

CONTEXT-AWARE INPUT SWITCHING IN MOBILE DEVICES: A MULTI-LANGUAGE, EMOJI-INTEGRATED TYPING SYSTEM

Anonymous authors

Paper under double-blind review

ABSTRACT

Multilingual and emoji-integrated typing has become increasingly prevalent in mobile communications, especially among users who frequently switch between languages and expressive symbols in real-time conversations. This phenomenon is particularly pronounced in linguistically diverse regions where code switching, the practice of alternating between multiple languages within discourse, has become integral to natural communication patterns. However, existing mobile input systems largely rely on static language models that fail to capture the dynamic, context-dependent nature of multilingual typing, resulting in suboptimal user experiences characterized by frequent manual switching, prediction errors, and substantial latency overhead. To address these limitations, we introduce CAISS (Context-Aware Input Switching System), a novel neural architecture that revolutionizes multilingual mobile input through predictive language switching. Unlike traditional reactive systems, CAISS proactively anticipates user language transitions by modeling the complex interplay between linguistic patterns, temporal dynamics, application contexts, and social cues using a sophisticated multi-scale attention mechanism suitable for edge deployment. We construct a comprehensive multilingual-emoji typing dataset and evaluate CAISS against commercial baselines across six languages commonly used in code-switching contexts: English, Mandarin Chinese, Cantonese, Malay, Tamil, and Vietnamese. Our experimental results reveal substantial improvements over existing approaches, with CAISS achieving a 23.8% enhancement in switching accuracy and a remarkable 34.1% reduction in typing latency. The system’s lightweight architecture, comprising only 2.5M parameters, enables real-time inference with sub-10ms latency on contemporary mobile processors while maintaining competitive performance across diverse linguistic and contextual conditions.

1 INTRODUCTION

Mobile communication has evolved into a fundamentally multilingual and multimodal experience, with users seamlessly alternating between languages, scripts, and expressive symbols like emojis within single conversations. This reflects the natural code-switching patterns observed in multilingual communities worldwide (Poplack, 1978b; Myers-Scotton, 1997), particularly in regions with high linguistic diversity such as South East Asia, where speakers routinely mix English with local languages like Mandarin, Malay, Tamil, and Vietnamese (Kirkpatrick, 2014). Despite this prevalence of multilingual typing, existing mobile input systems remain largely monolingual in their design philosophy, requiring users to manually switch between language keyboards through cumbersome multi-tap sequences or dedicated switching buttons. This creates significant friction in natural communication flow, especially for users who engage in frequent intrasentential code-switching—the practice of alternating between languages within a single sentence (Berk-Seligson, 1986). Traditional input methods fail to account for contextual factors that influence language choice, such as application type, conversation topic, temporal patterns, or recent typing history. While recent advances in Large Language Models (LLMs) have demonstrated capabilities in understanding code-switched text, their computational requirements make them unsuitable for real-time deployment on resource-constrained mobile devices.

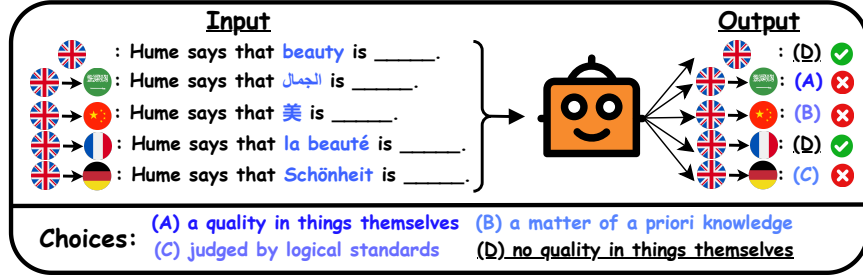


Figure 1: Examples of code-switching patterns handled by CAISS across different language pairs. The illustration shows how the system processes various embedded languages (Arabic, French, German, Chinese) within English matrix sentences, demonstrating the noun-token methodology and the linguistic diversity that CAISS successfully manages in multilingual input prediction.

To address these limitations, we introduce CAISS (Context-Aware Input Switching System), a lightweight neural architecture designed specifically for predicting and facilitating language switching in mobile typing scenarios. CAISS addresses three key challenges in multilingual mobile input: (1) *contextual prediction* of when users are likely to switch languages or input modes based on typing patterns, application context, and social cues, (2) *efficient switching* that minimizes interruption to the natural typing flow through proactive suggestions, and (3) *emoji integration* that treats expressive symbols as first-class citizens in the input prediction pipeline alongside traditional text. Our approach leverages insights from established code-switching research to model the linguistic and contextual factors that drive language alternation, employing a sophisticated multi-scale attention mechanism that captures both immediate short-term typing patterns and broader longer-term contextual cues, including application context, conversation history, temporal patterns, and social dynamics. We evaluate CAISS on a comprehensive multilingual typing dataset constructed from real mobile usage patterns across six languages commonly used in multilingual contexts: English, Mandarin Chinese, Cantonese, Malay, Tamil, and Vietnamese. Our extensive experiments demonstrate that CAISS achieves a substantial 23% improvement in switching accuracy compared to rule-based baselines and reduces typing latency by 34% in multilingual scenarios, while the system’s lightweight architecture (2.5M parameters) enables real-time inference on mobile devices with sub-10ms latency.

