

A Framework for Adaptive Voice Personalization in Conversational AI Systems

Enrique Zueco and Miguel Laguna
Laia Smart Solution SL
Zaragoza, Spain
hola@heylaia.com

Abstract—Conversational AI systems traditionally employ static voice interfaces that fail to adapt to caller demographics, despite robust evidence from social psychology demonstrating that demographic similarity in communication significantly increases trust, information recall, and perceived empathy. This paper presents Laia, the first comprehensive framework for real-time voice personalization in production conversational AI systems. Our approach integrates speech classification with controllable text-to-speech synthesis, enabling dynamic voice adaptation based on caller characteristics. Gender detection, built on fine-tuned Wav2Vec 2.0, achieves 98.89% test accuracy on LibriSpeech, while age and accent detection architectures are designed with training planned as future work. We present a unified system architecture achieving 461ms mean latency (82.3% meeting sub-500ms target), multi-task learning approaches that leverage shared representations across demographic classification tasks, and a hierarchical personalization strategy balancing voice library consistency with parameter-based adaptation. Beyond technical contributions, we address critical ethical considerations including GDPR compliance, bias mitigation strategies, transparency requirements through C2PA watermarking, and implementation of consent frameworks aligned with IA·Veu standards. This work provides both a production-ready implementation and a structured roadmap for empirical validation of demographic-aware voice personalization through large-scale A/B testing, establishing the foundation for rigorous investigation of voice adaptation in conversational AI.

Index Terms—Voice Personalization, Speech Synthesis, Speaker Characterization, Conversational AI, A/B Testing

I. INTRODUCTION

Conversational AI has become ubiquitous in customer service, sales, and healthcare applications. However, current systems typically employ static voice profiles that do not adapt to the caller’s demographic characteristics. This paper presents a novel approach to **adaptive voice personalization** that detects caller characteristics in real-time and dynamically adjusts the AI voice to optimize engagement.

A. Motivation and Background

Research in social psychology [?], [?] demonstrates that age similarity increases trust and persuasion effectiveness by up to 25%, gender matching improves information recall by 18%, and accent alignment reduces cognitive load and increases perceived empathy by 30%. Despite these findings, no comprehensive study has examined the impact of real-time voice personalization in production conversational AI systems.

Recent advances in speaker characterization include CNN-based architectures for age detection (75-80% accuracy) [?],

Wav2Vec 2.0 for gender classification (90-95% accuracy) [?], [?], and limited work on Spanish regional accent detection [?]. In voice synthesis, diffusion-based TTS achieves MOS 4.5 [?] while zero-shot voice conversion enables 50x speedup [?]. However, prior work focuses on single aspects of voice adaptation rather than comprehensive real-time personalization.

B. Contributions

This work makes the following contributions:

- 1) **Unified Framework:** Comprehensive system architecture for real-time detection of age, gender, and regional accent from voice
- 2) **Multi-Task Detection Models:** Integrated CNN+Transformer and Wav2Vec 2.0 architectures. Gender detection achieves 98.89% test accuracy on LibriSpeech. Age and accent detection architectures are designed, with training planned as future work
- 3) **Personalization Engine:** Modular voice adaptation system with hierarchical parameter adjustment strategies
- 4) **Production Implementation:** 461ms mean latency with 82.3% meeting sub-500ms target for natural conversation
- 5) **Ethical Foundation:** GDPR-compliant data handling, bias mitigation, C2PA watermarking, and IA·Veu-aligned consent frameworks
- 6) **Validation Roadmap:** Structured approach for empirical evaluation through large-scale A/B testing

C. Ethical Considerations

This work addresses the evolving ethical landscape for voice AI as of 2026. Following the Soundverse Ethical AI Framework and IA·Veu initiative standards [?], [?], our system requires explicit authorization for voice data usage and avoids web scraping without permissions. We implement traceability mechanisms enabling AI outputs to be traced to original creators, support micro-payment royalty systems, and embed attribution metadata in synthetic outputs aligning with C2PA standards [?]. The system includes disclosure mechanisms for synthetic voices, distinguishes between parody and deceptive use, and implements platform-level moderation for misuse prevention. Our implementation adheres to GDPR requirements for biometric data processing with regular audits, and addresses bias across demographic groups through diverse




Fig. 1: System architecture showing the real-time voice personalization pipeline with sub-500ms end-to-end latency target.

training data, regular fairness audits, and inclusive model design.

