




**TOKENIZED CURRICULUM LEARNING FOR
MULTI-EFFECT NEURAL RESTORATION OF
FULL-MIX MUSIC**
M. YUNG-HAO KHOO

HONORS COMPUTER ENGINEERING THESIS
SUBMITTED TO THE FACULTY OF DARTMOUTH COLLEGE
IN FULFILLMENT OF THE REQUIREMENTS FOR
THE HONORS DEGREE
OF
BACHELORS OF ENGINEERING

ADVISED BY: PETER CHIN, Ph.D.
JUNE 2025

ABSTRACT

Music mastering is traditionally the final and most nuanced phase of audio production, requiring expert engineers, costly software, and access to multitrack stems to refine tonal balance, dynamics, and spatialization. While recent deep learning tools have advanced stem-based workflows, no general-purpose system exists for high-quality enhancement of bounced, fully mixed audio—particularly amateur recordings affected by noise, distortion, or environmental artifacts.

This thesis introduces Token U-Net, a deep learning architecture for high-fidelity restoration of real-world, non-stem music. Our method operates in a discrete latent space: compressed token sequences produced by Meta AI’s EnCodec codec. Instead of modeling raw waveforms or spectrograms, the network learns token-level transformations that reverse degradations such as reverb, gain mismatch, compression artifacts, EQ imbalance, stereo collapse, and echo—without requiring source separation.

Paired training data are created from the Free Music Archive (FMA) medium dataset (25,000 tracks, 16 genres), with degradations procedurally applied using a parametric simulation engine spanning ten controllable variables. To improve convergence and generalization, the model is trained via curriculum learning—a five-stage degradation schedule of increasing difficulty—using a composite loss that combines token cross-entropy, time-domain L1, STFT and Mel losses, and auxiliary objectives for perceptual, dynamic, and spatial enhancement.

Token U-Net achieves consistent improvements across metrics including SNR, PESQ, STOI, and parameter recovery, outperforming baseline pipelines that rely on spectrogram models with Griffin-Lim or HiFi-GAN for reconstruction. In contrast, our model operates fully in the token space, avoiding lossy waveform synthesis

and yielding greater fidelity. Visualizations of token usage, spectral restoration, and training dynamics further demonstrate interpretability and robustness.

This work presents a comprehensive study of full-mix music restoration using discrete token representations and curriculum-aware, automated training. While not a replacement for professional mastering, the system offers a scalable framework for accessible music enhancement and lays groundwork for future research in token-based audio restoration.

PREFACE

ACKNOWLEDGEMENTS

I would first like to express my deepest gratitude to my thesis advisor and mentor in the engineering department, Dr. Peter Chin, for his unwavering support, encouragement, and faith in my ability to pursue such a challenging endeavor. I also thank Dr. Michael Casey and Professor Sang-Wook Nam for their invaluable guidance: Dr. Casey for his pioneering work at the intersection of music and AI, which continually inspired new directions, and Professor Nam for his years of hands-on instruction in music production and mastering, shaping both my technical fluency and artistic sensitivity.

I am equally grateful to the many professors whose lessons across disciplines shaped the thinking behind this project. Every course, whether directly related to this thesis or not, contributed to the skills and mindset that made it possible.

A heartfelt thank you to my classmates and friends in the Engineering Sciences and Music departments for their enthusiasm, feedback, collaboration, and moral support throughout this journey. Special thanks to Harry Leiter, my partner on our first deep learning and signal processing project, the technical foundation on which this thesis was built.

Most of all, I am deeply grateful to my family, Miya, Kris and Larry, for their love, patience, and steadfast belief in me. Their own achievements, past and present, inspired many of the long nights I devoted to this work. Their encouragement gave me the strength to chase something bigger: the unification of my technical and artistic passions into my second and final undergraduate thesis.

I see this work not as a conclusion but as a beginning—a foundation on which I hope to build a career at the intersection of signal processing, deep learning, and music research. Although I hope this thesis contributes to the growing field of music and artificial intelligence, it also marks the beginning of what I believe will be a lifelong journey of exploration, invention, and sound.

ON SOUND

My foundation in music was not built in conservatories or classrooms; it began on the bedroom floor, with a laptop, a MIDI keyboard, and a hunger to create. Before I studied signal processing, I was shaping soundscapes in Logic Pro and GarageBand, obsessing about balance, distortion, and feel. I grew up chasing sonics that I did not yet understand: layering, looping, and mixing without the vocabulary to explain what I was doing.

Years later, those experiments evolved into formal engineering. Under the mentorship of Sang-Wook Nam, I studied the anatomy of professional postproduction. VU metering, frequency masking, dynamic range sculpting, and began to see how subtle changes in signal flow could shape a listening experience. At Igloo Music, I worked on industry productions at a world-class level, but the spark remained the same: sound as an expressive, malleable material.

My research began with a question: could we train a machine to understand music not just as signal, but as intention? Could we give a small model at least the basics of the intuition of a seasoned engineer and let it recover clarity, not just fidelity?

This thesis reflects that journey—equal parts theory and emotion, mathematics, and memory. It is rooted in the belief that music is not merely content to be processed, but rather a structure to be understood, shaped, and respected.

CONTENTS

1	Introduction	1
1.1	Motivation	1
1.1.1	The Democratization of Music Creation	1
1.1.2	The Problem with Enhancement in the Wild	1
1.1.3	My Dual Lens: Engineering and Sound	2
1.1.4	Vision for Accessible Audio Intelligence	2
1.2	Problem Statement	3
1.3	Research Objectives	4
1.4	Thesis Contributions	5
1.5	Structure of the Thesis	6
2	Background & Related Work	8
2.1	What Is Music Mastering?	8
2.1.1	Historical Context	9
2.1.2	Mastering in the Modern Production Pipeline	10
2.1.3	The Role of Mastering Engineers	10
2.1.4	Mastering as a Restoration Process	12
2.2	Common Audio Degradations	12
2.2.1	Spectral Imbalance and Equalization Errors	13
2.2.2	Dynamic Range Compression Artifacts	13
2.2.3	Reverb Tails and Echo Mismanagement	14
2.2.4	Gain Staging Errors and Clipping	14
2.2.5	Stereo Imaging and Phase Incoherence	15
2.2.6	Stacked, Compound Degradations	16
2.2.7	Summary and Implications	18
2.3	Deep Learning in Audio Processing	18
2.3.1	Time-Domain vs. Frequency-Domain Representations	18
2.3.2	The Limitations of Spectrogram-Based Approaches	19
2.3.3	Survey of Recent Architectures	20
2.3.4	Domain-Specific Constraints for Music vs Speech	22
2.3.5	Implications for This Thesis	22
2.4	Speech vs. Music in Restoration Tasks	23
2.4.1	Signal Structure and Spectral Complexity	23
2.4.2	Failure Modes of Speech Models on Music	24
2.4.3	Visual Comparison of Spectral Energy	25
2.4.4	Perceptual vs. Objective Metrics	25
2.4.5	Toward Music-Specific Restoration Paradigms	26

2.5	Toward Token-Based Restoration and Curriculum Learning	27
2.5.1	Why Tokens?	28
2.5.2	Why Curriculum Learning?	28
3	Dataset Design & Audio Pair Generation	30
3.1	Dataset Selection and Preparation	30
3.1.1	Why FMA?	31
3.1.2	Why <code>fma_medium</code> ?	31
3.1.3	Track Properties and Preprocessing	32
3.1.4	Track Metadata and Usage	33
3.1.5	License and Reproducibility	33
3.1.6	Motivation for Paired Audio Learning	34
3.2	STFT Demastering: Simulating Audio Degradations	35
3.2.1	Equalization	36
3.2.2	Gain Mismatch	36
3.2.3	Echo	36
3.2.4	Reverb	37
3.2.5	Compression	37
3.2.6	Effect Combination Strategy	37
3.2.7	Examples of Degraded Outputs	38
3.3	Parameter Ranges and Realism	
	Justification	40
3.3.1	Overview of Randomly Applied Effects and Ranges	40
3.3.2	Motivating the Design of Each Range	40
3.3.3	Justification via Perceptual and Engineering Context	42
3.3.4	Illustration of Degradation Effects	43
3.3.5	Ensuring Robustness and Generalization	43
3.4	Preprocessing Tools and Libraries	44
3.4.1	Library Stack Overview	44
3.4.2	Waveform I/O and Normalization	45
3.4.3	Effect Application Pipeline	45
3.4.4	Mel-Spectrogram Generation	46
3.4.5	Dataset Pair Construction and Metadata Logging	46
3.4.6	Integration with Dataset Loader	47
3.4.7	Computational Considerations	48
3.4.8	Conclusion: A Scalable Degradation Framework	48
4	Early Approaches and Experimental Baselines	49
4.1	Spectrogram-Based Architectures	49
4.1.1	Motivation and Setup	49
4.1.2	U-Net Architectures: 1D vs. 2D	50
4.1.3	Explorations in Mel Band Resolution and Depth	50
4.1.4	Inference via Griffin–Lim Reconstruction	51

4.1.5	Key Failure Modes	51
4.1.6	Conclusion	52
4.2	Neural Vocoder Attempts	53
4.2.1	The Promise of HiFi-GAN	53
4.2.2	Initial Setup and Adaptation	54
4.2.3	Failure Modes Observed	54
4.2.4	Conclusion	56
4.3	Parameter Inversion Networks	56
4.3.1	Forecasting-Based Parameter Estimation	57
4.3.2	LSTM Architecture and Mathematical Formulation	58
4.3.3	Bridging Two Domains	58
4.3.4	Failure Mode and Illusion of Low Loss	59
4.3.5	Conclusion	59
4.4	Loss-Reality Gap: Why These Failed	60
4.4.1	Loss Functions Used	60
4.4.2	Sources of Discrepancy	61
4.4.3	Visual Examples of Failure	63
4.4.4	Beyond Conventional Losses	63
4.5	Lessons Learned: Toward	
	Token-Based Restoration	64
4.5.1	Why Existing Models Fell Short	65
4.5.2	Reframing the Problem	66
4.5.3	Conclusion and Outlook	67
5	Tokenization via Meta AI's	
	EnCodec	69
5.1	Why Tokens?	69
5.2	EnCodec Internals	71
5.2.1	Encoder and Decoder Pathways	72
5.2.2	Residual Vector Quantization (RVQ)	73
5.2.3	Bandwidth, Frame Rate, and Compression	73
5.2.4	Training Procedure and Discriminators	74
5.2.5	Summary Configuration	74
5.3	Precomputing Tokens for	
	Curriculum Training	75
5.3.1	Curriculum Degradation Stages	75
5.3.2	Comprehensive Feature Bundling in .pt Files	77
5.3.3	Speed Gains and Training Efficiency	78
5.3.4	SHA-256 Hashes and Reproducibility	78
5.3.5	Auxiliary Spectral Domains for Loss Functions	79
5.3.6	Error Handling and Fault Tolerance	80
5.4	.pt File Format and Metadata	80
5.4.1	Structure Overview	81
5.4.2	Audio Hashing for Provenance	82
5.4.3	Spectral Configuration Block	82

5.4.4	Metadata Dictionary	83
5.4.5	Example Entry	83
5.4.6	Training Integration via <code>TokenPairDataset</code>	84
5.5	Visualizing Token Usage	84
5.5.1	Histograms of Token Distributions	85
5.5.2	Token Heatmaps	85
5.5.3	Token Entropy and KL Divergence	86
5.5.4	Training Interpretability Roles	87
5.5.5	Justification for Token-Space Modeling	87

6 The Token U-Net

	Architecture	89
6.1	Overview	89
6.1.1	Motivating a Token-Level Architecture	89
6.1.2	End-to-End Objective	90
6.1.3	Input and Output Dimensions	91
6.1.4	Token Flow Through the U-Net	91
6.1.5	Skip Connection Gating and Token Alignment	92
6.1.6	Sequence Characteristics and Token Receptive Fields	93
6.1.7	Example Input/Output Flow	93
6.1.8	Transition to Component-Level Explanation	94
6.2	U-Net Architectural Components	95
6.2.1	Residual Blocks with GroupNorm and GELU	95
6.2.2	Convolutional Block Attention Module (CBAM)	96
6.2.3	Feature-wise Linear Modulation (FiLM)	98
6.2.4	Learnable Scalar Skip Gates	98
6.2.5	Temporal Context Block (Dilated Convolution)	100
6.2.6	Optional Bottleneck Layer	100
6.2.7	Summary of Shape Transformations	100
6.2.8	Primary Logit Prediction	101
6.2.9	Perceptual Spectrum Head	101
6.2.10	Gain Head	102
6.2.11	Compression Head (Dynamics Estimation)	102
6.2.12	Stereo Width and Phase Head	103
6.2.13	Mask Head (Token Drop Prediction)	103
6.2.14	Total Loss Function	104
6.2.15	Theoretical Justification for Multi-Effect Recovery	104
6.3	Comparison to Other U-Net	
	Variants	105
6.3.1	Deep Residual U-Net (5-Layer)	106
6.3.2	Token Residual U-Net (No Attention)	106
6.3.3	Attention-Only U-Nets	107
6.3.4	Residual + Naive Attention (No CBAM, No FiLM)	107
6.3.5	Two-Stage U-Net (Token + Spectral)	107
6.3.6	2D Spectrogram U-Nets	108