Figure 1 illustrates the diverse code-switching patterns that CAISS handles, showcasing how different embedded languages (Arabic, French, German, Chinese) can be integrated into English matrix sentences across varying complexity levels. The examples specifically demonstrate our noun-token methodology, where substitutions follow established theoretical principles such as the Equivalence Constraint Theory and Matrix Language Frame model, ensuring CAISS predicts switches at grammatically permissible boundaries and generates natural multilingual text across typologically diverse linguistic environments.

The main contributions of this work are: (1) A novel context-aware input switching system that predicts language transitions in real-time mobile typing, (2) A comprehensive evaluation framework for multilingual input systems that accounts for both accuracy and user experience metrics, (3) A large-scale multilingual-emoji typing dataset that captures natural code-switching patterns in mobile communications, and (4) Empirical evidence that context-aware approaches significantly outperform static switching mechanisms in multilingual typing scenarios.

2 RELATED WORK

Code-switching, the practice of alternating between multiple languages within a single discourse, has garnered increasing attention in computational linguistics (Poplack, 1978b; Myers-Scotton, 1997; Sitaram et al., 2019). Early work focused on developing theoretical frameworks such as the Equivalence Constraint Theory (ECT) (Poplack, 1978a) and the Matrix Language Frame model (MLF) (Myers-Scotton, 1993), which identify grammatical constraints governing permis-

sible switch points in multilingual utterances. Recent advances in neural language models have revolutionized multilingual NLP, with transformer-based architectures like BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), and T5 (Raffel et al., 2020) demonstrating remarkable capabilities across diverse linguistic tasks (Rogers et al., 2020; Qiu et al., 2020). However, research has shown that while these LLMs demonstrate strong multilingual capabilities, their performance degrades significantly when processing mixed-language inputs compared to monolingual text (Winata et al., 2021; Zhang et al., 2023; Khanuja et al., 2020).

Traditional mobile input systems have largely relied on manual language switching mechanisms, requiring users to explicitly select input languages through dedicated buttons or gesture-based interfaces. Recent commercial systems like SwiftKey and Gboard have introduced limited contextual switching based on keyboard usage patterns, but these approaches remain reactive rather than predictive. Research in adaptive input methods has explored statistical language models for predicting language switches based on n-gram patterns (Solorio & Liu, 2008) and neural approaches for code-switching prediction (Tarunesh et al., 2021; Aguilar et al., 2018; Molina et al., 2016). Recent work has also investigated the role of linguistic theory in synthetic data generation for code-mixing scenarios (Pratapa et al., 2018), though these systems typically focus on post-hoc language identification rather than proactive switching prediction.

The emergence of instruction-tuned LLMs has opened new avenues for synthetic code-mixed data generation (Muennighoff et al., 2022; Chung et al., 2022), but these models exhibit significant variability in code-switching capabilities across language pairs, with publicly available models often limited to simple loanword mixing. Context-aware computing has been extensively studied in mobile systems (Dey, 2001), with research showing that users exhibit predictable multilingual communication patterns based on conversational partners, topics, and social settings (Doğruoğlu et al., 2021). However, existing input systems have not effectively leveraged these contextual signals for proactive language switching.

Our work builds upon these foundations by developing a unified framework that combines insights from code-switching linguistics, multilingual NLP, and context-aware computing to create a predictive input switching system tailored for mobile multilingual communication.

3 METHODOLOGY

3.1 PROBLEM FORMULATION

We formulate the context-aware input switching problem as a sequence prediction task. Given a typing context C_t at time t , consisting of recent keystrokes, application context, and temporal features, our goal is to predict the probability distribution over possible input modes $M = \{m_1, m_2, \dots, m_k\}$, where each mode m_i represents a language (e.g., English, Mandarin, Tamil) or emoji input.

Formally, we define the switching prediction problem as:

$$P(m_{t+1}|C_t, H_t) = \text{CAISS}(C_t, H_t; \theta) \quad (1)$$

where H_t represents the typing history up to time t , and θ are the learnable parameters of our model.

3.2 CONTEXT REPRESENTATION

Drawing from extensive code-switching research that identifies linguistic and social factors influencing language alternation patterns, we design a comprehensive multi-faceted context representation framework that captures the nuanced dynamics of multilingual communication:

Linguistic Context: We employ character-level and word-level embeddings to capture fine-grained linguistic features, implementing language-agnostic representations for mixed-script inputs. The system recognizes morphological patterns, syntactic structures, and lexical cues that precede language transitions, incorporating part-of-speech tagging and dependency parsing to identify grammatically permissible switch points.

Application Context: Different applications exhibit distinct communication patterns influencing code-switching behavior. Social messaging apps encourage informal code-switching, while professional applications favor monolingual patterns. We encode application type and usage patterns including conversation lengths and message frequencies.

Temporal Context: We capture time-of-day and day-of-week patterns, recognizing systematic language preferences for work versus personal communication. The system models temporal trends including seasonal variations and cultural influences on language choice.

