj Surendran and Gina-Anne Levow. 2004. [The functional load of tone in Mandarin is as high as that of vowels](#). In *Speech Prosody 2004*, pages 99–102. ISCA.
- Hiroyuki Suzuki. 2011. Preaspiration and tonal development in Tibetan dialects of Khams, Shar and Amdo. In *Tone, accent and intonation in eastern Eurasian languages. The 18th Meeting of the Linguistic Circle for the Study of Eastern Eurasian Languages*, Aoyama Gakuin University, Tokyo, pages 9–17.
- Lilja Maria Sæbø, Eitan Grossman, and Steven Moran. 2025. [Tonogenesis: A diachronic typology](#). *Diachronica*, 42(3-4):451–478.
- Graham Thurgood. 2002. [Vietnamese and tonogenesis: Revising the model and the analysis](#). *Diachronica*, 19(2):333–363.
- Nicolas Tournadre. 2014. The Tibetic languages and their classification. *Trans-Himalayan linguistics: Historical and descriptive linguistics of the Himalayan area*, 266(1):105–29.
- Jiuju Wang, Hua Shu, Linjun Zhang, Zhaoxing Liu, and Yang Zhang. 2013. [The roles of fundamental frequency contours and sentence context in Mandarin Chinese speech intelligibility](#). *The Journal of the Acoustical Society of America*, 134(1):EL91–EL97.
- Turrell Wylie. 1959. A standard system of Tibetan transcription. *Harvard journal of Asiatic studies*, 22:261–267. Publisher: JSTOR.
- Maira Jean Winsland Yip. 2002. *Tone*. Cambridge textbooks in linguistics. Cambridge University Press, Cambridge ; New York.
- Jiahong Yuan, Neville Ryant, Xingyu Cai, Kenneth Church, and Mark Liberman. 2021. [Automatic recognition of suprasegmentals in speech](#). *arXiv preprint*. ArXiv:2108.01122 [cs].
- Yubin Zhang and James Kirby. 2020. [The role of F0 and phonation cues in Cantonese low tone perception](#). *The Journal of the Acoustical Society of America*, 148(1):EL40–EL45.
- Yue Zhao, Xiaona Xu, Jianjian Yue, Wei Song, Xiali Li, Licheng Wu, and Qiang Ji. 2020. [An open speech resource for Tibetan multi-dialect and multitask recognition](#). *Int. J. Comput. Sci. Eng.*, 22(2-3):297–304.

A Appendix

This appendix provides additional details on our fine-tuning hyperparameters for XLS-R 300m in both experiments.

A.1 XLS-R Fine-Tuning Hyperparameters

All training runs (for both Common Voice and TIBMD@MUC data) used the same set of essential hyperparameters, with only minor adjustments for batch size depending on GPU memory:

- **Model:** facebook/wav2vec2-xls-r-300m
- **CTC loss reduction:** *mean*
- **Batch Size:** 4 or 8 per device (gradient accumulation steps adjusted to keep effective batch size at 16)
- **Learning Rate:** 3×10^{-4}
- **Optimizer:** AdamW with $\beta_1 = 0.9$, $\beta_2 = 0.999$
- **Warmup Steps:** 500
- **Max Steps:** 2000
- **Vocabulary Size:** based on unique characters in the training corpus (including space or | as word delimiter).