6.3.7	LSTM Regression Heads (Parameter Estimation)	108
6.3.8	Memory Footprint and Generalization Tradeoffs	108
6.3.9	Representation Differences: Spectrograms vs. tokens	109
6.3.10	Conclusion	109
6.4	Model Size, Parameter Counts, and Depth Variations	110
6.4.1	Final Parameter Count and Dimensional Flow	110
6.4.2	Why Four Layers Were Optimal	111
6.4.3	Parameter Growth with Base Dimension	112
6.4.4	Spectrogram vs. Token Runtime and Memory Footprint	112
6.4.5	Training and Batch Scaling Performance	113
6.4.6	Conclusion	113
6.5	Final Architecture Justification	114
6.5.1	Modular Components Recap: Designed for Multi-Effect Recovery	114
6.5.2	The Role of Auxiliary Heads	115
6.5.3	Temporal Context Block: Long-Range Dependency Modeling	115
6.5.4	Curriculum-Aware Modularity	116
6.5.5	Why This Architecture Succeeds	116
6.5.6	From Architecture to Optimization	117
7	Curriculum Training	
	Pipeline & Losses	118
7.1	Why Curriculum Learning?	118
7.1.1	Mathematical Justification: Progressive Distribution Shift	119
7.1.2	Empirical Failures of Flat Training	119
7.1.3	Stagewise Design: Linking Degradations to Learning Phases	120
7.1.4	Architectural Synergy with Curriculum Staging	122
7.1.5	Throughput Gains via Token Preprocessing	124
7.1.6	Conclusion: A Scaffold for Multi-Effect Restoration	125
7.2	Curriculum Stages & Loss Functions Defined	125
7.2.1	Token Cross-Entropy Loss	126
7.2.2	Mask Head Loss: Dereverberation Mask	127
7.2.3	Perceptual Head Loss: Quality Embedding	128
7.2.4	Gain Head Loss: Volume Restoration	129
7.2.5	Stereo Field Head Loss: Spatial Re-expansion	130
7.2.6	Compression Head Loss: Dynamics Restoration	131
7.2.7	Audio-Domain Losses	131
7.2.8	Combined Loss Function	132
7.2.9	Stage-Wise Activation	133
7.2.10	Stage 0: Identity Mapping ($\mathcal{D}_0$)	134
7.2.11	Stage 1: Single Effect ($\mathcal{D}_1$)	135
7.2.12	Stage 2: Double Effects ($\mathcal{D}_2$)	135
7.2.13	Stage 3 and 3_Stronger: Triple Effects	135
7.2.14	Stage 4 and 4_Stronger: Full FX Sampling	136
7.2.15	Token Bundle Format	137

7.2.16	Conclusion: A Curriculum of Gradual Supervision	137
7.3	Stage Advancement Logic	137
7.3.1	Validation Plateau Detection via EMA	138
7.3.2	Fallback Criterion: Training Loss Stagnation	139
7.3.3	Stage 0 Exception Handling	139
7.3.4	Dynamic Feature Activation Per Stage	139
7.3.5	Failure Recovery and Monitoring	140
7.3.6	Summary	141
7.4	Training Infrastructure	142
7.4.1	Hardware Environment	142
7.4.2	Mixed Precision with AMP (FP16)	143
7.4.3	Batch Adaptation and OOM Recovery	143
7.4.4	Gradient Clipping and Accumulation	144
7.4.5	Training Duration and Compute Cost	144
7.4.6	Debug Tools and Safeguards	145
7.4.7	Reproducibility and Version Control	145
7.5	Logging, Checkpoints, and Debug Tools	146
7.5.1	Epoch-Level Logging and Monitoring	146
7.5.2	Checkpointing for Resilience and Restart	147
7.5.3	Adaptive Batch Scaling and Gradient Accumulation	147
7.5.4	Audio Logging and Sanity Checks	148
7.5.5	Summary and Transition to Evaluation	149
8	Evaluation and Results	150
8.1	Quantitative Benchmarking	150
8.1.1	Evaluation Metrics	150
8.1.2	Metric Summary Tables	152
8.1.3	Discussion and Justification	153
8.2	Degradation Combination	
	Performance Across Models	153
8.2.1	Degradation Formulas and Parameters	153
8.2.2	Scoring Methodology	154
8.2.3	Results on 10 Multi-Effect Tracks	155
8.2.4	Interpretation and Comparative Analysis	155
8.3	Aggregate Performance Metrics and Effect-Specific Trends	157
8.3.1	Metric 1: Average Restoration Score	157
8.3.2	Metric 2: Effect-Specific Histograms	157
8.3.3	Metric 3: Standard Deviation and Robustness Index	158
8.4	Qualitative Listening Tests	159
8.5	Token Flow and Attention Visualization	160
8.6	Ablation Studies and Generalization Limits	162

9	Discussion of Strengths & Limitations	164
9.1	Strengths of the Token U-Net	164
9.1.1	Unified, Multi-Effect Restoration via Token Space	164
9.1.2	1.08 Billion Parameter Architecture with Attention & Conditioning	165
9.1.3	Curriculum-Aware Learning for Generalization	166
9.1.4	Token-Based Interpretability & Diagnostics	166
9.1.5	Robust Training Infrastructure	166
9.2	Limitations of the Current System	167
9.2.1	No Effect-Type Awareness	167
9.2.2	Lack of Stem or Source Separation	168
9.2.3	Token Quantization Bottlenecks	168
9.2.4	Overgeneralization and Failure on Edge Cases	168
9.3	Summary of Capabilities and Tradeoffs	169
10	Future Work	170
10.1	Genre-Aware Modeling	170
10.2	Stem-Aware Restoration	171
10.3	User-Controlled Enhancement	171
10.4	Prompt-Based and Transformer Remixing	172
10.5	Real-Time Deployment	173
11	Conclusion	175
11.1	Reflecting on the Problem and Contributions	175
11.2	Final Thoughts and Broader Impact	176
11.3	Toward the Future of Neural Audio	177
11.3.1	Bridging Creativity and Computation	177
11.3.2	Summary of Key Innovations	178
11.3.3	Positioning and Impact	178
11.3.4	Personal Reflection	179
A	List of Acronyms	180
B	Glossary of Terms	184
C	Extended Equations and Derivations	190
C.1	Short-Time Fourier Transform (STFT)	190
C.2	Mel Spectrogram	191
C.3	EnCodec Quantization and Tokenization	191
C.3.1	Vector Quantization	191

C.3.2	Multi-Codebook Tokenization	191
C.4	Griffin-Lim Reconstruction	192
C.5	HiFi-GAN Generator Losses	192
C.6	Token-UNet Output Shapes	193
C.6.1	Token Prediction (Cross-Entropy)	193
C.6.2	Auxiliary Heads	193
C.7	Perceptual and Token Losses	193
C.8	KL Divergence Between Token Distributions	194
C.9	Curriculum Advancement Condition	195
C.10	Entropy of Token Distributions	195
C.11	Quantized Latent Reconstruction Error	195
C.12	Token Coverage Ratio	196
C.13	Curriculum-Aware Loss Scheduling	196
C.14	Streaming Inference Constraints (Real-Time)	197
Bibliography		198

LIST OF TABLES

2.1	Representative Deep Learning Architectures for Audio	22
2.2	Comparative signal properties of speech and music.	23
3.1	Degradation parameters for synthetic audio impairment generation. .	41
5.1	EnCodec Configuration Used for Token U-Net Training	75
6.1	Comparison of U-Net variants across metrics (higher is better unless noted).	109
6.2	Parameter scaling with base dimension (depth fixed at 4)	112
6.3	Mapping of architectural components to target degradation classes. .	114
8.1	SNR, PESQ, and STOI on mild degradations (Stage 1)	152
8.2	ERLE and De-reverb Effectiveness (Echo/Room Tests)	152
8.3	Dynamic Range Compression (DRC) Recovery Comparison	152
8.4	Multi-Degradation Generalization (Strong Degradation Test)	152
8.5	Per-track restoration scores across models for compound degradations	155
8.6	Average restoration score across all tracks	157
8.7	Per-effect average restoration scores	158
8.8	Standard deviation and robustness index	159
8.9	Impact of architectural components on validation PESQ (Stage 4 Stronger)	162

LIST OF FIGURES

2.1	A schematic of Izhaki’s hybrid mastering chain [88]	11
2.2	Illustration of compounded degradations in non-professional mixes. Each processing stage introduces entanglement, increasing restoration complexity.	17
2.3	Mel spectrograms: Speech (left) shows narrow, localized formants; Music (right) displays sustained, broadband harmonic structure.	25
3.1	Waveform and mel-spectrogram examples from the FMA Medium dataset. Note the diversity in timbral structure, genre, and energy profile across tracks.	33
3.2	Example of degraded waveform with EQ boost and reverb tail; amplitude distribution is greatly affected. Top: original waveform. Bottom: modified waveform.	39
3.3	Mel spectrogram comparison. Degraded version exhibits unideal frequency distribution from EQ and smearing of high frequency bands.	39
3.4	Side-by-side visual comparison of effects applied with different parameters. Left column: clean audio. Right column: degraded versions. Top to bottom: EQ, reverb, compression.	43
4.1	Examples of low resolution output via Griffin-Lim Inference	52
4.2	Hifi-GAN model components [103]	53
4.3	Brunet’s proposed system	57
4.4	Example of low resolution output via Griffin-Lim Inference	63
4.5	Another of low resolution output via HiFi-GAN (Pre-Trained Neural Vocoder) Inference	63
5.1	Visualization of tokens from a waveform in the latent space [28]	69
5.2	Schematic of EnCodec’s architecture from [51], illustrating convolutional encoders, RVQ quantization, and decoder pathways. Tokens are derived from quantized latent embeddings at fixed frame rates.	71
5.3	Schematic of a Stage 3 Training pair and the curriculum stage progression	76
5.4	Visual sample of contents of feature bundling in .pt files	77
5.5	Visualexmaple of auxiliary STFT sum for token-spectrum alignment	79
5.6	Our pipeline to achieve the frozen lunchbox .pt token bundle	81
5.7	Token heatmaps for the same track before (left) and after (right) degradation. Repetition and banding in the degraded version signal token collapse and loss of structure.	86
6.1	Recap token vs. spectral advantage	90

6.2	The composite Convolution Block Attention Module) is a simple yet effect attention block for feedforward convolutional neural network as proposed by Woo et. al [191]	96
6.3	The channel attention block as discussed in <i>Attention is all you Need</i> Vaswani et. al [183] [191]	96
6.4	Inclusion of temporal/spatial attention as part of CBAM contributes greatly to the high-fidelity de-reverberation, dynamic range compression, and echo cancellation in support of the bottleneck temporal block [141, 114]	97
6.5	FiLM as proposed by Perez et. al assists in passing extra encoding information	98
6.6	U-Net architecture prior to integration of parameter inversion techniques, 741M vs. 1.08B parameters with integration of auxiliary heads and temporal block	99
6.7	A full diagram of Token U-Net architecture with temporal context block in bottleneck, CBAM, FiLM, auxiliary heads, etc.	105
7.1	Promising Loss Curves yet limited ability to invert reverb and echo degradations at inference	122
7.2	Left: Clean 3s signal (Right) Reverbed Signal. With just 3 seconds of music, we see the significant temporal and spectral alterations made.	124
7.3	Left: Degraded Signal from above right (Right) Restored signal after incorporating auxiliary losses.	124
7.4	Performance of CE loss graphed independently with all auxiliaries.	126
7.5	Mask Loss independently graphed.	127
7.6	Perceptual Head loss independently graphed.	128
7.7	Gain head loss graphed independently	129
7.8	One of our most vital losses graphed independently: stereo field	130
7.9	Compression loss graphed independently	131
7.10	All Audio Domain losses graphed. While not all employed the entire time in aggregate training loss, they were all independently tracked.	132
7.11	All auxiliary losses graphed, whether active or not	136
7.12	Evidence of 1cycle LR, Cosine Annealing OOM events, NaN events, Gradient Accumulation, and continuity due to automated training from fallback routines	142
8.1	Codebook token histogram comparison for Stage 2 vs Stage 4, while difficult to illustrate, has more utilization and distribution amongs codebooks	161

CHAPTER 1

INTRODUCTION

1.1 MOTIVATION

1.1.1 The Democratization of Music Creation

Over the past decade, the landscape of music production has undergone a profound transformation. Affordable digital tools and accessible platforms like TikTok, whose user base grew by approximately 13 million monthly active users in 2024, and SoundCloud, with over 140 million active users in 2025, have empowered countless bedroom producers and hobbyists to create and share music globally. Despite this increased accessibility, achieving a polished, professional sound remains a significant challenge. Mixing and especially finalizing audio for clarity, loudness, and playback consistency remains a black box for many. Most creators work with a final bounced mix, where all effects are flattened into a single stereo file, leaving no room for detailed post-production.

1.1.2 The Problem with Enhancement in the Wild

Conventional post-production workflows often depend on access to individual stems and the expertise of trained engineers using specialized hardware or software. However, the majority of modern creators only have access to a single stereo mix—often recorded in non-ideal conditions—and lack the tools or knowledge to perform nuanced

adjustments. Addressing issues such as muddiness, external noise, transient harshness, or dynamic imbalance becomes extremely difficult without access to isolated tracks. Commercial tools like iZotope Ozone [90] and LANDR [109] offer automated assistance, but are often opaque, assume clean input, and produce inconsistent results when applied to real-world, non-studio material. In addition, most deep learning research on audio enhancement has focused on speech [141], which lacks the complexity of the polyphonic structure of music, the variety of timbres and the aesthetics specific to the genre.

This thesis addresses a critical gap: *Can we design a system that restores and enhances fully mixed music audio, even when it is degraded and stemless, in a generalizable, accessible way?*

1.1.3 My Dual Lens: Engineering and Sound

As both a musician and an engineer, I have personally encountered the challenges of working with non-studio recordings and the frustrations of attempting to restore music captured under suboptimal conditions. This dual perspective drives the core of this thesis, blending practical insights from music production with advanced techniques in signal processing and machine learning [48]. This interdisciplinary experience highlights the need for models that are both perceptually effective and interpretable in real-world scenarios.

1.1.4 Vision for Accessible Audio Intelligence

If enhancement and restoration remain locked behind expert knowledge, expensive plugins, and access to clean stems, the creative potential of millions of musicians will remain limited. This work imagines a future where intelligent audio cleanup, handling things like echo, muddiness, or reverb, is available to any creator, regardless of environment or experience. Interpretability is reinforced through discrete token

modeling and auxiliary output heads that target specific effects such as gain correction, dynamic compression, and stereo field expansion. Our system, which combines token representations [51], curriculum learning [14], and perceptually aligned objectives, aims to move beyond hand-crafted signal chains and towards learned audio enhancement built for realism and accessibility.