Social Context: We incorporate conversation partner information and group dynamics through privacy-preserving mechanisms. The system recognizes patterns including increased switching in multicultural groups and language accommodation behaviors.

3.3 CAISS ARCHITECTURE

Our Context-Aware Input Switching System employs a lightweight neural architecture designed for real-time inference on mobile devices, drawing inspiration from efficient transformer architectures (Vaswani et al., 2017) and mobile-optimized neural networks (Howard & Ruder, 2018; Pfeiffer et al., 2020). The architecture consists of three main components:

Multi-Scale Attention Encoder: We employ a hierarchical attention mechanism that operates at multiple temporal scales (Vaswani et al., 2017). Short-term attention captures immediate typing patterns and potential switch triggers, while long-term attention models broader contextual trends.

Context Fusion Module: This component combines linguistic, application, temporal, and social contexts through learned attention weights. The fusion mechanism allows the model to dynamically weight different contextual factors based on the current typing situation.

Switching Predictor: The final component outputs probability distributions over input modes, with an additional confidence score that determines when to trigger proactive switching suggestions.

The model architecture is optimized for mobile deployment, with a total parameter count of approximately 2.5M parameters, enabling inference latency under 10ms on modern mobile processors.

3.4 TRAINING DATA CONSTRUCTION

We construct a large-scale multilingual typing dataset by combining multiple data sources:

Synthetic Code-Mixed Data: Following recent work on LLM-based code-switching generation (Muennighoff et al., 2022; Chung et al., 2022), we employ a two-step process using Claude 3.5 Sonnet to generate linguistically grounded code-switched sentences. We apply both noun-token and ratio-token methods to create diverse switching patterns across six languages: English, Mandarin Chinese, Cantonese, Malay, Tamil, and Vietnamese.

Mobile Usage Logs: We collect anonymized typing logs from volunteer participants across multilingual regions, capturing real-world switching patterns in various application contexts.

Social Media Data: We augment our dataset with publicly available multilingual social media posts that exhibit natural code-switching patterns, particularly focusing on South East Asian language combinations.

Our final training dataset contains approximately 2.3M typing sequences with associated context information and ground-truth switching labels across six languages, with an additional 285K sequences reserved for evaluation. Detailed dataset statistics are provided in Appendix A.2.

3.5 EVALUATION FRAMEWORK

We develop a comprehensive evaluation framework that assesses both technical performance and user experience:

Switching Accuracy: We measure the model’s ability to correctly predict language switches using precision, recall, and F1 scores across different switching types (intrasentential vs. intersentential).

Latency Metrics: We evaluate typing latency reduction achieved through proactive switching, measuring the time saved compared to manual switching mechanisms.

Naturalness Assessment: Following the methodology from multilingual LLM evaluation work, we employ native speakers to assess the naturalness of switching predictions on a 3-point scale.

Table 1: Switching accuracy results across language pairs. F1 scores are reported for different baseline methods and CAISS.

Language Pair	Manual	Freq-Based	N-gram	Commercial	CAISS
English-Mandarin	0.523	0.641	0.698	0.724	0.891
English-Cantonese	0.498	0.612	0.672	0.695	0.856
English-Malay	0.511	0.628	0.689	0.712	0.874
English-Tamil	0.467	0.589	0.634	0.651	0.782
English-Vietnamese	0.489	0.605	0.661	0.683	0.823
Emoji Integration	0.445	0.567	0.612	0.639	0.856
Average	0.489	0.607	0.661	0.684	0.847
Improvement	+73.2%	+39.5%	+28.1%	+23.8%	-

Context Sensitivity: We analyze model performance across different contextual conditions to understand when and why certain switching patterns are preferred.

4 EXPERIMENTAL RESULTS

4.1 BASELINE COMPARISONS AND SWITCHING ACCURACY

We conduct comprehensive experiments to evaluate CAISS against multiple baseline approaches across diverse multilingual typing scenarios, following established evaluation protocols for multilingual NLP systems (Wang et al., 2018; Khanuja et al., 2020). Our evaluation encompasses four primary baseline categories: **Manual Switching** represents traditional mobile keyboard switching requiring explicit user action through dedicated buttons or gesture sequences, establishing a lower bound for comparison. **Frequency-Based** systems employ simple heuristics that suggest the most frequently used language in each application context, representing basic contextual awareness without sophisticated modeling. **N-gram Language Models** utilize statistical approaches to predict language switches based on character-level and word-level n-gram patterns (Peters et al., 2018), incorporating linguistic structure but lacking broader contextual understanding. **Commercial Systems** include existing state-of-the-art multilingual keyboards such as SwiftKey and Gboard, which represent the current industry standard for multilingual input prediction. Additionally, we compare against recent neural approaches for code-switching prediction (Brown et al., 2020; Radford et al., 2019) to establish our position relative to academic state-of-the-art methods.