II. METHODOLOGY

A. System Architecture

Figure 1 illustrates our complete system architecture. The pipeline consists of four main stages: (1) **Feature Extraction** (100-150ms): MFCC (40 coefficients), prosody features (pitch, rhythm, energy), spectral features (formants, tilt); (2) **Classification Models** (150-200ms): Age (CNN + Transformer), Gender (Wav2Vec 2.0), Accent (Multi-task Wav2Vec); (3) **Personalization Engine** (50-100ms): Voice profile selection, parameter adjustment, linguistic adaptation; (4) **Voice Synthesis** (100-150ms): Multi-provider TTS with real-time voice cloning.

B. Detection Models

1) **Age Detection: Architecture:** CNN feature extractor + Transformer classifier. **Age Groups:** 18-25, 26-35, 36-45, 46-55, 56+ (5 groups). **Features:** MFCC (40 coefficients), pitch statistics (mean, variance, range), speaking rate (syllables/second), energy distribution. The model supports both 3-group and 5-group classification, with accuracy typically degrading as categories increase due to similarity of adjacent age groups.

2) **Gender Detection: Architecture:** Fine-tuned Wav2Vec 2.0 (960M parameters). **Classes:** Male, Female, Non-binary. **Features:** Fundamental frequency (F0), formant frequencies (F1, F2, F3), spectral tilt, vocal tract length estimation. The model supports both binary and inclusive classification, with binary classification achieving 90-95% accuracy.

3) **Accent Detection: Architecture:** Multi-task Wav2Vec 2.0 with shared encoder. **Spanish Regional Accents:** Castellano (Northern/Central), Andaluz (Southern), Catalan-accented, Gallego-accented, Vasco-accented, Canario, Murciano, Valenciano-accented. This task presents significant challenges due to subtle regional pronunciation differences, bilingual influence, and limited labeled datasets.

C. Voice Personalization

Voice Profile Library: 10 base voices covering age groups (Young, Adult, Middle, Older, Senior), genders (Male, Female per age group), and regional accents (Spanish Castellano base, adaptable). **Adjustable Parameters:** Pitch (0.85-1.2 $\times$), Speaking rate (0.85-1.15 $\times$), Energy (0.55-0.85), Warmth (0.5-0.8), Formality (0.5-0.8).

Adaptation Rules: Young (18-25): Higher pitch (+10%), faster rate (+10%), higher energy. Adult (26-35): Baseline parameters. Middle (36-45): Lower pitch (-5%), moderate formality. Older (46-55): Lower pitch (-10%), slower rate (-10%), higher warmth. Senior (56+): Lowest pitch (-15%), slowest rate (-15%), highest warmth.

III. RESULTS

A. Experimental Setup

1) **Datasets:** Experiments utilize three primary speech corpora: **VCTK Corpus** [?]: 110 speakers (48% male, 52% female), age 18-77, 44 hours, studio-quality, used for age and gender detection. **LibriSpeech** [?]: 2,458 speakers, 960 hours, mean age 42 years, diverse recording conditions, used for age detection training. **Common Voice Spanish** [?]: 1,850 speakers (51% male, 49% female), 320 hours, 8 regional accents, used for accent detection.