1.2 PROBLEM STATEMENT

Although modern deep learning has advanced capabilities in speech enhancement [141, 117], dereverberation, and source separation, these systems typically assume idealized inputs: clean, isolated sources or narrow-bandwidth, monophonic speech. In the polyphonic, nonlinear world of music production, these assumptions often fail.

Research Problem

How can we restore and enhance fully mixed music audio - without access to stems, without expert supervision, and without manual DSP intervention - using a deep learning system that generalizes between genres, degradations, and real-world recording conditions?

Fully mixed music tracks, or “bounced” files, embed not just the instruments and vocals but also every EQ curve, compression setting, reverb trail, and gain adjustment applied during mixing. Once exported, these effects are irreversibly entangled in the final waveform. Unlike speech, music is diverse in genre, instrumentation, tempo, and timbre [166], making the problem more difficult and much less studied when approached without source separation.

Compounding Challenges:

1. **No ground truth stems:** Most users have no access to isolated instrument or vocal tracks.

2. **Entangled effects:** Artifacts are non-linear and inseparable, invalidating additive noise models.
3. **Non-professional conditions:** Audio may be clipped, distorted, or captured on consumer-grade equipment.
4. **No definitive target:** Multiple “clean” versions may be equally perceptually valid, complicating supervision.

Most importantly, restoration must be *feasible*: it must operate on accessible hardware, provide interpretable outputs, and outperform existing approaches in perceptual quality and objective metrics [161, 179].

The Research Gap

While speech enhancement and stem-based mastering have received substantial attention, **there is no state-of-the-art system specifically designed to restore fully mixed, stemless music recordings under real-world conditions [158]**. Most related work - including DeepVQE [84], Moliner et al. [133], and Li et al. [117]—target isolated effects or speech; none operate together in all five effects on degraded music without stems.

1.3 RESEARCH OBJECTIVES

The goal of this thesis is to develop an efficient and interpretable deep learning system for enhancing fully mixed music audio, especially recordings created by non-professionals without access to studio tools, stems, or prior cleanup. This work spans signal processing, machine learning, and perceptual audio engineering.

The objectives are as follows:

1. **Build a Deep Learning Model for Full-Mix Enhancement:** Design and

train a model that operates directly on stereo, full-mix audio to reduce artifacts, improve clarity, and simulate mastering-style improvements such as EQ balance, dynamic smoothing, reverb cleanup, and stereo widening [105].

2. **Learn from Artificially Degraded Data:** Construct a synthetic dataset by systematically degrading clean tracks with a range of audio effects [134, 163].
3. **Use Compact, Token-Based Audio Representations:** Replace waveform or spectrogram input with discrete compressed token sequences using Meta AI’s EnCodec [28].
4. **Implement Curriculum-Aware Training with Multi-Objective Losses:** Progressively increase restoration difficulty using a five-stage schedule with diverse losses that reflect both perceptual and structural restoration goals [22].

Final Aim

This work seeks to demonstrate that perceptually aligned, multi-effect restoration is achievable for real-world, fully mixed music—without stems, handcrafted DSP, or expert oversight.

1.4 THESIS CONTRIBUTIONS

This thesis introduces a new deep learning framework for full-mix music restoration that challenges prevailing assumptions in both speech and music signal processing. Key contributions include:

- **Reframing mastering as restoration:** This work recasts mastering-style enhancement as a learned restoration problem rather than rule-based processing [48, 22].
- **Multi-effect prediction for compound degradations:** The model uses

specialized output heads for gain, stereo width, dynamic compression, reverb masking, and perceptual embedding alignment.

- **Discrete latent modeling:** EnCodec token sequences eliminate the need for phase reconstruction and support lossless audio compression in training and inference [28].
- **Curriculum-aware learning:** A five-stage training protocol gradually increases task difficulty, stabilizing optimization and improving generalization [133].
- **Focus on accessibility and interpretability:** The system prioritizes independent creators and practical usability over studio constraints or handcrafted heuristics [17].

1.5 STRUCTURE OF THE THESIS

This thesis is organized as follows:

- **Chapter 1: Introduction** describes motivation, problem statement, research goals, and contributions.
- **Chapter 2: Background and Related Work** situates this project within the broader fields of music production, audio restoration, and deep learning.
- **Chapter 3: Dataset Design & Audio Pair Generation** discusses our choice of dataset, the data preparation pipeline, and parameter degradation processing logistics
- **Chapter 4: Early Approaches and Experimental Baselines** discusses the preliminary experiments and failures in spectral spaces and with neural vocoders, where inspiration had been derived from applications not directly applicable to the full-mix music problem

- **Chapter 5: Tokenization via EnCodec** discusses our shift to the latent token space and a major reduction in computation power consumption through precomputing all necessary data in a compressed form
- **Chapter 6: The Token U-Net Architecture** focuses in on the central focus of this thesis, the architecture and all of its components, supplied with their mathematical foundation. Each has a purpose and is justified clearly with ample references to state of the art papers
- **Chapter 7: Curriculum Training Pipeline & Losses** walks through the process by which we were able to make the model learn all possible degradations, and how we were able to automate training across diverse degradation scenarios.
- **Chapter 8: Evaluation & Results** provides a multitude of metrics and comparisons by industry and Token U-Net specific standards
- **Chapter 9: Discussion of Strengths & Limitations** discusses the extent of our model's capabilities
- **Chapter 10: Future Work** proposes further ideas and implementations that arose from the research around Token U-Net
- **Chapter 11: Conclusion** reflects on this project and re-highlights the key contributions

The next chapter provides the historical and technical foundation for this work, including an overview of music enhancement, types of degradation, and related systems in deep learning for audio restoration.

CHAPTER 2

BACKGROUND & RELATED WORK

2.1 WHAT IS MUSIC MASTERING?

Music mastering is the final and most nuanced phase of the audio production pipeline [17], traditionally responsible for shaping the overall sonic character and optimizing playback readiness [145]. Unlike mixing, which balances individual stems, mastering operates on a fully mixed stereo file, applying holistic signal processing to ensure tonal balance, dynamic consistency, and translation across all listening environments [88]. A well-mastered track should maintain clarity and presence, whether played on earbuds, car stereos, or reference monitors [46].

Mastering spans both technical compliance and creative refinement: it ensures adherence to loudness and delivery specifications (e.g., LUFS normalization for streaming platforms) while preserving or enhancing musical intent through subtle signal manipulation [71]. Engineers typically employ equalization (EQ), multiband compression, stereo field shaping, limiting, harmonic enhancement, and dynamic control [124] to finalize the audio.

2.1.1 Historical Context

Mastering originated in the analog era, when engineers prepared mixes for physical media such as vinyl or cassette. This process involved compensatory EQ to accommodate vinyl cutting limitations (e.g., minimizing excessive low-end that could dislodge a stylus), dynamic control to prevent distortion, and mono compatibility for radio broadcast.

Engineers used iconic analog gear such as the Manley Massive Passive EQ, Fairchild 670 compressor, SSL G-Series bus compressors, and Studer A800 tape machines—devices prized not only for their technical function but also for their sonic coloration and nonlinearity [48]. These tools introduced harmonics, phase shifts, and compression artifacts that became part of the “analog sound.”

With the advent of digital audio workstations (DAWs) in the late 1990s, software-based mastering became widespread. Tools such as Pro Tools, Logic Pro, and Ableton Live democratized access to professional workflows [139]. Plugin suites such as iZotope Ozone, Waves L2, FabFilter Pro-Q, and Slate Digital emulated analog chains with digital precision. Yet even with such tools available, the human expertise behind mastering remains critical. Many professionals now use hybrid chains—combining analog gear with digital precision—to blend musical character and control.

Historical Context Takeaway

The shift from analog to DAW-based mastering reduced the barrier to tools but not the barrier to knowledge. As a result, many creators now have access to powerful plugins but lack the signal processing, mixing, and mastering fluency to use them effectively.

2.1.2 Mastering in the Modern Production Pipeline

A typical mastering session involves multiple stages of enhancement and correction, each contributing to the final sonic polish:

1. **Broadband Equalization:** Adjusting the overall tonal balance, often using high-shelf boosts ($\approx 10\text{kHz}$) to add air or cuts $\approx 3\text{--}5\text{kHz}$ to reduce harshness [131].
2. **Dynamic Range Control:** Applying multiband compression to smooth peaks across frequency bands while preserving transients.
3. **Stereo Imaging:** Using mid-side EQ or stereo widening to optimize spatial perception without introducing phase issues [151].
4. **Harmonic Enhancement:** Introducing warmth or saturation through analog emulation or harmonic exciters.
5. **Limiting and Clipping:** Raising perceived loudness using brickwall limiters or soft clippers while avoiding digital distortion.
6. **Metering and Platform Compliance:** Ensuring final delivery meets platform standards [109], such as:
 - **Spotify:** -14 LUFS, -1.0 dBTP
 - **Apple Music:** -16 LUFS, -1.0 dBTP
 - **YouTube:** -13 LUFS, -1.0 dBTP

2.1.3 The Role of Mastering Engineers

Despite increasing automation, mastering remains a craft rooted in critical listening. Renowned engineers like Emily Lazar, Mandy Parnell, and Bernie Grundman are


Digital Audio Mastering FlowChart

Figure 2.1: A schematic of Izhaki’s hybrid mastering chain [88]

valued for their ability to retain musical intent while subtly enhancing loudness and clarity. They operate in acoustically tuned environments with transparent monitors, allowing for perceptual judgments that go beyond what meters can show.

However, for most independent musicians, such resources are out of reach. This raises an important research question: can a model learn to understand and possibly simulate these perceptual enhancements using only a degraded stereo mix?

Key Insight

Mastering is not simply about applying signal chains—it is the human perception of musical clarity, loudness, and spatial intent that defines professional-grade audio. Any neural system must therefore model both transformation and intention.

2.1.4 Mastering as a Restoration Process

Mastering is often conflated with enhancement, but in many real-world cases—especially for amateur or archival recordings—it serves as a form of restoration. Tracks may arrive with clipped transients, poor spectral balance, over-compression, noise, or unintended reverberation. Engineers address these problems using tools like transient shapers, de-essers, de-clippers, and surgical EQ [87, 118].

In this thesis, we adopt this restorative view. We argue that mastering-style enhancement can be reframed as a high-fidelity restoration task: one that seeks to infer the perceptual and spectral clarity a human engineer might aim for, without requiring stem access or handcrafted DSP rules.

Thesis Framing

We define music restoration as the inverse problem of recovering clarity, dynamic range, and spectral balance from entangled, fully mixed audio—approximating the corrective and perceptual actions of a mastering engineer without explicit stem information.

2.2 COMMON AUDIO DEGRADATIONS

Real-world music recordings—especially those created outside of professional studios—often contain a wide range of sonic degradations. These artifacts can stem from poor mixing practices, inadequate room acoustics, consumer-grade hardware, or a lack of post-production expertise [136]. Unlike synthetic degradations used in speech enhancement (typically additive, stationary, and statistically separable), music degradations tend to be nonlinear, temporally entangled, and context-sensitive [38].

2.2.1 Spectral Imbalance and Equalization Errors

Equalization (EQ) shapes the frequency response of a track. When misused, it can introduce tonal imbalances that make mixes translate poorly across playback systems.

Examples include:

- *Muddiness*: Excess buildup around 200–500 Hz due to overlapping sources (e.g., kick, bass, and piano).
- *Harshness*: Overemphasized bands near 2–4 kHz to push vocals forward, often causing fatigue.
- *Over-darkening*: Broad cuts above 10 kHz, sometimes introduced by amateur de-essing or poor mic response.

Such errors may arise from genre misalignment, monitoring inaccuracies, or a misunderstanding of how frequency masking works across instruments.

2.2.2 Dynamic Range Compression Artifacts

Compression is essential for loudness control and punch [22], but misuse—especially in post-mix workflows—can result in flattened dynamics and unnatural behavior [132].

Let $x[n]$ be the input signal and $y[n]$ the compressed output. A soft-knee compressor may apply:

$$y[n] = x[n] \cdot \left(1 - \alpha \cdot \tanh \left(\frac{x[n]^2}{T^2} \right) \right)$$

where α controls compression strength and T is the threshold.

Improper settings lead to:

- *Transient smearing*: Fast attack times suppress peaks (e.g., snare hits), reducing impact.

- *Pumping*: Slow release times drag energy across phrases.
- *Ducking*: Poor sidechain configuration unintentionally suppresses key elements.

Compression misuse is especially detrimental in genres like hip-hop and EDM where groove and energy are heavily reliant on tight dynamics.

2.2.3 Reverb Tails and Echo Mismanagement

Artificial reverb and delay simulate space and depth, but excessive use can obscure details [4]; we will investigate the magnitude of such effects through the spectral domain in later chapters. In bounced stereo tracks, these effects are convolved, and thus irreversibly imprinted:

$$y[n] = x[n] * h[n] = \sum_{m=0}^M x[m]h[n-m]$$

where $h[n]$ is the room or reverb impulse response.

Common issues [130]:

- *Overlong decay*: Reverb tails muddy dense, highly-detailed arrangements.
- *High-frequency wash*: Poor damping produces spectral build-up.
- *Echo phase clashes*: Poorly time delays clash rhythmically with the content.

These are frequently committed during production and rendered into the file, leaving no clean reference for post-processing.

2.2.4 Gain Staging Errors and Clipping

Improper gain staging results in loudness mismatch or hard clipping. Without sufficient headroom, summing across tracks may exceed 0 dBFS [47]:

$$\text{clip}(x[n]) = \begin{cases} T, & x[n] > T \\ x[n], & |x[n]| \leq T \\ -T, & x[n] < -T \end{cases}$$

This introduces aliasing and nonlinear distortion that is often irreversible [167].

Typical issues:

- *Peak clipping*: Causes brittle transients and IMD (intermodulation distortion).
- *Gain drift*: Imbalanced sections (e.g., verse too soft relative to chorus).
- *LUFS noncompliance*: Tracks are rejected or down-leveled by streaming platforms.