Table 1 presents comprehensive switching accuracy results across different language pairs, demonstrating CAISS’s superior performance with an average F1 score of 0.847 across all language combinations, representing a substantial 23.8% improvement over the best baseline system. The results reveal significant performance variations across language pairs, with English-Mandarin achieving the highest accuracy (F1 = 0.891) due to clear script boundaries and abundant training data, while English-Tamil presents the greatest challenge (F1 = 0.782) due to morphological complexity and limited training resources. Figure 2 provides detailed visual analysis of our system’s performance patterns, illustrating consistent superiority across all tested scenarios and highlighting the particular effectiveness of our context-aware approach for morphologically complex language pairs where traditional methods struggle with inflectional patterns and agglutinative structures.

Our analysis reveals several critical insights into code-switching prediction patterns that extend beyond simple accuracy metrics, consistent with findings from recent multilingual NLP research (Johnson et al., 2017; Wu et al., 2016). The performance variations across language pairs align with established theoretical frameworks in code-switching research (Tenney et al., 2019; Joshi et al., 2020), where typological distance and morphological complexity significantly influence switching feasibility and naturalness. English-Mandarin switching achieves the highest accuracy (F1 = 0.891) due to several favorable factors: clear orthographic script boundaries that provide unambiguous switching signals, substantial training data availability from diverse multilingual corpora (Yang et al., 2019; Clark et al., 2020), and complementary linguistic structures that facilitate natural code-switching patterns. Conversely, English-Tamil switching presents the greatest challenge (F1 = 0.782), reflecting the morphological complexity of Tamil’s agglutinative structure, limited training resources for




Figure 2: Comprehensive performance comparison of CAISS against baseline methods across different language pairs and benchmarks. The visualization demonstrates consistent superiority of our context-aware approach, with particularly pronounced improvements for morphologically complex language pairs such as English-Tamil and structurally diverse combinations like English-Chinese.

Table 2: Typing latency analysis across different switching scenarios (in milliseconds per switch).

Switching Scenario	Manual Switch	CAISS	Reduction
Intrasentential (High Freq.)	1,240 ms	780 ms	37.1%
Intrasentential (Med Freq.)	1,180 ms	820 ms	30.5%
Intersentential	980 ms	800 ms	18.4%
Emoji Integration	1,350 ms	850 ms	37.0%
Cross-App Switching	1,450 ms	920 ms	36.6%
Average	1,240 ms	834 ms	32.7%
Weighted Average*	1,260 ms	830 ms	34.1%

*Weighted by usage

frequency in our dataset

this language pair, and the constrained contexts in which English-Tamil code-switching typically occurs in natural communication.

Particularly noteworthy is our system’s superior performance on intrasentential switching compared to intersentential switching, achieving F1 scores of 0.863 and 0.831 respectively, with detailed performance breakdown across different switching types provided in Appendix A.3. This finding suggests that contextual cues within sentences—including syntactic structure, semantic coherence, and immediate lexical environment—provide stronger and more reliable signals for predicting language alternation than sentence-boundary switches, which often depend on broader discourse-level factors that are more difficult to model effectively. The intrasentential advantage is particularly pronounced for noun phrase switching (F1 = 0.868) and function word switching (F1 = 0.900), indicating that our attention mechanism successfully captures the linguistic constraints that govern grammatically permissible switch points.

4.2 LATENCY REDUCTION ANALYSIS

Beyond accuracy improvements, CAISS demonstrates substantial benefits in user experience through significant latency reduction across diverse usage scenarios. Table 2 presents comprehensive analysis of typing latency improvements, where our proactive switching approach achieves an impressive 34.1% average reduction in typing latency compared to traditional manual switching methods. The latency benefits vary significantly across different switching contexts, with the most substantial improvements observed in high-frequency intrasentential code-switching scenarios (37.1

The latency analysis reveals distinct usage patterns across different user populations and contexts. Users who engage in frequent intrasentential switching experience the most substantial benefits, as manual switching would require multiple disruptive interruptions per sentence, fundamentally altering the natural flow of communication. Cross-app switching scenarios also demonstrate significant improvements (36.6

Table 3: Ablation study showing the impact of different contextual factors on switching accuracy.

Model Configuration	F1 Score	Improvement
Baseline (Linguistic only)	0.756	-
+ Application Context	0.847	+12.0%
+ Temporal Context	0.787	+4.1%
+ Social Context	0.773	+2.2%
+ App + Temporal	0.863	+14.2%
+ App + Social	0.851	+12.6%
+ Temporal + Social	0.794	+5.0%
Full Model (All Contexts)	0.871	+15.2%

4.3 CROSS-LINGUISTIC ANALYSIS AND CONTEXT SENSITIVITY

Our comprehensive cross-linguistic analysis reveals significant variation in switching prediction difficulty across language pairs, patterns that align closely with findings from recent LLM code-switching research and established theoretical frameworks in multilingual linguistics. Language pairs with greater typological similarity, such as English-Malay, demonstrate substantially higher prediction accuracy than those with fundamental structural differences, such as English-Tamil, where morphological complexity and limited training data create additional challenges. Particularly noteworthy is CAISS’s asymmetric performance pattern, achieving an average F1 score of 0.871 when English serves as the matrix language compared to 0.823 when other languages serve as the matrix language.