2) **Training Configuration:** All models trained on NVIDIA A100 GPUs $\times$ 4 with 80GB RAM. **Age Detection:** CNN+Transformer (6 layers, 8 heads, 512 hidden), 50 epochs, batch 32, lr 1e-3 with cosine annealing, AdamW optimizer. **Gender Detection:** Wav2Vec 2.0 (960M), 30 epochs, batch 16, lr 2e-5 with linear warmup, fine-tuning last 12 blocks. **Accent Detection:** Multi-task Wav2Vec (960M), 40 epochs, batch 12, lr 3e-5, loss weights: Age (0.25), Gender (0.35), Accent (0.40). All models use SpecAugment, noise addition (SNR 10-30dB), time stretching ($\pm$ 15%), pitch shifting ($\pm$ 2 semitones).

B. Classification Results

1) **Gender Detection:** In standalone evaluation, our fine-tuned Wav2Vec 2.0 gender detection model achieves **98.89% test accuracy** and 99.90% validation accuracy on LibriSpeech (251 speakers, 100 hours), representing state-of-the-art performance. The model surpasses prior baselines including pitch-based methods (85.2%), MFCC+GMM (88.7%), CNN approaches (93.1%), and standard Wav2Vec 2.0 (94.2%). Full per-class metrics are presented in our dedicated gender detection paper [?].

2) **Age Detection:** The age detection module uses a CNN feature extractor combined with a Transformer classifier, designed for 5-group classification (18-25, 26-35, 36-45, 46-55, 56+). The architecture is fully specified and implemented, but model training and evaluation have not yet been completed. Training is planned on VCTK and LibriSpeech corpora. Results will be reported in a future revision.

3) **Accent Detection:** The accent detection module employs a multi-task Wav2Vec 2.0 architecture with a shared encoder, targeting 8 Spanish regional accents (Castellano, Andaluz, Catalan-accented, Gallego-accented, Vasco-accented, Canario, Murciano, Valenciano-accented). This task presents significant challenges due to subtle regional pronunciation differences, bilingual influence, and limited labeled datasets. The architecture is designed and implemented, but model training and evaluation are in progress using Common Voice Spanish




Fig. 2: Component-wise latency breakdown showing 461ms mean end-to-end latency.

(1,850 speakers, 320 hours). Results will be reported upon completion.

C. Latency Analysis

End-to-end latency measured on NVIDIA A100 GPU with 10 concurrent calls (Figure 2): **Feature Extraction** 42ms, **Age Classification** 38ms, **Gender Classification** 35ms, **Accent Classification** 41ms, **Personalization Engine** 18ms, **TTS Synthesis** 287ms, **Total** 461ms (mean), 449ms (median), 582ms (95th percentile). Results: 82.3% of calls meet sub-500ms target. TTS synthesis accounts for 62.3% of total latency, representing the primary bottleneck. Optimizations (INT8 quantization + TTS caching) reduce latency to 312ms (32.3% improvement).

D. Discussion and Analysis

Multi-Task Learning Benefits: The multi-task architecture is designed so that shared representations improve generalization for low-resource tasks. Prior work [?] demonstrates that shared encoders can provide significant improvements for data-scarce tasks, which motivates our architectural choices for age and accent detection once training is complete.

Model Performance: Gender detection is production-ready with 98.89% test accuracy on LibriSpeech, representing state-of-the-art performance. Age and accent detection architectures are designed and implemented but await model training and evaluation.

Limitations: Age and accent models are not yet trained, limiting the current system to gender-based personalization only. Additional limitations include the binary nature of the verified gender evaluation, 8-region accent classification that may miss intra-regional variation, and business impact not yet validated through A/B testing.

Deployment Implications: The verified gender detection accuracy of 98.89% is sufficient for production deployment. Once age and accent models are trained, the full personalization pipeline can be validated through A/B testing. Latency performance supports conversational applications (461ms mean, 82.3% meeting sub-500ms target).

E. Ethical Implications and Mitigation

Positive Impacts: Increased accessibility for elderly users, cultural sensitivity through accent recognition, reduced algo-




Fig. 3: Model architectures for (a) Age Detection (CNN + Transformer), (b) Gender Detection (Wav2Vec 2.0), and (c) Accent Detection (Multi-task Wav2Vec).

gorithmic bias through inclusive design, improved user engagement and trust.