2.2.5 Stereo Imaging and Phase Incoherence

Stereo imaging defines the spatial experience of a mix. Poor imaging, especially phase-based widening or improper panning, can create mono-incompatibility and listener fatigue [151, 62].

Common degradations:

- *Pseudo-width*: Delay-based widening causes phase cancellation when summed to mono.
- *Balance skew*: Left- or right-dominant mixes from uneven panning.
- *Decorrelated reverb*: Spatial reverb artifacts that collapse or shimmer unpredictably.

Stereo field mismanagement is difficult to detect without specialized meters, which is absent from most amateur workflows.

2.2.6 Stacked, Compound Degradations

Real-world music often exhibits multiple interacting degradations. Unlike additive noise in speech, musical artifacts cascade in non-linear stages [57]. A poorly EQ'd, compressed, and reverberated vocal track baked into a stereo mix cannot be salvaged with a single-effect fix.




Figure 2.2: Illustration of compounded degradations in non-professional mixes. Each processing stage introduces entanglement, increasing restoration complexity.

2.2.7 Summary and Implications

Implication for Enhancement Systems

Music degradation is not simply additive noise; it is a complex, non-linear, genre-sensitive entanglement. Restoration systems must address multiple effects jointly, rather than in isolation, requiring new representations and learning frameworks.

Full-mix music restoration cannot rely on assumptions valid for speech enhancement. This motivates architectural innovations such as tokenized audio modeling, perceptual objectives, and curriculum-based training, explored in later chapters.

2.3 DEEP LEARNING IN AUDIO PROCESSING

The past decade has witnessed a seismic shift in audio signal processing with the advent of deep learning [14]. Moving beyond handcrafted filters and DSP heuristics, modern neural architectures now learn hierarchical representations of sound directly from raw or transformed data [178]. From source separation and speech enhancement to synthesis and mastering, these models span a wide range of tasks and modalities. However, most models have been trained and evaluated under assumptions specific to speech or stem-isolated music, making their adaptation to fully mixed audio a significant research challenge.

2.3.1 Time-Domain vs. Frequency-Domain Representations

Audio signals are inherently temporal, but frequency analysis offers interpretability and structure, leading to two dominant representation paradigms: time-domain and frequency-domain models [165].

Time-domain models such as WaveNet [100] and Wave-U-Net [174] operate directly on raw waveforms $x[n]$. WaveNet, an autoregressive model, demonstrated that deep stacks of dilated causal convolutions could generate high-quality speech audio sample by sample. However, its high computational cost limits the practicality of inference. Wave-U-Net, in contrast, uses an encoder-decoder structure with 1D convolutions to estimate target waveforms (e.g., isolated vocals or denoised outputs) from mixed inputs.

Frequency-domain models instead transform the waveform via Short-Time Fourier Transform (STFT) [66]:

$$X(k, m) = \sum_{n=0}^{N-1} x[n + mH] \cdot w[n] \cdot e^{-j2\pi kn/N}$$

where N is the FFT window size, H is the hop size, $w[n]$ is a window function and (k, m) indexes frequency and time, respectively. These models typically operate on magnitude spectrograms, discarding phase entirely or estimating it through algorithms like Griffin–Lim [66].

Although frequency domain models are easier to train due to their compactness and perceptual salience (e.g., formant bands, harmonics), they suffer from phase loss, spatial collapse, and spectral smoothing, especially problematic for polyphonic music. [129].

2.3.2 The Limitations of Spectrogram-Based Approaches

Despite their prevalence, spectrogram-based models introduce three major bottlenecks in music restoration:

1. **Phase Recovery:** Griffin–Lim reconstruction often introduces swirling artifacts or temporal smearing. Even advanced vocoders like HiFi-GAN [103] can-

not guarantee phase-coherent outputs in long musical passages.

2. **Bandwidth Bottlenecks:** Many models operate on 16kHz or 22.05kHz audio, sufficient for speech but insufficient for the 20Hz–20kHz fidelity range needed in music [51].
3. **Stereo Collapse:** 2D convolutional operations applied on magnitude spectrograms typically ignore spatial cues and stereo phase alignment—critical in modern music mastering [145].

Spectrogram Tradeoffs

While mel-spectrograms are efficient for training and capture perceptual features, their use requires separate models such as pre-trained vocoders for inversion and results in poor stereo reproduction and phase distortion in full-mix music.

2.3.3 Survey of Recent Architectures

We now examine key architectures that have shaped the field [165]:

U-Net. A foundational architecture for image and audio tasks, U-Net features symmetric encoder-decoder paths with skip connections. It was adapted by Jansson et al. and others for music source separation (e.g., Spleeter, Demucs), and later used in audio enhancement tasks like speech and tape denoising [92, 52, 134]. Although effective, most standard U-Nets operate on spectrogram magnitudes and lack phase handling.

Wave-U-Net. An adaptation of U-Net to the raw waveform domain, introduced by Stoller et al., this model replaces 2D convolutions with 1D, capturing temporal dependencies directly [174]. It demonstrated that raw audio modeling is viable for

tasks such as stem separation and drum mixing, albeit at the cost of increased memory usage [105].

HiFi-GAN. A generative adversarial network trained to synthesize waveform audio from mel spectrograms. It outperforms predecessors like MelGAN [107] and WaveGlow [153] in both fidelity and speed. However, HiFi-GAN was trained predominantly in speech and does not generalize to complex music textures, often producing harmonic artifacts, collapsed stereo, or incorrect transients when fed music mel spectrograms not distributed.

DDSP (Differentiable Digital Signal Processing). Proposed by Google Magenta, DDSP integrates traditional DSP elements like filters, oscillators, and reverb into a differentiable graph. It achieves fine control in instrument synthesis tasks, but assumes clean input signals, which limits the utility for restoration [55].

Diffusion Models. Recent approaches such as DiffWave [104] and Noise2Music [82] employ denoising diffusion probabilistic models (DDPMs) to model waveform distributions via iterative refinement. Although extremely powerful, these models are often too slow or unstable for real-time use and require extensive computation.

Transformers and Token Models. High-capacity sequence models such as AudioLM [20], SoundStorm [21], MusicGen [36], and EnCodec-based [28] systems have pushed the frontier of audio token modeling. These models embed audio as discrete codebook sequences and generate long-form music using self-attention [183]. While transformative in generation tasks, they have yet to be adapted at scale to full-mix audio enhancement or restoration.

Table 2.1: Representative Deep Learning Architectures for Audio

Model	Domain	Task	Limitation
U-Net	Frequency (STFT)	Separation, denoising	No phase modeling
Wave-U-Net	Time (waveform)	Mixing, source sep.	Memory-intensive
HiFi-GAN	Spectrogram $\rightarrow$ waveform	Vocoding	Music hallucination
DDSP	Hybrid (waveform + parameters)	Synthesis	Needs clean signals
AudioLM	Token-based (EnCodec)	Generation	Not used for enhancement
Demucs	Hybrid (waveform + STFT)	Source separation	Needs stems

2.3.4 Domain-Specific Constraints for Music vs Speech

Speech-focused models are often evaluated on intelligibility metrics like PESQ and STOI, trained on narrowband signals (e.g., 8-16kHz) and assume quasi-stationary single-source content [161]. In contrast, music is polyphonic, broadband, and exhibits a hierarchical structure, containing motifs, structures, time signatures, a spectrum of timbres, etc. [135]. Therefore, models that excel at speech denoising such as DeepVQE [84] and DPCRNN [114] perform quite poorly when applied to music.

2.3.5 Implications for This Thesis

The failures of conventional deep learning architectures on full-mix music motivate a new direction: operating on compressed, latent token representations like EnCodec, allowing for phase-aware, stereo-preserving, and memory-efficient training. The system proposed in this thesis builds on this insight, using curriculum training on EnCodec token pairs to handle a wide range of audio degradations in music mastering rather than focusing on one particular effect.

Architectural Implication

In order to handle real-world, bounced music inputs in deep learning, models must preserve stereo, operate across full bandwidth, tolerate high dynamic range, and model non-stationary harmonic structure, all while remaining computationally tractable.

2.4 SPEECH VS. MUSIC IN RESTORATION TASKS

Restoring speech and music audio involves fundamentally different challenges due to divergent signal characteristics, perceptual goals, and degradation profiles. While deep learning has advanced speech enhancement—addressing dereverberation, denoising, and acoustic echo cancellation—these same models often degrade musical quality when applied indiscriminately. This section outlines structural differences, explains model failure modes, and motivates the need for music-specific restoration strategies.

2.4.1 Signal Structure and Spectral Complexity

Table 2.2 summarizes key contrasts between speech and music:

Table 2.2: Comparative signal properties of speech and music.

Property	Speech	Music
Bandwidth	300 Hz – 4 kHz	20 Hz – 20 kHz
Spectral Density	Sparse, formant-driven	Dense, harmonic, overlapping
Temporal Texture	Monophonic or dialogic	Polyphonic, rhythmic, hierarchical
Phase Importance	Secondary to intelligibility	Crucial for tone and stereo realism
Structure	Predictable (phonemes)	Syncopated, layered, genre-specific
Eval Metrics	Intelligibility (STOI, PESQ)	Fidelity, musicality, expressiveness

Speech signals are typically narrowband, formant-based, and temporally sparse. Music, by contrast, spans the full human auditory range, exhibits rich harmonic layering, and encodes intent in timbre, space, and timing. These distinctions render speech-optimized models brittle when applied to music.

2.4.2 Failure Modes of Speech Models on Music

Most speech systems assume additive noise:

$$y[n] = x[n] + n[n]$$

and employ magnitude masking in the STFT domain:

$$|\hat{X}(k, m)| = M(k, m) \cdot |Y(k, m)|$$

where $M(k, m)$ is a learned mask and $|Y(k, m)|$ is the noisy input magnitude spectrogram.

This works well for speech due to its sparse formant structure and short reverberation tails [140]. However, musical degradations are non-linear, aesthetic (e.g. intentional reverb, clipping, or saturation), and entangled. The application of a learned magnitude mask can remove musically meaningful content, suppress transients, or destroy reverb trails—leading to perceptual degradation [81].

In addition, models such as DeepVQE [84], SEGAN [147], and DPCRN [110] prioritize intelligibility and often approximate or ignore phase - acceptable for speech, but catastrophic for stereo and tonal realism in music [62]. This leads to:

- **Stereo collapse:** Interchannel phase cues are lost or distorted.
- **Harmonic aliasing:** Over-suppression leads to unnatural harmonic content.
- **Transient smearing:** Compression artifacts or masking smooth musical attacks.
- **Reverb destruction:** Natural decay tails are removed as if “noise.”

2.4.3 Visual Comparison of Spectral Energy

Figure 2.3 compares mel spectrograms of clean speech and music.




Figure 2.3: Mel spectrograms: Speech (left) shows narrow, localized formants; Music (right) displays sustained, broadband harmonic structure.

Speech signals are temporally sparse and spectrally constrained. Music exhibits continuous harmonic content, broadband energy, and high modulation density. These properties are misaligned with the assumptions built into most speech enhancement pipelines.

2.4.4 Perceptual vs. Objective Metrics

Speech systems are often evaluated using intelligibility or distortion metrics: PESQ, STOI, SI-SDR [161]. These do not capture subjective audio quality or musical expressiveness.

Music restoration requires perceptual evaluation focused on:

- **Timbre preservation:** Accurate instrument/vocal identity.
- **Spatial integrity:** Stereo localization, phase realism.
- **Dynamic contour:** Avoiding over-compression or unnatural leveling.
- **Genre-sensitive fidelity:** Expectations differ across styles.

Standardized metrics like PEAQ [179], ViSQOL [78], or subjective listening tests are more appropriate but underused in restoration research.

2.4.5 Toward Music-Specific Restoration Paradigms

Given these distinctions, speech-centric pipelines are ill-suited for music. Full-mix music restoration requires:

1. **Phase-aware processing:** To preserve stereo imaging and transients.
2. **Capacity for compound degradations:** EQ, compression, reverb, etc.
3. **Aesthetic alignment:** Training objectives that reflect musical perception.
4. **Robustness across styles:** From ambient to metal to classical.

This thesis addresses these needs by operating in a tokenized audio space using EnCodec [51], enabling discrete, stereo-preserving representations. Unlike spectrogram masking or waveform regression, the proposed Token-UNet directly learns to map degraded token sequences to enhanced ones. Chapter 6 details the architecture and motivation.

Preview

Speech enhancement removes noise. Music restoration preserves emotion. This thesis reframes enhancement as a perceptual fidelity task—leveraging tokens, curriculum learning, and multi-effect supervision to recover clarity from stemless, degraded music recordings.

2.5 TOWARD TOKEN-BASED RESTORATION AND CURRICULUM LEARNING

Despite decades of work on automated mixing and mastering, current commercial and academic systems remain limited in transparency, flexibility, and generalizability. Many tools (e.g., iZotope Ozone, LANDR, CloudBounce) offer parametric mastering via preset-based heuristics or proprietary black-box models, with limited access to intermediate states or degraded-clean pairs. Although such systems are effective for basic loudness normalization or tonal adjustments, they rarely support nuanced correction between interacting degradations or genre-specific mastering intent [172, 105].

Academic approaches have sought to model specific mastering tasks, such as equalization prediction [132], drum mixing [157], or tape denoising [133], typically using supervised learning on synthetic data. However, these models operate on either individual stems or full waveforms with narrow effect targets, and few can address real-world compound degradations within bounced stereo audio.

Gap in Prior Work

No existing system offers fine-grained, generalizable mastering-style restoration on fully mixed audio, with support for stereo preservation, compound degradations, and perceptually coherent enhancement.

To address this, we propose a shift from parametric enhancement to perceptual restoration: treating the mastering process as a learned inverse problem, not of recovering stems, but of restoring fidelity and clarity from an entangled stereo mix. This requires models capable of:

- Operating on compact but expressive audio representations;

- Handling phase-aware stereo output;
- Managing multiple degradation types simultaneously;
- Learning from progressively more complex data.