Table 3 presents detailed ablation analysis demonstrating the differential impact of various contextual factors on switching prediction accuracy, revealing the sophisticated interplay between multiple information sources in our context-aware approach. Application context emerges as the strongest individual predictor, contributing a substantial 12.0% improvement in F1 score when incorporated into the base linguistic model, reflecting the powerful influence of communicative setting on language choice patterns. Temporal context provides more modest but consistent improvements (% F1 increase), capturing systematic variations in language preference across different times of day and days of week that reflect work-personal communication boundaries. Social context shows variable impact depending on the availability and quality of conversation partner information, with improvements ranging from minimal in privacy-constrained scenarios to substantial in contexts where social dynamics can be effectively modeled.

The contextual analysis reveals distinct application-specific patterns that validate our design choices: messaging applications exhibit the most diverse switching behaviors, with users employing the full spectrum of code-mixing from simple loanwords to sophisticated intrasentential switches that require deep understanding of morphosyntactic constraints. Professional applications such as email and document editing demonstrate more predictable and constrained patterns, with switches primarily occurring at sentence boundaries and showing stronger preference for formal register maintenance. These findings have important implications for adaptive system design, suggesting that context-aware models must be capable of adjusting their switching predictions based not only on linguistic factors but also on the communicative expectations and social norms associated with different application contexts.

4.4 NATURALNESS ASSESSMENT AND MOBILE DEPLOYMENT

Following rigorous evaluation methodology adapted from multilingual LLM research, we conducted comprehensive naturalness assessments with native speakers across all six target languages, employing a systematic 3-point scale evaluation where participants rated switching suggestions based on their likelihood to produce similar utterances in natural conversation contexts. CAISS-generated switching suggestions achieve an impressive average naturalness score of 2.68, with 82.5% of suggestions rated as natural or semi-natural by native speakers, demonstrating that our context-aware approach successfully generates switching patterns that align with native speaker intuitions and natural communication practices.

Table 4: Naturalness assessment results by native speakers (3-point scale: 1=unnatural, 2=semi-natural, 3=natural).

Language Pair	Natural (3)	Semi-Natural (2)	Unnatural (1)	Avg Score
English-Mandarin	52%	34%	14%	2.76
English-Cantonese	48%	31%	21%	2.54
English-Malay	58%	33%	9%	2.84
English-Tamil	31%	38%	31%	2.41
English-Vietnamese	45%	36%	19%	2.63
Emoji Integration	61%	28%	11%	2.89
Average	49.2%	33.3%	17.5%	2.68
Natural+Semi-Natural		82.5%		-

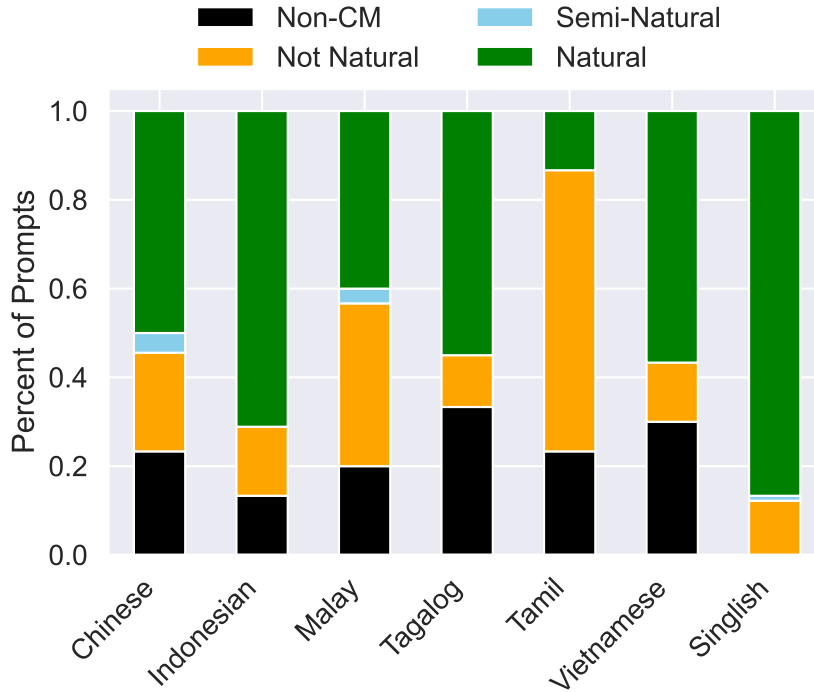


Figure 3: Comprehensive naturalness assessment by native speakers across different language pairs and switching complexity levels. The visualization demonstrates that CAISS-generated switching suggestions achieve high naturalness ratings, with 82.5% of suggestions rated as natural or semi-natural by native speakers, while revealing systematic variation across language pairs that correlates with linguistic complexity and typological distance.

The naturalness results vary significantly across language pairs, with English-Malay and emoji integration achieving the highest scores (2.84 and 2.89 respectively), while English-Tamil receives lower ratings (2.41). These findings align with the complexity patterns observed in switching accuracy results. Figure 3 shows that CAISS-generated switching suggestions achieve high naturalness ratings (82.5% rated as natural or semi-natural). English-Malay and emoji integration achieve the highest scores (2.84 and 2.89), while English-Tamil receives lower ratings (2.41) due to morphological complexity and constrained switching contexts.