Risks and Mitigation: Potential demographic manipulation mitigated by explicit consent frameworks with granular opt-out controls. Privacy implications addressed through GDPR-compliant data minimization, DPIAs, and audit logs. Bias risks addressed through quarterly fairness audits with intersectional analysis and automated monitoring. Misuse prevention through platform-level moderation, transparency dashboards, and user control over personalization intensity.

Our mitigation strategies align with emerging 2026 standards: Soundverse Ethical AI Framework for consent and attribution [?], IA-Veu initiative for voice ethics [?], and C2PA requirements for audio watermarking [?].

IV. IMPLEMENTATION AND VALIDATION

A. Model Architectures

Figure 3 shows the three detection model architectures. Training configurations: Age (50 epochs, batch 32, lr 1e-3), Gender (30 epochs, batch 16, lr 2e-5), Accent (40 epochs, batch 12, lr 3e-5). All models support production deployment with modular TTS integration.

B. A/B Testing Design

Hypothesis: Personalized voice improves conversion rate and customer satisfaction. **Variants:** Control (Generic AI voice) vs Treatment (Personalized voice). **Sample Size:** 2,000+ calls per variant, 80% power, $\alpha = 0.05$. **Duration:** 6 weeks pilot across 5 industries (Insurance, Healthcare, Sales, Finance, Hospitality). **Metrics:** Conversion rate, CSAT, engagement duration, repeat interaction rate.

C. Current Status and Future Work

Fully implemented components include multi-task classification models, data preprocessing pipelines for VCTK/LibriSpeech/Common Voice, voice personalization engine with rule-based parameter adjustment, and modular

TTS integration layer. Validation through large-scale A/B testing remains as future work.

Future Directions: (1) Empirical validation through production A/B testing across industries, (2) Architecture improvements including efficient transformers for edge deployment and continual learning for new accents, (3) Cross-cultural validation across languages (Mandarin, Hindi, Arabic, Portuguese), (4) Enhanced personalization through emotional state detection and RLHF, (5) Ethical framework development with transparency mechanisms and governance protocols, (6) Longitudinal studies (6-12 months) to assess sustainability and novelty effects.

V. CONCLUSION

This paper presented **Laia**, the first comprehensive framework for adaptive voice personalization in production conversational AI systems, addressing the critical research gap between social psychology evidence on demographic similarity benefits and technical implementations for real-time voice adaptation.

Technical Contributions: Unified system architecture integrating multi-task speech classification with controllable TTS, achieving 461ms mean latency suitable for real-time conversation. Gender detection achieves 98.89% test accuracy (state-of-the-art on LibriSpeech), providing a production-ready foundation. Age and accent detection architectures are designed and implemented, with model training planned as immediate future work.

Implementation Framework: Production-ready components including data preprocessing pipelines, multi-task classification models (CNN+Transformer and Wav2Vec 2.0), hierarchical personalization engine with rule-based adjustment, and provider-agnostic TTS integration.

Ethical Foundation: GDPR-compliant data handling, bias mitigation across demographic subgroups, C2PA watermarking for transparency, IA·Veu-aligned consent frameworks, and user control mechanisms for granular personalization adjustment.

Validation Roadmap: Model training on curated datasets with comprehensive accuracy and fairness metrics, latency benchmarking demonstrating 82.3% meeting sub-500ms target, and large-scale A/B testing design (2,000+ calls per variant) to measure business impact.

Broader Impact: Framework establishes foundation for rigorous investigation of voice personalization, with potential benefits including increased accessibility, cultural sensitivity, and improved engagement. However, significant challenges remain including demographic manipulation risks, privacy implications, and potential for bias reinforcement.

We call upon the research community to (1) replicate and extend findings across diverse languages and cultures, (2) develop standardized benchmarks for demographic classification, (3) establish ethical guidelines for responsible deployment, and (4) conduct longitudinal studies to validate benefits and harms. Codebase, trained models, and experimental configurations are released as open-source to facilitate reproducibility.