2.5.1 Why Tokens?

Token-based modeling, popularized in music generation [21, 36], enables high-fidelity, discrete representations using quantized audio codecs such as EnCodec [51]. These systems convert waveform audio into latent codebook tokens, which preserve phase, stereo, and bandwidth without the overhead of raw waveform modeling. In contrast to spectrograms, these representations are more amenable to sequence learning, allow for direct comparison in discrete space, and can be decoded with high fidelity.

2.5.2 Why Curriculum Learning?

The complexity of full-mix audio restoration spans a broad spectrum, from simple EQ mismatches to non-linear, compounding degradations across time and frequency. A fixed-difficulty training scheme risks instability or overfitting. Instead, we adopt curriculum learning [14], introducing degradations in stages: from identity mappings to stacked non-linear distortions. This structured progression stabilizes training, enables representation reuse across conditions, and helps the model develop generalization without catastrophic forgetting.

Chapter Transition

The remainder of this thesis presents our system for token-level mastering restoration: a Token-UNet trained over curriculum-based degradation pairs. Chapter 3 outlines the dataset design, token pipeline, and degradation strategy that underpin this architecture.

CHAPTER 3

DATASET DESIGN & AUDIO PAIR GENERATION

While this thesis is rooted in the domain of music mastering, the methodology and system design are fundamentally reframed as a restoration task. In contrast to mastering, which assumes high-quality stems and creative intent, our approach treats real-world bounced mixes as degraded signals requiring not just enhancement, but generalizable restoration strategies that can recover intent from flawed or entangled input signals. This shift allows us to model not only loudness and balance, but also the removal of unintended artifacts, poor processing, and stylistic mismatches, especially in nonprofessional contexts.

3.1 DATASET SELECTION AND PREPARATION

A cornerstone of any machine learning pipeline is the quality and relevance of the training data. In audio enhancement—particularly for the mastering and restoration of non-stem real-world musical tracks—most publicly available datasets are either (a) designed for speech processing, (b) curated under studio-grade stem conditions, or (c) do not reflect the entangled degradation patterns observed in user-generated content [136]. To bridge this gap, we adopted the **Free Music Archive (FMA)** dataset, and specifically, the `fma_medium` subset [40].

3.1.1 Why FMA?

The Free Music Archive (FMA) is a Creative Commons-licensed corpus of more than 100,000 tracks contributed by independent artists and netlabels. Unlike data sets such as MUSDB18 [155] or NSynth [126], which focus on studio-grade multitrack or synthetic content, FMA contains naturally produced or recorded *bounced* stereo tracks of music, but often lacks professional mastering [39]. This makes it a rare fit for studying enhancement techniques in the context of realistic (fully mixed or even casually recorded) content.

In addition, FMA’s diversity spans genres, instrumentation, production styles, and loudness profiles. Many tracks are not even traditional songs in the commercial sense: Some are field recordings, instrumental etudes, lo-fi demos, or generative sound experiments. This heterogeneity supports the thesis goal: to build a model that generalizes to low-resource, highly varied music content.

Key Justification

Most enhancement systems assume stem-level access or work exclusively on speech. FMA provides stereo waveforms that reflect the real-world scenario: flattened, mixed-down audio with no access to separated sources.

3.1.2 Why fma_medium?

FMA is released in four main versions:

- **FMA Small:** 8,000 tracks (30s) across 8 genres
- **FMA Medium:** 25,000 tracks (30s) across 16 genres
- **FMA Large:** 106,000 tracks (30s) across 161 genres
- **FMA Full:** 106,000 full-length tracks

We selected `fma_medium` for three critical reasons:

1. **Scale and Feasibility:** At 22 GiB and 25,000 tracks, it strikes a practical balance between diversity and computational tractability, enabling robust training and validation without overwhelming disk or memory usage [41].
2. **Fixed-Length Segments:** Each track is cut to exactly 30 seconds, standardizing the input length for batch training and aligning with the tokenization and curriculum structures in the following [6].
3. **Genre and Production Diversity:** While unbalanced, `fma_medium` spans 16 broad genres and includes a mix of lo-fi, amateur, and semi-professional productions [53]. The lack of professional mastering on many tracks further matches our enhancement objectives.

3.1.3 Track Properties and Preprocessing

Each track in `fma_medium` is encoded as a 30-second stereo MP3 at 22.05 kHz. Before any degradation or tokenization, all tracks were:

- Resampled to **22.05kHz** for efficient training and model compatibility. Encoder vocoder inference later upsamples all outputs to 48kHz for high-fidelity audio synthesis.
- Normalized to a fixed peak amplitude of -1.0 dBFS to eliminate volume artifacts across tracks [79].
- Saved in `.wav` format using 16-bit PCM encoding [143].




Figure 3.1: Waveform and mel-spectrogram examples from the FMA Medium dataset. Note the diversity in timbral structure, genre, and energy profile across tracks.

3.1.4 Track Metadata and Usage

FMA provides rich metadata via the `tracks.csv` and `genres.csv` files [40]. For this work, we parsed only track ID and audio file path in order to stay true to our end-to-end pipeline goal. The genre label was cached externally.

We explicitly did not use any genre information during training, to avoid injecting bias or overfitting to style-dependent features. However, we preserve the genre labels for future subgroup evaluation (e.g., genre-wise PESQ performance).

3.1.5 License and Reproducibility

All tracks in FMA are licensed under Creative Commons, making them suitable for academic use and publication. We include scripts to preprocess the dataset and convert it into EnCodec token sequences in our public GitHub repository, ensuring full reproducibility.

Summary

The FMA Medium dataset offers a rare combination of:

1. Realistic, fully mixed musical content (non-stem, stereo bounces)
2. High genre and production diversity
3. Standardized input lengths and Creative Commons license

These properties make it an ideal benchmark for training and evaluating full-mix music enhancement systems.

3.1.6 Motivation for Paired Audio Learning

To train a supervised learning system capable of enhancing full-mix musical audio, it is essential to construct data sets that contain both clean and degraded versions of the same recording. This pairing allows a model to learn the mapping from noisy or poorly mastered tracks to their high-quality counterparts. Unlike traditional supervised data sets in computer vision or speech enhancement, where paired data are often available natively, music master data sets rarely contain the same recording in raw and professionally mastered forms [46].

As such, we adopt a synthetic approach. Beginning with relatively high-fidelity musical recordings from the FMA Medium dataset, we generated degraded variants by applying controlled audio effects. This process—which we coined *demastering*—enables the creation of parallel datasets where each degraded waveform is temporally aligned with its clean version, allowing frame-level losses and reconstruction objectives to be meaningfully computed.

Importantly, this synthetic degradation process allows for:

- **Controlled ground truth:** By knowing the exact transformation applied to each file, we can perform deterministic evaluation.

- **Task diversity:** Different degradations simulate various mastering challenges (e.g., poor EQ vs. excessive reverb).
- **Data scaling:** We can generate as many paired examples as needed for model pretraining or ablation studies.

This methodology mimics the structure of paired learning systems found in areas like image denoising and neural translation, while tailoring it to the complexities of polyphonic audio [165].

Design Principle

The goal is not to perfectly model real-world degradation pipelines, but rather to simulate a broad enough distribution of plausible impairments such that generalizable enhancement behavior can be learned.

3.2 STFT DEMASTERING: SIMULATING AUDIO DEGRADATIONS

Each track in the FMA Medium dataset is 30 seconds long and has been preprocessed to mono waveforms sampled at 22.05 kHz resampled from their original 44.1 kHz stereo format. This sample rate balances perceptual fidelity with computational efficiency and matches the model training assumptions. Starting from these "clean" audio files, we synthetically apply one or more degradative effects from a defined pool of signal processing operations. These include equalization, gain perturbation, dynamic range compression, reverberation, and echo. The output of this process is a degraded waveform of identical length and alignment to the original.

Let $x[n]$ denote the clean waveform. A degraded version $\tilde{x}[n]$ is produced as:

$$\tilde{x}[n] = \mathcal{D}(x[n]; \theta)$$

where $\mathcal{D}$ is a nonlinear transformation function parameterized by effect-specific hyperparameters θ . Each effect is sampled independently with parameter ranges described in the following.

3.2.1 Equalization

We simulate poor tonal balance by applying a second-order IIR peak filter with randomized parameters:

$$f_c \sim \mathcal{U}(500, 2000) \quad (\text{center frequency in Hz})$$

$$Q \sim \mathcal{U}(0.9, 1.2) \quad (\text{quality factor})$$

$$g \sim \mathcal{U}(-0.8, 0.8) \text{ dB} \quad (\text{gain})$$

These simulate scenarios like excessive low-mid boost or missing high-end sparkle.

3.2.2 Gain Mismatch

Simple global loudness shifts are simulated by scaling the amplitude [71]:

$$\tilde{x}[n] = 10^{g/20} \cdot x[n], \quad g \sim \mathcal{U}(-1, 1) \text{ dB}$$

This emulates inconsistent normalization between tracks.

3.2.3 Echo

To replicate poor delay settings or slapback effects [3]:

$$d \sim \mathcal{U}(0.1, 0.2) \quad (\text{delay in seconds})$$

$$a \sim \mathcal{U}(0.2, 0.3) \quad (\text{attenuation factor})$$

The echoed signal is added to the original waveform after a delay of $d \cdot \text{sr}$ samples.

3.2.4 Reverb

Artificial room responses are created using exponential impulse responses:

$$\tau \sim \mathcal{U}(0.2, 0.4) \quad (\text{decay constant})$$

$$T \sim \mathcal{U}(0.05, 0.15) \quad (\text{impulse duration in seconds})$$

The reverberated signal is computed via FFT convolution:

$$\tilde{x}[n] = x[n] * h[n], \quad h[n] = e^{-\tau t}, \quad 0 \leq t \leq T$$

3.2.5 Compression

Poor dynamics processing is simulated via soft-knee downward compression [22]:

$$\theta \sim \mathcal{U}(-12, -8) \text{ dB (threshold)}$$

$$r \sim \mathcal{U}(1.5, 2.5) \text{ (ratio)}$$

$$m \sim \mathcal{U}(0, 0.5) \text{ dB (makeup gain)}$$

Sample-level compression is applied with conditional attenuation:

$$\text{if } |x[n]| > \theta \rightarrow \tilde{x}[n] = \theta + \frac{|x[n]| - \theta}{r}$$

3.2.6 Effect Combination Strategy

For dataset diversity, we randomly apply 1-5 degradations per track; this quantity is specified in later chapters when we introduce the curriculum learning pipeline in

more depth. Each track is independently assigned a random combination of effects:

$$\mathcal{D} = \{\text{gain, EQ, compression, reverb, echo}\}$$

This ensures that the resulting training data reflect real-world entanglement between dynamic, spectral and temporal degradation [57].

Output Format

For every track i , we store:

- Clean waveform: `{id}_original.wav`
- Degraded waveform: `{id}_modified.wav`
- Effect parameters: `{id}.txt`
- Diagnostic mel spectrograms for before/after visualization.

This provides strong supervision while maintaining data transparency.

3.2.7 Examples of Degraded Outputs

These simulated degradations serve as the core mechanism by which paired training data is constructed, setting the stage for architectural designs in Chapters 4 and 6, where we begin training neural networks to learn inverse mappings from these degradations to perceptual restoration.




Figure 3.2: Example of degraded waveform with EQ boost and reverb tail; amplitude distribution is greatly affected. Top: original waveform. Bottom: modified waveform.




Figure 3.3: Mel spectrogram comparison. Degraded version exhibits unideal frequency distribution from EQ and smearing of high frequency bands.

3.3 PARAMETER RANGES AND REALISM

JUSTIFICATION

The integrity of any supervised learning approach to audio restoration is only as strong as the realism of its training data. In this work, we simulate audio degradations by applying signal processing effects with parameter values carefully sampled from empirically grounded ranges. The goal is to strike a balance between diversity and plausibility: to simulate challenging real-world audio degradation without veering into unrealistic or uninformative extremes.

In this section, we present the specific parameter distributions used to generate degraded waveforms for paired training and evaluation, and justify their relevance from a music production and perceptual audio standpoint.

3.3.1 Overview of Randomly Applied Effects and Ranges

We denote each degradation effect by its transformation $\mathcal{D}_e$ with the associated parameter space Θ_e . Let $x[n]$ be a clean waveform, and let $\tilde{x}_e[n] = \mathcal{D}_e(x[n]; \theta_e)$ be its degraded version. Table 3.1 summarizes the ranges of parameter samples that are generally used throughout this work. Later chapters will fine-tune these for our curriculum pipeline.

3.3.2 Motivating the Design of Each Range

Equalization. The range of center frequencies from 300 Hz to 5 kHz captures several common spectral problem areas [131]. Boosting or cutting in this range can simulate issues such as honky mids (500–800 Hz), nasal vocals (1–2 kHz), or harsh consonants (4–5 kHz). Gains of ± 6 dB allow tonal shifts that are audible but not so extreme as to render the content unrecognizable. The Q-factor is allowed to vary to

Table 3.1: Degradation parameters for synthetic audio impairment generation.

Effect	Parameter	Range	Description
Equalization (EQ)	Center freq f_c	300–5000 Hz	Targets midband clutter or vocal harshness
	Q-factor Q	0.5–2.0	Controls peak sharpness
	Gain g	–6 to +6 dB	Attenuate or boost
Global Gain Shift	Gain g	–3 to +3 dB	Simulates inconsistent loudness
Echo (Slapback Delay)	Delay d	100–250 ms	Audible repetitions
	Attenuation α	0.1–0.5	Echo decay strength
Reverb	Decay rate τ	0.2–1.0	Simulated room reverb
	IR duration T	50–400 ms	Impulse response length
Compression	Threshold θ	–24 to –6 dB	Controls onset of compression
	Ratio r	1.5:1–4.0:1	Controls amount of gain reduction
	Makeup gain m	0–3 dB	Restores output loudness

produce both broadband and narrowband effects, modeling everything from subtle tonal tilt to surgical notches.

Global Gain. Small shifts in global amplitude (± 3 dB) are typical of poor normalization workflows, especially in user-generated content [48]. Larger deviations would trigger clipping or inaudibility, which would be outside the scope of mastering-related degradations.