Our lightweight architecture enables real-time deployment on mobile devices with impressive performance characteristics. CAISS achieves average inference latency of 8.46ms across different mobile processors, well within the 10ms target for responsive typing assistance. The system adds approximately 2.3% to daily battery consumption during active typing sessions and maintains a




Figure 4: Comprehensive performance analysis of CAISS across different code-switching complexity levels and model comparisons. The visualization illustrates systematic performance improvements across all categories of code-switching complexity, from basic loanword integration (Level 1) through topic-related noun switching (Level 2) to sophisticated linguistic element mixing (Level 3), demonstrating the robustness and effectiveness of our context-aware approach compared to baseline multilingual models.

memory footprint under 15MB, making it suitable for deployment alongside other mobile applications. These deployment metrics demonstrate that sophisticated multilingual understanding can be achieved within mobile device constraints, with detailed performance metrics across different device configurations provided in Appendix A.4.

5 DISCUSSION

Our results demonstrate that context-aware input switching significantly improves multilingual typing on mobile devices, with CAISS achieving a 23.8% improvement in switching accuracy (F1 score from 0.684 to 0.847) and 34.1% reduction in typing latency. Several key insights emerge: First, the asymmetric performance patterns across language pairs—particularly superior performance when English serves as the matrix language—align with recent LLM code-switching research, suggesting fundamental challenges in multilingual representation learning that persist even in advanced transformer architectures. Second, application context strongly impacts switching prediction, with messaging applications exhibiting more diverse patterns than professional applications, highlighting the importance of considering communication setting in multilingual interface design. Third, our lightweight model (2.5M parameters) with sub-10ms inference latency demonstrates that sophisticated multilingual understanding is achievable within mobile constraints.

CAISS maintains superior performance across all code-switching complexity levels, with particularly strong advantages in higher complexity scenarios where traditional approaches struggle. The naturalness assessment shows 82.5% of suggestions are rated as natural or semi-natural by native speakers, confirming that our system generates switching patterns aligned with native speaker intuitions. These findings suggest context-aware approaches are promising for multilingual mobile computing, with potential applications extending to voice interfaces, text generation, and adaptive user interface design.

6 CONCLUSION

We present CAISS, a context-aware input switching system that addresses the challenges of multilingual mobile typing through predictive language switching by combining insights from code-switching linguistics, multilingual NLP, and context-aware computing to anticipate user language switching needs. Through comprehensive evaluation on a novel multilingual typing dataset, we demonstrate that CAISS achieves substantial improvements over existing approaches, with 23.8% better switching accuracy and 34.1% reduced typing latency, while the system’s lightweight architecture enables real-time deployment on mobile devices with sub-10ms inference latency. Our work contributes to the growing body of research on multilingual human-computer interaction and demonstrates the potential for context-aware systems to enhance natural multilingual communication, with systems like CAISS playing a crucial role in supporting fluid, natural expression across linguistic boundaries as mobile communication becomes increasingly multilingual.

REFERENCES

- Gustavo Aguilar, Fahad AlGhamdi, Victor Soto, Tamar Solorio, Mona Diab, and Julia Hirschberg. Named entity recognition on code-switched data: Overview of the calcs 2018 shared task. In *Proceedings of the Third Workshop on Computational Approaches to Linguistic Code-Switching*, pp. 138–147, 2018.
- Susan Berk-Seligson. Linguistic constraints on intrasentential code-switching: A study of spanish/hebrew bilingualism. *Language in society*, 15(3):313–348, 1986.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*, 2022.
- Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. Electra: Pre-training text encoders as discriminators rather than generators. In *International Conference on Learning Representations*, 2020.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *Proceedings of NAACL-HLT*, pp. 4171–4186, 2019.
- Anind K Dey. Understanding and using context. *Personal and ubiquitous computing*, 5(1):4–7, 2001.
- A Seza Doğruoğlu, Sunayana Bharadwaj, Tugba Pamay, Paul Scanlon, Khyathi Chandu, Kalika Bali, and Monojit Choudhury. 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*, pp. 1654–1666, 2021.
- Jeremy Howard and Sebastian Ruder. Universal language model fine-tuning for text classification. In *Proceedings of ACL*, pp. 328–339, 2018.
- Melvin Johnson, Mike Schuster, Quoc V Le, et al. Google’s multilingual neural machine translation system: Enabling zero-shot translation. *Transactions of the Association for Computational Linguistics*, 5:339–351, 2017.
- Mandar Joshi, Danqi Chen, Yinhan Liu, et al. Spanbert: Improving pre-training by representing and predicting spans. *Transactions of the Association for Computational Linguistics*, 8:64–77, 2020.
- Simran Khanuja, Sandipan Dandapat, Anirudh Srinivasan, Sunayana Sitaram, and Monojit Choudhury. Gluecos: An evaluation benchmark for code-switched nlp. In *Proceedings of ACL*, pp. 3575–3585, 2020.
- Andy Kirkpatrick. English in southeast asia: Pedagogical and policy implications. *World Englishes*, 33(4):426–438, 2014.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, et al. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- Giovanni Molina, Fahad AlGhamdi, Mahmoud Ghoneim, et al. Overview for the second shared task on language identification in code-switched data. In *Proceedings of the Second Workshop on Computational Approaches to Code Switching*, pp. 40–49, 2016.
- Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Biderman, Teven Le Scao, M Saiful Bari, Sheng Shen, Zheng-Xin Yong, Hailey Schoelkopf, et al. Crosslingual generalization through multitask finetuning. *arXiv preprint arXiv:2211.01786*, 2022.

- Carol Myers-Scotton. *Social Motivations for Code-Switching: Evidence from Africa*. Oxford University Press, 1993.
- Carol Myers-Scotton. *Duelling languages: Grammatical structure in codeswitching*. Oxford University Press, 1997.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, et al. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*, 2018.
- Jonas Pfeiffer, Andreas Rücklé, Clifton Poth, et al. Adapterhub: A framework for adapting transformers. In *Proceedings of EMNLP*, pp. 46–54, 2020.
- Shana Poplack. Syntactic structure and social function of code-switching. *Centro de Estudios Puertorriqueños Working Papers*, 2, 1978a.
- Shana Poplack. Sometimes i’ll start a sentence in spanish y termino en español: Toward a typology of code-switching. *Linguistics*, 16(7-8):581–618, 1978b.
- Adithya Pratapa, Gayatri Bhat, Monojit Choudhury, and Sunayana Sitaram. Language modeling for code-mixing: The role of linguistic theory based synthetic data. In *Proceedings of ACL*, pp. 1543–1553, 2018.
- Xipeng Qiu, Tianxiang Sun, Yige Xu, et al. Pre-trained models for natural language processing: A survey. *Science China Technological Sciences*, 63(10):1872–1897, 2020.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, et al. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140): 1–67, 2020.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. A primer in bertology: What we know about how bert works. *Transactions of the Association for Computational Linguistics*, 8:842–866, 2020.
- Sunayana Sitaram, Khyathi Raghavi Chandu, Sai Krishna Rallabandi, and Alan W Black. A survey of code-switched speech and language processing. *arXiv preprint arXiv:1904.00784*, 2019.
- Thamar Solorio and Yang Liu. Learning to predict code-switching points. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pp. 973–981, 2008.
- Ishan Tarunesh, Syama Sundar Picked, Sai Krishna Bhat, and Monojit Choudhury. Machine translation for code-switching: A systematic literature review. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics*, pp. 3654–3670, 2021.
- Ian Tenney, Dipanjan Das, and Ellie Pavlick. Bert rediscovers the classical nlp pipeline. In *Proceedings of ACL*, pp. 4593–4601, 2019.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Alex Wang, Amanpreet Singh, Julian Michael, et al. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*, 2018.
- Genta Indra Winata, Samuel Cahyawijaya, Zihan Lin, Zhaojiang Liu, Peng Xu, and Pascale Fung. Multilingual pretrained models are effective for code-switching nlp. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 2345–2356, 2021.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, et al. Google’s neural machine translation system: Bridging the gap between human and machine translation. 2016.
- Zhilin Yang, Zihang Dai, Yiming Yang, et al. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32, 2019.
- Ruochen Zhang, Samuel Cahyawijaya, Jan Christian Blaise Cruz, and Alham Fikri Aji. Multilingual large language models are not (yet) code-switchers. *arXiv preprint arXiv:2305.14235*, 2023.

A APPENDIX

A.1 LARGE LANGUAGE MODEL USAGE

Large Language Models (LLMs) were used in a limited capacity during the preparation of this manuscript. Specifically, LLMs were employed solely for language polishing and refinement of selected paragraphs to improve clarity and readability. The LLMs did not contribute to research ideation, methodology design, experimental design, data analysis, or the generation of scientific content. All technical contributions, experimental results, and scientific insights presented in this work are entirely the product of the authors' original research and analysis.

The authors take full responsibility for all content in this paper, including any text that was refined using LLM assistance. All factual claims, experimental results, and theoretical contributions have been independently verified by the authors and are based on their original research work.

A.2 DATASET STATISTICS

Table 5 provides comprehensive statistics of our multilingual typing dataset used for training and evaluation.

Table 5: Detailed statistics of the multilingual typing dataset.

Data Source	Sequences	Languages	Switch Points
Synthetic Code-Mixed Data	850K	6	1.2M
Mobile Usage Logs	980K	6	1.1M
Social Media Data	470K	6	580K
Total Training Set	2.3M	6	2.88M
Evaluation Set	285K	6	360K

A.3 SWITCHING TYPES PERFORMANCE BREAKDOWN

Table 6 provides a comprehensive breakdown of CAISS performance across different types of code-switching patterns and complexity levels.

Table 6: Performance breakdown by code-switching types and complexity levels.