Echo. Delays in the 100–250 ms range mimic poorly configured slapback effects or default plugin settings in digital audio workstations [24]. These delays are short enough to fall within the temporal coherence window of musical phrases but long enough to be perceptually disruptive. Attenuation values from 0.1 to 0.5 replicate subtle to moderate reflections, commonly introduced by non-professional producers seeking spatial depth.

Reverb. Decay constants up to 1.0 emulate small to mid-sized artificial spaces [125]. Impulse durations (50–400 ms) correspond to impulse responses found in convolutional reverb units and DAWs such as Logic Pro, Ableton Live, and Reaper. These effects produce reverberant tails that interfere with mix clarity, particularly in rhythmically dense sections.

Compression. Thresholds down to -24 dB with ratios up to 4:1 simulate aggressive dynamic range squashing [124]. This mimics settings misused in amateur mastering chains or automatic loudness maximization tools. The optional makeup gain accounts for user compensation often applied post-compression, frequently without monitoring peak-to-loudness ratios.

3.3.3 Justification via Perceptual and Engineering Context

- **Plausibility:** All ranges reflect realistic values encountered in hobbyist production tools and tutorials [145], while still accounting for the possibility of mistakes made by novice users. This ensures that our degraded data remain grounded in real-world error distributions.
- **Variety:** The overlapping but non-identical parameter ranges allow for diverse outputs. For example, EQ + gain + reverb can emulate everything from washed-out lo-fi textures to harsh amateur master exports.
- **Temporal Alignment:** Because all degradations are applied nondestructively and without time warping, clean and degraded waveforms remain perfectly aligned, crucial for supervised training without dynamic time warping (DTW) penalties [112].

3.3.4 Illustration of Degradation Effects




Figure 3.4: Side-by-side visual comparison of effects applied with different parameters. Left column: clean audio. Right column: degraded versions. Top to bottom: EQ, reverb, compression.

3.3.5 Ensuring Robustness and Generalization

Finally, these parameter choices allow future models to learn not merely to invert effects, but to generalize across a range of musical genres, timbral content, and degradation combinations. This generalization capacity is central to building enhancement

systems that function in real-world production scenarios—where distortions may be unknown, variable, and entangled.

Engineering Summary

Rather than simulate rare or extreme degradations, we center our parameter choices on the "sweet spot" of plausible but impactful audio impairments. This provides both difficulty and diversity for downstream learning without veering into implausible signal alterations.

3.4 PREPROCESSING TOOLS AND LIBRARIES

To implement a robust and reproducible degradation pipeline, we leveraged a suite of Python-based signal processing and machine learning libraries. These tools enabled high-fidelity audio transformations, precise control over effect parameters, batch processing over thousands of clips, and structured preparation of training-ready datasets [143]. This section details the core components of our preprocessing workflow, from raw waveform manipulation to structured audio pair output.

3.4.1 Library Stack Overview

- **Librosa** [128] — Feature extraction, filterbank design, STFT, and perceptual representations (e.g., mel spectrograms).
- **Torchaudio** — Tensor-compatible waveform I/O, especially for integration with PyTorch training loops.
- **Soundfile** [13] (**pysoundfile**) — High performance reading and writing of .wav files, supporting 32-bit float precision and multichannel I/O.
- **SciPy** [128] — Signal modeling primitives, convolution, IIR filter design (e.g., peaking EQ), and mathematical utility functions.

- **Matplotlib** / **Librosa.display** - Visualization of waveforms and spectrograms for data set audit and debugging.
- **NumPy** — Vectorized operations for waveform degradation and metadata tracking.

Each of these libraries plays a targeted role in the degradation and pairing process. Librosa provides high-level operations such as Mel filterbank transforms, while ‘scipy.signal’ enables fine-grained DSP primitives [10]. Together, they form a hybrid pipeline that is both musically informed and computationally efficient.

3.4.2 Waveform I/O and Normalization

The raw audio files of the FMA Medium dataset are first loaded using `librosa.load()` or `torchaudio.load()` at their native sampling rate. Mono conversion is enforced for all waveforms to ensure consistency in downstream processing [115]. All waveforms are normalized to have the peak amplitude within $[-1, 1]$, preventing saturation of the dynamic range during the application of the effect.

Peak Normalization Rule

We avoid global loudness normalization (e.g. LUFS alignment) at this stage to preserve dynamic range variability. Peak normalization ensures safety during degradation without removing realistic amplitude variation.

3.4.3 Effect Application Pipeline

As detailed in Section 3.3, a randomly selected combination of degradation effects (e.g., EQ, reverb, echo, gain, compression) is applied to each clean waveform. Each effect is implemented as a composable function that takes a waveform and produces a degraded version, ensuring modularity and reproducibility [12].

Each function is parameterized by sampled values within the effect’s prescribed range and applied in a randomized order. For example:

```
y = apply_eq(x, sr, fc=1100, Q=1.2, gain=-3.5)
y = apply_echo(y, sr, delay=0.15, attenuation=0.3)
y = apply_compression(y, threshold=-12, ratio=3.0, makeup_gain=2.0)
```

After degradation, we store both the clean waveform $x[n]$ and the degraded waveform $\tilde{x}[n]$, together with a text log of applied transformations and their parameters for traceability.

3.4.4 Mel-Spectrogram Generation

In addition to waveform output, mel-spectrograms are computed and saved for each degraded audio file to support future qualitative evaluation and visualization. We use:

- `n_fft = 2048, hop_length = 512`
- `n_mels = 128`
- `power = 2.0` followed by dB conversion via `librosa.power_to_db`

This representation is used for dataset audits but not directly for model input, as Chapter 5 will introduce a token-based formulation instead of a spectrogram-based model [43].

3.4.5 Dataset Pair Construction and Metadata Logging

All waveform pairs are stored in structured directories under:

```
experiments/
  curriculums/
```

```
stage3_triple/  
  | output_audio/  
  |   123456_original.wav  
    123456_modified.wav  
  output_txt/  
    123456.txt  
  output_spectrograms/
```

Each pair is uniquely identified by the original FMA track ID. Metadata logs include the effects applied and their parameters, allowing reproducible subset regeneration and ablation studies.

3.4.6 Integration with Dataset Loader

We encapsulated the audio-pairing logic in a custom PyTorch-compatible `TokenPairDataset` class, housed in `dataset.py` [99] later reused in Chapter 6 for curriculum-aware token learning. While the current chapter only discusses waveform-level pairing, this loader supports flexible back ends:

- Automatic matching of clean/degraded waveform paths.
- On-the-fly or precomputed transformation support (see Chapter 6).
- Return of waveform pairs or token representations depending on mode.

Forward Compatibility

Although this chapter focuses on raw waveform degradation, all audio pairs are stored in a format compatible with later token-based architectures. This foresight allows plug-and-play experimentation with future models without re-generating datasets.

3.4.7 Computational Considerations

To support efficient data set construction:

- All audio pairs are pre-processed and saved to disk in .wav format to avoid recomputation.
- Spectrograms are cached as .png files to support visual audits during training.
- Logging and progress bars are included in the degradation script for reproducibility and debugging.

Processing 25,000 tracks with multiple effect chains took approximately 15 hours on a 12-core CPU system with parallelism. GPU acceleration is not required for pre-processing, reducing hardware dependencies and enabling reproducibility on consumer-grade systems.

3.4.8 Conclusion: A Scalable Degradation Framework

This preprocessing framework enables the scalable and audibly realistic construction of paired audio data sets necessary for training and evaluating restoration models. By grounding each transformation in real-world production practices and storing rich metadata alongside the audio pairs, we ensure that future training regimes can be both flexible and perceptually relevant [135]. With this pipeline in place, we now turn to the model architectures designed to learn from these pairs, beginning with our first explorations of U-Net in Chapter 4.

CHAPTER 4

EARLY APPROACHES AND EXPERIMENTAL BASELINES

4.1 SPECTROGRAM-BASED ARCHITECTURES

The initial phase of this research investigated neural architectures based on spectrograms for music enhancement to predict restored spectrograms aimed at reversing common degradations. This approach was inspired by numerous successful applications of U-Nets in audio source separation, speech enhancement, and image restoration [165, 22, 134, 23, 172]. However, as this section will demonstrate, applying these techniques to full-mix music restoration, especially using traditional inversion schemes from the spectral domain like Griffin-Lim, presents unique challenges.

4.1.1 Motivation and Setup

The intuition behind using U-Nets is their ability to extract hierarchical multiscale representations while preserving fine-grained temporal frequency information via skip connections and, in some cases, the addition of attention blocks [183]. Thus, we entered this stage with the assumption that U-Net architecture could enable both local transient-level and global tonal corrections on our dataset. In our pipeline, the degraded audio $x[n]$ was first transformed to a time-frequency representation using a Short-Time Fourier Transform (STFT) or mel spectrogram. A neural network

was trained to predict the corresponding “restored” spectrogram of the clean version $\hat{X}(f, t)$, which was then inverted back to waveform audio [66].

4.1.2 U-Net Architectures: 1D vs. 2D

The U-Net model consists of a series of downsampling and upsampling layers, with concatenations that pass high-resolution features from corresponding encoder to decoder blocks. We experimented with two architectural variants:

- **2D U-Nets** [165, 92] operating on magnitude spectrograms $X \in \mathbb{R}^{F \times T}$, where F is the frequency axis and T is time.
- **1D U-Nets** applied across the time axis for each mel frequency band separately, inspired by Wave-U-Net variants [174].

Mathematically, a 2D U-Net encoder at layer l computes:

$$\mathbf{h}^{(l)} = \sigma(\text{Conv2D}(\mathbf{h}^{(l-1)}))$$

with corresponding decoder:

$$\hat{\mathbf{h}}^{(l-1)} = \sigma(\text{Deconv2D}(\hat{\mathbf{h}}^{(l)}) + \mathbf{h}^{(l-1)})$$

The skip connection adds back the encoder feature map $\mathbf{h}^{(l-1)}$ to the corresponding decoder layer, enabling precise reconstruction of localized structures such as harmonic ridges or transients.

4.1.3 Explorations in Mel Band Resolution and Depth

A key experimental axis was the number of mel bins. Our initial (pre-neural vocoder) models used 128 mel bins, which we increased to 256 and then to 512 when we obtained

poor results in inference. Although these larger models improved perceptual clarity in some cases, they also introduced overfitting and slowed training due to higher dimensionality. We then tuned our model down to 80 mel bins in order to align with HiFi-GAN’s convention. These results are discussed in 4.2

We also varied base channel widths (`base_dim`) from 32 to 256, doubling per layer. Deeper networks captured longer-range musical structures (e.g., reverb tails, drum patterns), but with diminishing returns. In particular, models with `base_dim` = 64 offered a sweet spot between learning capacity and overfitting risk, if not in audio at least on paper.

4.1.4 Inference via Griffin–Lim Reconstruction

After predicting the enhanced magnitude spectrogram, we attempted to reconstruct the waveform using the Griffin–Lim algorithm. Griffin–Lim [66] is an iterative algorithm that estimates the phase $\phi(f, t)$ given a magnitude spectrogram $|X(f, t)|$, initializing with random phase and iteratively refining:

$$x^{(n+1)}[t] = \text{ISTFT} \left(|X(f, t)| \cdot e^{j \cdot \angle \text{STFT}(x^{(n)}[t])} \right)$$

Despite being standard in classical audio enhancement pipelines, Griffin–Lim introduces artifacts, including phase jitter, pre-echo, and spectral smearing, especially in complex, polyphonic content like music. These limitations became especially pronounced with degraded inputs that already suffered from harmonic distortion or time-smearing effects.

4.1.5 Key Failure Modes

Several core limitations of this spectrogram + U-Net + Griffin–Lim pipeline emerged:

1. **Noisy Reconstruction:** Even with accurate magnitude estimation, phase ar-

tifacts dominated output perception.

2. **Inability to Learn Dynamics:** Spectrograms discard temporal envelope information, harming recovery of compression or gain-related degradations [22].
3. **Spectral Overlap:** Music’s polyphonic nature (e.g., chords, layered synths) confused the U-Net’s frequency locality assumptions. The resultant output spectrograms resembled low-resolution versions of the target.




Figure 4.1: Examples of low resolution output via Griffin-Lim Inference

Observation

Although U-Nets were able to produce visually plausible spectrograms, audio output revealed that magnitude alone was insufficient. The core issue was not the structure of the network, but the limitations of the spectrogram domain and its inversion mechanisms with respect to complex full-mix audio.

4.1.6 Conclusion

This phase of experimentation revealed a core insight: the bottleneck in quality was not just architectural depth or mel bin count, but the representational mismatch between spectrograms and high-fidelity music. U-Nets lacked the capacity to infer phase or resolve overlapping harmonic content when phase recovery was not part of

training. These failures motivated a deeper investigation into vocoder-based inference (Chapter 4.2) and, ultimately, alternative representations such as discrete latent tokens explored in later chapters.

4.2 NEURAL VOCODER ATTEMPTS

4.2.1 The Promise of HiFi-GAN




Figure 4.2: Hifi-GAN model components [103]

Generative adversarial networks (GANs) have become a popular approach to audio waveform generation, with HiFi-GAN [103] standing out as a state-of-the-art vocoder capable of synthesizing high-fidelity speech from mel spectrograms in real time. Unlike autoregressive models like WaveNet, HiFi-GAN leverages multiple discriminators with specialized receptive fields and periodic convolutions to model both local and periodic audio structure efficiently [97]. Its generator G_θ learns a mapping from a

mel-spectrogram $\mathbf{M}$ to a waveform $\hat{\mathbf{x}}$:

$$\hat{\mathbf{x}} = G_{\theta}(\mathbf{M}) \in \mathbb{R}^T,$$

where $\mathbf{M} \in \mathbb{R}^{N \times F}$ is a 2D mel spectrogram with F frequency bins and N time steps. The training objective combines adversarial losses $\mathcal{L}_{\text{adv}}$, feature matching losses $\mathcal{L}_{\text{FM}}$, and mel reconstruction losses $\mathcal{L}_{\text{mel}}$:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{mel}} + \lambda_{\text{adv}}\mathcal{L}_{\text{adv}} + \lambda_{\text{FM}}\mathcal{L}_{\text{FM}}.$$

4.2.2 Initial Setup and Adaptation

We attempted to repurpose HiFi-GAN for reconstructing music audio from degraded mel spectrograms generated by our early enhancement networks (Section 4.1). We repurposed the architecture of our original implementation [103] 30-second stereo tracks, assuming that the sample rate we had implemented was sufficient to generate mel spectrograms and Griffin-Lim inference. However, complete retraining and resampling of our original dataset pairs when we had to regenerate the dataset with HiFi-GAN’s reduced quantity of 80 mel bins.

The inference code (`inference.py`) follows a straightforward process: load degraded WAV files, compute mel spectrograms, and synthesize waveforms using a trained HiFi-GAN generator. Despite correct architectural setup, all trials suffered from sample rate issues or normalization mismatches leading to major reconstruction flaws when applied to full-mix music inputs.

4.2.3 Failure Modes Observed

Despite HiFi-GAN’s strength in speech synthesis, applying it to music enhancement failed for multiple structural reasons:

1. **Genre Generalization Failures:** HiFi-GAN was trained on speech corpora such as VCTK [27]. When applied to full-bandwidth, polyphonic music, it hallucinated high frequencies and introduced harsh aliasing in transient regions. Guitar transients, drums, and stereo reverbs were especially problematic.
2. **Stereo Collapse:** The model consistently outputs mono or phase-incoherent stereo, even when fed stereo mel spectrograms derived from stereo mixes. This was likely due to implicit mono assumptions in its training regime.
3. **Inference Latency and Quality Tradeoffs:** While HiFi-GAN boasts low latency on speech, the batch inference pipeline used with 30-second music tracks caused memory spikes on both the RTX 3070 Ti and RTX 6000 GPUs, requiring patchwork memory management. The output was distorted and did not align in time with the original content, often misaligned by orders of magnitude.
4. **Fundamental Misalignment:** Most critically, HiFi-GAN’s training assumes clean mel targets. Our use case required denoising or enhancement *prior* to waveform synthesis. Any enhancement errors magnified during vocoder reconstruction and retraining HiFi-GAN on degraded input distributions failed to stabilize perceptual quality.

Limitation Summary

HiFi-GAN failed to deliver meaningful reconstruction quality when applied to enhanced or degraded music spectrograms, regardless of fine-tuning efforts. The lack of phase-aware, stereo-consistent, and high-bandwidth generalization made it unsuitable as a final decoding step in our pipeline.

4.2.4 Conclusion

The limitations of HiFi-GAN in this context are not failures of the model per se, but rather mismatches in data domain and application. Vcoders like HiFi-GAN excel in narrow-band, monophonic, speech-centric applications. Their adversarial training is not optimized for high-dimensional, polyphonic, effect-heavy full mixes, especially when input spectrograms are derived from enhanced or restored sources rather than pristine ground truth.

These challenges ultimately motivated a pivot toward tokenized audio modeling (Chapter 5), which bypassed the need for unstable spectrogram inversion altogether. Instead of relying on waveform reconstruction from continuous representations, our system began operating directly in discrete EnCodec token space, where both training and inference remained entirely differentiable and invertible by design.

4.3 PARAMETER INVERSION NETWORKS

A novel early hypothesis during the exploratory phase of this project was whether a neural network could learn to predict explicit parameters of audio degradation and simply reverse them during inference. Rather than learning an end-to-end mapping between degraded and clean signals, the idea was to approach the task as a parameter regression problem: estimate the degradation vector Θ (e.g., EQ gain, compressor threshold, echo delay), then analytically apply the inverse transformation.

This section outlines the motivation, structure, and failure of such an approach, particularly the use of LSTM-based time series forecasting networks inspired by previous work in audio dynamics modeling, such as the paper *Application of ML-Based Time Series Forecasting to Audio Dynamic Range Compression* [22].

4.3.1 Forecasting-Based Parameter Estimation

The aforementioned paper by Brunet et al. [22] proposes the use of LSTM-based recurrent neural networks to forecast compressor parameters in the dynamic range from raw audio waveforms. Their model predicts frame-wise values of threshold, ratio, attack, and release, relying on handcrafted features and domain-specific transformations. Specifically, their formulation minimizes the loss between true and estimated compressor parameters:

$$\mathcal{L}_{\text{params}} = \sum_{t=1}^T \left\| \Theta_t - \hat{\Theta}_t \right\|_2^2, \quad (4.1)$$

where Θ_t denotes the true parameter vector at time t , and $\hat{\Theta}_t$ is the LSTM output.

Inspired by this, our architecture accepted either a mel spectrogram or a raw waveform of the degraded signal and attempted to regress a global vector Θ representing the degradation configuration. These parameters were then passed into inverse transformations (e.g., inverse gain, EQ reversal, compressor expansion), aiming to recreate the clean signal.




Figure 4.3: Brunet's proposed system

4.3.2 LSTM Architecture and Mathematical Formulation

The LSTM network used in our experiments followed the standard formulation [32]:

$$f_t = \sigma(W_f x_t + U_f h_{t-1} + b_f) \quad (\text{forget gate}) \quad (4.2)$$

$$i_t = \sigma(W_i x_t + U_i h_{t-1} + b_i) \quad (\text{input gate}) \quad (4.3)$$

$$o_t = \sigma(W_o x_t + U_o h_{t-1} + b_o) \quad (\text{output gate}) \quad (4.4)$$

$$\tilde{c}_t = \tanh(W_c x_t + U_c h_{t-1} + b_c) \quad (4.5)$$

$$c_t = f_t \odot c_{t-1} + i_t \odot \tilde{c}_t \quad (4.6)$$

$$h_t = o_t \odot \tanh(c_t) \quad (4.7)$$

Here, x_t is the feature vector at timestep t , h_t is the hidden state, and c_t is the memory cell. Our implementation fed in either STFT magnitudes or mel features over time as x_t .

We minimized a hybrid loss:

$$\mathcal{L}_{\text{total}} = \lambda_1 \cdot \mathcal{L}_1(\hat{\Theta}, \Theta) + \lambda_2 \cdot \mathcal{L}_{\text{spec}}(x_{\text{inv}}, x_{\text{clean}}) \quad (4.8)$$

where $\mathcal{L}_1$ is the regression loss on degradation parameters, and $\mathcal{L}_{\text{spec}}$ is an L_2 loss in the spectral domain between the restored signal x_{inv} (after inverting predicted parameters) and the original clean signal x_{clean} .

4.3.3 Bridging Two Domains

To increase modeling capacity, a hybrid system was constructed where the encoder of a U-Net operated over 2D spectrogram representations, followed by global average pooling and an attached LSTM layer responsible for regressing Θ . However, this

introduced severe complications.

The U-Net naturally operates in a spatially structured frequency-time representation. The LSTM, in contrast, assumes ordered temporal sequences. Bridging these required flattening or aggregating the U-Net feature maps into sequences—a lossy transformation that eroded critical frequency structure and local correlations.

Even when this mismatch was handled carefully, the true problem emerged: *degradation inversion is not analytically reliable*. This highlighted a broader principle: inverse parameterization is ill-posed and often non-differentiable in practice, rendering regression-based recovery ineffective for perceptual tasks.

4.3.4 Failure Mode and Illusion of Low Loss

Despite promising validation curves, with $\mathcal{L}_1(\hat{\Theta}, \Theta)$ quickly decreasing, inference revealed a deeper issue. Many degradations (especially reverb, echo, compression) are nonlinear, non-invertible, or dependent on signal history. For instance:

- Reverb smears transients and cannot be reversed by simple IR subtraction.
- Compressors introduce gain reduction dynamics that depend on both instantaneous and historical amplitudes [132].
- EQ alterations cannot be reliably reversed when frequency masking or aliasing has occurred.

As a result, even perfect parameter regression did not guarantee signal recovery. Furthermore, the network began to "game" the loss: discovering $\hat{\Theta}$ vectors that led to slightly better audio, but did not resemble the true degrading parameters.

4.3.5 Conclusion

Ultimately, this architecture taught a critical lesson: minimizing parameter error is not equivalent to maximizing perceptual quality. The regression task optimized a

proxy that was neither fully observable nor fully correctible. Combined with the domain mismatch and inference unreliability, this led to the abandonment of the parameter inversion idea.

This failure marked a conceptual turning point, leading toward direct signal-space restoration strategies and away from analytically-defined reversals. It also justified a future shift toward latent token transformation, as described in Chapter 5.

Key Insight

Not all degradation effects are analytically invertible. Optimizing parameter regression is insufficient when signal corruption involves nonlinear, irreversible processes.

4.4 LOSS-REALITY GAP: WHY THESE FAILED

Despite early experimentation with U-Nets, LSTMs, and hybrid regression-based models, a recurring problem began to emerge: quantitative training losses consistently decreased, yet perceptual quality during inference often remained poor or even deteriorated. This discrepancy between numerical loss minimization and subjective audio restoration quality is known as the *loss-reality gap*, a phenomenon frequently encountered in generative and signal enhancement tasks [43].

4.4.1 Loss Functions Used

Across multiple architectural iterations, we employed the following objective functions:

- **Time-domain L_1 Loss:**

$$\mathcal{L}_{\text{time}} = \frac{1}{N} \sum_{n=1}^N |x_n - \hat{x}_n|$$

Direct sample-wise error, used for waveform predictions.

- **Spectral Convergence Loss: [10]**

$$\mathcal{L}_{\text{sc}} = \frac{\| |\text{STFT}(x)| - |\text{STFT}(\hat{x})| \|_F}{\| |\text{STFT}(x)| \|_F}$$

Measures normalized magnitude difference in the frequency domain.

- **Mel Spectrogram Loss:**

$$\mathcal{L}_{\text{mel}} = \| M(x) - M(\hat{x}) \|_1$$

Where $M(\cdot)$ denotes the mel filterbank transformation, emphasizing perceptually relevant bands.

- **Parameter Regression Loss:**

$$\mathcal{L}_{\Theta} = \left\| \Theta - \hat{\Theta} \right\|_2^2$$

Used exclusively in the parameter inversion experiments.

These losses—particularly when used in combination—provided a seemingly robust training signal. In several runs, validation curves plateaued at impressively low levels, suggesting successful convergence. However, side-by-side listening of reconstructions revealed harsh distortions, stereo collapse, reverb smearing, and loss of instrumental separation.

4.4.2 Sources of Discrepancy

The misalignment between loss curves and perceptual quality is multifaceted:

- 1. Phase Insensitivity** Most frequency-based losses (e.g., mel loss, spectral convergence) rely solely on magnitude spectra. This ignores phase, which is essential to reconstructing transients, spatial width, and stereo imaging in music [162].
- 2. Ambiguity of L_1/L_2 in Waveform Space** While pointwise errors are mathematically simple, they fail to capture perceptual similarity. A small L_1 loss between waveforms may correspond to a perceptually drastic shift—e.g., a misaligned beat, flanged vocal, or collapsed stereo panning.
- 3. Griffin-Lim Limitations** For spectrogram-based architectures, audio reconstruction often relied on the Griffin–Lim algorithm [66] to infer missing phase. This iterative method introduces artifacts, and in practice, reconstructions often sounded robotic, hollow, or unstable in pitch.
- 4. Overfitting to Proxy Objectives** In the LSTM parameter prediction experiments, networks learned to “guess” parameter combinations that produced spectrally similar audio—but not necessarily correct or invertible configurations. This was a form of shortcut learning: the network minimized loss without learning to solve the actual restoration problem.
- 5. Lack of Structure in Spectrogram Space** While U-Nets perform well on visually structured domains like images, magnitude spectrograms of music—especially when distorted—exhibit low structural redundancy and high frequency variance. This weakens the value of convolutional inductive bias and makes training unstable.

4.4.3 Visual Examples of Failure

In the figures that follow, notice the low quality outputs when compared to their inputs. There is lots of smearing indicative of issues caused by high frequency variance and detail in the spectral space.




Figure 4.4: Example of low resolution output via Griffin-Lim Inference




Figure 4.5: Another of low resolution output via HiFi-GAN (Pre-Trained Neural Vocoder) Inference

4.4.4 Beyond Conventional Losses

Given these challenges, we concluded that:

- No combination of L_1 , mel, or spectral convergence loss was sufficient to align training signal with perceptual goals.

- Existing vocoders (Griffin–Lim, HiFi-GAN) were not designed for restoration from non-ideal spectrograms.
- A more structured, constrained representation space was needed—one that preserved phase implicitly, allowed discrete transformations, and could be trained on robust, pairwise data.

This led to a paradigm shift: instead of predicting waveforms or spectrograms directly, we transitioned to operating in a *discrete latent audio token space*, introduced in Chapter 5. Token-based systems enabled sharper learning signals, faster training, and more interpretable transformations.

Key Insight

Loss functions are not neutral—they impose a worldview on what constitutes “good” audio. In music restoration, perceptual alignment is more important than numerical convergence.

4.5 LESSONS LEARNED: TOWARD TOKEN-BASED RESTORATION

The culmination of these early-stage experiments—spanning spectrogram-based U-Nets, vocoder-driven waveform reconstructions, and LSTM-based parameter regression—revealed a set of fundamental limitations in conventional restoration strategies. These limitations, while initially frustrating, ultimately motivated a more robust, efficient, and perceptually aligned approach: modeling directly in a discrete, tokenized latent space.