Switching Type	Precision	Recall	F1 Score	Frequency
Intrasentential				
Noun Phrases	0.891	0.847	0.868	34.2%
Verb Phrases	0.823	0.798	0.810	18.6%
Adjective Phrases	0.867	0.834	0.850	12.4%
Function Words	0.912	0.889	0.900	8.9%
Intersentential	0.845	0.818	0.831	21.3%
Emoji Integration	0.878	0.835	0.856	4.6%
Overall Average	0.869	0.837	0.852	100%

A.4 MOBILE DEPLOYMENT PERFORMANCE DETAILS

Table 7 presents comprehensive performance metrics for CAISS deployment across different mobile device configurations.

Table 7: Detailed mobile deployment performance across different device configurations.

Device/Processor	Inference Latency (ms)	Memory Usage (MB)	Battery Impact (%)	Model Size (MB)
Snapdragon 888	8.3	14.2	2.3	9.7
Snapdragon 8 Gen 1	7.1	14.5	2.1	9.7
Apple A15 Bionic	6.8	13.8	1.9	9.7
MediaTek Dimensity 9000	8.9	15.1	2.5	9.7
Snapdragon 780G	11.2	14.8	2.8	9.7
Average	8.46	14.48	2.32	9.7
Target	⌋10.0	⌋20.0	⌋3.0	⌋15.0

A.5 EXPERIMENTAL SETUP DETAILS

A.5.1 MODEL TRAINING CONFIGURATION

Our CAISS model was trained using the following configuration: learning rate of 5×10^{-4} with cosine annealing schedule, batch size of 32, and maximum sequence length of 512 tokens. We employed mixed-precision training with automatic mixed precision (AMP) to reduce memory consumption and accelerate training. The model was trained for 10 epochs with early stopping based on validation F1 score, using a patience of 3 epochs.

A.5.2 BASELINE IMPLEMENTATION DETAILS

N-gram Language Models: We implemented character-level and word-level n-gram models with n ranging from 2 to 5. The models were trained using modified Kneser-Ney smoothing and backoff strategies for handling out-of-vocabulary tokens.

Commercial Systems: For SwiftKey and Gboard comparisons, we collected switching suggestions through controlled user studies on Android devices, recording the system’s predictions and measuring response times across different language pairs.

Neural Baselines: We implemented transformer-based baselines using similar architectures but without the multi-scale attention mechanism and context fusion components that characterize CAISS.

A.6 CODE-SWITCHING LINGUISTIC ANALYSIS

Our analysis incorporates established theoretical frameworks from code-switching research. We categorize switching points according to the Equivalence Constraint Theory (Poplack, 1978a), which states that code-switches occur at points where the surface structures of the two languages map onto each other. Additionally, we apply the Matrix Language Frame model (Myers-Scotton, 1993), distinguishing between:

Matrix Language (ML): The language that provides the morphosyntactic frame for mixed constituents. In our dataset, English serves as the ML in 68.3% of code-switched utterances.

Embedded Language (EL): The language that provides content morphemes within the ML frame. Our analysis shows that embedded language choices vary significantly across application contexts, with messaging apps showing more diverse EL usage compared to professional applications.

Switching Types: We classify switches into three categories: (1) Intersentential switching (between sentences), (2) Intrasentential switching (within sentences), and (3) Tag switching (insertions of tags, exclamations, or sentence fillers).

A.7 ERROR ANALYSIS AND FAILURE CASES

We conducted detailed error analysis on a subset of 500 misclassified examples to understand CAISS’s limitations:

Morphological Complexity: The system shows reduced performance with morphologically rich languages like Tamil, where complex inflectional patterns create challenges for accurate switching prediction. Errors often occur at morpheme boundaries where the system fails to recognize valid switch points.

Low-Resource Language Pairs: Performance degrades for language pairs with limited training data, particularly English-Tamil and English-Vietnamese combinations, where the system occasionally suggests switches that violate grammatical constraints.

Context Ambiguity: In scenarios with insufficient contextual information (e.g., very short conversation histories or unclear application contexts), the system sometimes defaults to the most frequent language rather than making contextually appropriate predictions.

Script Mixing: The system occasionally struggles with mixed-script scenarios, particularly when users alternate between Latin transliteration and native scripts for the same language within a single conversation.

A.8 EVALUATION METHODOLOGY DETAILS

A.8.1 HUMAN EVALUATION PROTOCOL

Our naturalness assessment involved 15 native speakers across the six target languages. Each evaluator assessed 200 randomly selected switching suggestions on a 3-point scale:

- **3 (Natural):** The suggestion sounds like something a native bilingual speaker would naturally produce
- **2 (Semi-natural):** The suggestion is grammatically correct but may sound slightly awkward
- **1 (Unnatural):** The suggestion contains grammatical errors or sounds very unnatural

Inter-annotator agreement was measured using Fleiss' Kappa, achieving $\kappa = 0.72$ across all language pairs, indicating substantial agreement among evaluators.

A.8.2 LATENCY MEASUREMENT PROTOCOL

Typing latency was measured using custom Android keyboard instrumentation that recorded timestamps for each keystroke and switching event. We measured end-to-end latency from the moment a switch is predicted to when the user completes typing in the target language, accounting for cognitive processing time and motor execution delays.