4.5.1 Why Existing Models Fell Short

Our investigations reinforced several key weaknesses in prevailing architectures and workflows:

- **Spectrogram-based U-Nets**, as explored in works such as *Audio Segmentation with U-Net Architecture* [163], proved effective for segmentation and enhancement in clean domains. However, they failed to generalize when input spectrograms were contaminated with multiple interdependent degradations (e.g., reverb + echo + EQ skew). This breakdown was compounded by their reliance on Griffin–Lim reconstruction, which often failed to recover high-fidelity audio.
- **Wave-U-Nets** [157] and waveform-domain models did preserve phase but were prohibitively slow to train on high-resolution 48 kHz data. Additionally, their convolutional receptive fields lacked the capacity to capture long-range temporal dependencies necessary in music signals spanning tens of seconds.
- **HiFi-GAN**, while state-of-the-art for speech synthesis, failed to produce realistic reconstructions when applied to enhanced music spectrograms. Even after extensive adjustments to mel-band resolution and resampling the entire FMA dataset to 22.05 kHz to match training expectations, outputs suffered from persistent stereo collapse, transient smearing, and unnatural tonal colorations. This failure emphasized that HiFi-GAN is highly sensitive to the distribution of its training data and lacks robustness in general audio enhancement settings.
- **Parameter Inversion via LSTM Forecasting**, inspired by work such as *Application of ML-Based Time Series Forecasting to Audio Dynamic Range Compression* [22], demonstrated good regression performance on synthetic targets. However, our findings showed that minimizing loss on parameters (e.g.,

threshold, ratio, Q-factor) was orthogonal to producing high-quality audio. The inferred audio, even when the parameter vector appeared “accurate,” was perceptually poor—highlighting the inability of the regression path to capture context, artistic intent, or perceptual nuance.

These observations were not isolated—rather, they are consistent with a broader trend in the literature. Numerous restoration systems in the field of audio enhancement rely on assumptions that break down under real-world music conditions. For instance:

- **Smit and Lee** [172] used a U-Net+LSTM system to predict effect parameters for musical stems. But the model’s effectiveness depends on clean instrument isolation and ignores mixed, overlapping signals.
- **DeepVQE** [84] applies masking techniques suitable for formant-driven speech but cannot handle dense harmonic structures in full mixes.
- **iZotope’s RX suite** [90], while powerful, requires human fine-tuning and stems from empirical DSP rules—lacking the learnability and automation potential of deep learning.

Empirical Summary

Conventional audio restoration systems either assume access to clean stems, operate only on synthetic degradations, or lack perceptual alignment in their loss functions. These assumptions crumble under real-world, fully-mixed, degraded music.

4.5.2 Reframing the Problem

Based on these failures, a shift in perspective was necessary. Rather than attempt to denoise raw waveforms or recover STFTs with problematic phase components, the

core idea emerged:

What if we learned to translate from degraded audio to clean audio not in continuous waveform space, but in a discrete, structured latent space where high-quality audio already lives?

This reframing draws inspiration from several foundational advances:

- **Discrete Audio Tokens via EnCodec** [28]: Meta AI introduced a multi-codebook vector quantized autoencoder (VQ-VAE) capable of compressing audio into compact discrete representations while preserving quality. These tokens allow high-fidelity reconstruction while being much easier to model.
- **Token-to-Token Models for Audio Generation:** Papers such as *AudioLM* and *MusicLM* [20, 1] demonstrated that autoregressive models trained over EnCodec tokens could generate realistic and stylistically consistent audio. While those works focused on generation, the framework suggested a new path for enhancement and restoration.
- **Curriculum Learning:** Motivated by the difficulty of learning directly from heavily degraded inputs, we designed a progressive curriculum training pipeline (introduced in Chapter 5) where the model is first trained on identity mappings before facing increasing degradation complexity. This technique was essential for generalization in token space.

4.5.3 Conclusion and Outlook

These insights set the stage for a token-centric framework. Rather than minimize unstable losses in a poorly aligned continuous space, we propose to operate in a more stable and semantically rich discrete token space, enabling:

- Faster training and inference

- Robust generalization to mixed and unstructured audio
- Interpretable, modular transformations aligned with musical perception

The next chapter introduces **Token U-Net**—a novel architecture designed to enhance audio in token space. We detail the tokenization strategy using EnCodec, the curriculum stages, the loss functions used, and the architecture’s core components.

Transition to Chapter 5

After extensive experimentation in waveform and spectrogram domains, we present a fundamentally new approach: treating audio enhancement as a transformation over learned audio tokens. This shift forms the heart of this thesis.

CHAPTER 5

TOKENIZATION VIA META AI’S ENCODEC

5.1 WHY TOKENS?




Figure 5.1: Visualization of tokens from a waveform in the latent space [28]

Conventional audio restoration techniques operate in continuous domains, such as raw waveform space or time-frequency representations, such as mel spectrograms [176]. Although effective in narrow contexts, these representations suffer from critical drawbacks when applied to full-mix music enhancement. Waveform modeling is computationally expensive, unstable under noise, and difficult to optimize. Spectrograms, though more compact, discard phase information and impose resolution trade-offs that limit the fidelity of the reconstruction [10]. Our early baselines, including U-Net

spectrograms and HiFi-GAN reconstructions, clearly demonstrated these weaknesses (see Chapter 4): models often overfit, introduced aliasing artifacts, or collapsed stereo and transient detail.

To overcome these limitations, we adopt a discrete representation: compressed audio tokens derived from Meta AI’s *EnCodec* model [51]. EnCodec transforms audio into sequences of quantized latent vectors using residual vector quantization (RVQ) [28], resulting in a compact and semantically rich token stream. Each code represents an abstract, learned basis of audio structure, capturing everything from pitch to texture to temporal dynamics.

This shift allows us to frame music restoration as a token-to-token enhancement task. Rather than predicting waveforms or magnitude spectra, our model learns to map *degraded token sequences* to *clean token sequences*, all within a discrete, learnable, and phase-consistent latent space. The advantages are profound:

1. **Compression and Speed:** A 30-second stereo clip is reduced to only a $[16 \times 2250]$ token matrix, allowing fast training and GPU-efficient batching.
2. **Semantic Abstraction:** Tokens encode perceptual features such as timbre, attack, stereo width, and distortion. They provide an implicit basis for recognizing effects such as EQ, compression, and reverb, without the need for raw amplitude modeling.
3. **Multi-Effect Compatibility:** Because EnCodec tokens reflect *all aspects* of musical structure, they serve as a universal substrate for learning simultaneous restoration tasks: masking reverberation, undoing gain shifts, improving stereo field, and more.

Importantly, this discrete representation supports high-level architectural advances. In Chapter 7, we introduce a compact and shallow 1.08 billion parameter *Token U-Net* that operates directly on these token streams. It features gated skip connections,

FiLM modulation, a temporal context block, and five dedicated output heads trained via multi-effect curriculum learning.

This approach mirrors progress in adjacent fields: discrete latent modeling has revolutionized NLP (e.g. byte-pair encoding in Transformers), image compression (e.g. VQ-GANs) and audio generation (e.g., AudioLM, MusicLM) [12, 1]. By adopting a token-centric paradigm, we unlock structured modeling, tighter loss alignment with human perception, and greater restoration fidelity across musical degradations.




Figure 5.2: Schematic of EnCodec’s architecture from [51], illustrating convolutional encoders, RVQ quantization, and decoder pathways. Tokens are derived from quantized latent embeddings at fixed frame rates.

In this chapter, we dissect the internals of EnCodec, detail our token precomputation strategy, and visualize token behavior in clean and degraded examples. These discrete tokens serve as the learnable substrate for our restoration system, offering interpretability and efficiency.

5.2 ENCODEC INTERNALS

EnCodec [51] is a state-of-the-art neural audio codec developed by Meta AI for high-fidelity compression and decoding of raw audio. Its core innovation lies in the transformation of waveforms into compact quantized token sequences using residual vec-

tor quantization (RVQ), which allows discrete modeling of complex audio signals. This tokenization forms the backbone of our Token U-Net system: all enhancement, restoration, and prediction tasks are conducted entirely within this learned latent space.

EnCodec combines three key components:

- A **deep convolutional encoder**, which maps the waveform into a high-dimensional latent space.
- A **multi-stage RVQ quantizer**, which discretizes the latent space into a fixed number of codebook indices (tokens).
- A **decoder**, trained with adversarial losses, which reconstructs the waveform from the quantized tokens.

5.2.1 Encoder and Decoder Pathways

Let $x(t)$ be a raw waveform sampled at 22.05 kHz. The EnCodec encoder f_{enc} maps $x(t)$ to a continuous latent embedding $z(t)$:

$$z(t) = f_{\text{enc}}(x(t)) \in \mathbb{R}^{C \times T}$$

Here, C is the latent channel dimensionality and T is the number of temporal frames, which typically corresponds to 75 frames per second. These continuous embeddings are then quantized into discrete tokens using RVQ, and later passed to the decoder f_{dec} for reconstruction:

$$\hat{x}(t) = f_{\text{dec}}(\text{tokens})$$

5.2.2 Residual Vector Quantization (RVQ)

EnCodec performs K -stage residual vector quantization to discretize $z(t)$. At each stage k , a learned codebook $c_k \in \mathbb{R}^{N \times D}$ maps each latent vector z_t to the closest centroid:

$$z_t \approx \sum_{k=1}^K c_k[q_k(z_t)]$$

where $q_k(z_t)$ is the index of the codebook (token) in stage k , and N is the number of centroids per codebook (typically $N = 1024$). This produces a token matrix:

$$\text{tokens} \in \mathbb{Z}^{K \times T}$$

Each row in this matrix corresponds to a quantizer stage, and each column to a time frame. The resulting representation is highly compressed and semantically meaningful, enabling downstream learning of high-level signal corrections directly in token space.

5.2.3 Bandwidth, Frame Rate, and Compression

The effective bit rate of EnCodec depends on the number of quantizers K , the token frame rate r (e.g., 75 fps) and the number of bits per codebook index $\log_2 N$:

$$B = \frac{K \cdot \log_2(N) \cdot r}{1000}$$

For example, with $K = 16$ codebooks, $N = 1024$ centroids per codebook, and $r = 75$:

$$B = \frac{16 \cdot 10 \cdot 75}{1000} = 12.0 \text{ kbps}$$

Because this value is per channel, stereo audio encoded with EnCodec at 48 kHz yields an overall bit rate of:

$$B_{\text{stereo}} = 2 \times 12.0 = 24.0 \text{ kbps}$$

This 24 kbps representation is used throughout our training and evaluation pipeline. Thus, a 30-second stereo clip is represented as $[16 \times 2250]$ tokens per channel, compact enough for large-batch training, yet rich enough to preserve perceptual fidelity.

5.2.4 Training Procedure and Discriminators

EnCodec is trained end-to-end using:

- **Reconstruction losses** in waveform space (e.g. mean squared error),
- **STFT-based spectral losses** to align perceptual frequency content,
- **Adversarial losses** from multiscale discriminators operating in both waveform and spectrogram domains.

These adversarially trained discriminators allow the decoder to reconstruct realistic waveforms even when token inputs are slightly distorted or incomplete. This robustness is critical for our downstream use case: the Token U-Net is trained to *transform degraded token sequences into clean ones*, and must rely on EnCodec’s decoder to convert predicted tokens back into intelligible audio.

5.2.5 Summary Configuration

In the following sections, we detail how these token sequences are precomputed and bundled for high-efficiency, curriculum-aware training. We also show how auxiliary features (e.g., STFT, mel-spectrograms) are extracted and embedded into each sample for multi-objective loss functions.

Table 5.1: EnCodec Configuration Used for Token U-Net Training

Component	Description	Value
Sample rate	Input waveform rate	48,000 Hz
Frame rate r	Token frame resolution	75 fps
Codebooks K	Number of RVQ stages	16
Codebook size N	Entries per codebook	1024
Token shape	Final token matrix	$[16 \times 2250]$ (30s)
Total bandwidth	Stereo bitrate	24 kbps

5.3 PRECOMPUTING TOKENS FOR CURRICULUM TRAINING

To enable efficient and reproducible training of our curriculum-aware Token U-Net model, we precompute all EnCodec representations in advance. Each sample, consisting of a clean-degraded audio pair, is tokenized, analyzed, and serialized into a structured `.pt` file that contains all necessary tensors and metadata for multitask learning. This section details the degradation pipeline, curriculum stages, feature-bundling format, and technical infrastructure used to ensure speed, fidelity, and reproducibility.

5.3.1 Curriculum Degradation Stages

Our training curriculum is inspired by staged learning principles in vision and speech enhancement [14, 133], and simulates increasingly severe audio degradations typical of amateur or archival mixes. We define five curriculum stages:

- **Stage 0 – Identity:** Clean audio paired with itself; the model must reconstruct the tokens exactly.
- **Stage 1 – Single Effect:** A single mastering-related degradation (EQ, gain, compression, reverb, or echo) is applied.




Figure 5.3: Schematic of a Stage 3 Training pair and the curriculum stage progression

- **Stage 2 – Double Effects:** Two degradations applied in random order and parameter combination.
- **Stage 3 – Triple Effects:** Three randomly selected effects; interaction and masking effects increase task difficulty.
- **Stage 4 – Full Random:** Up to five simultaneous degradations, mimicking real-world low-fidelity bounce downs.

To enforce generalization and simulate extreme real-world degradations, we introduce *stronger variants* of Stages 3 and 4 by widening the parameter ranges of each degradation. For example, an EQ center frequency f_c originally chosen in the range $[f_{\text{lo}}, f_{\text{hi}}]$ is perturbed to:

$$f'_c \in [f_c - 0.3(f_{\text{hi}} - f_{\text{lo}}), \min(f_{\text{hi}} + 0.3(f_{\text{hi}} - f_{\text{lo}}), 20000)]$$

These wider intervals increase out-of-distribution variability and enforce generalization across adverse conditions. Stronger stages are denoted as `stage3_stronger`

and stage4_stronger in our logging and checkpoint structure.

5.3.2 Comprehensive Feature Bundling in .pt Files




Figure 5.4: Visual sample of contents of feature bundling in .pt files

Each pair of tracks is processed into a single PyTorch .pt file containing all the features required for training and evaluation. All audio inputs are resampled to 22.05 kHz before EnCodec encoding at 48 kHz to ensure consistency with the training distribution. The contents include:

- $\mathbf{X} \in \mathbb{Z}^{n_q \times T}$: Tokenized degraded input
- $\mathbf{Y} \in \mathbb{Z}^{n_q \times T}$: Tokenized clean target
- $\mathbf{s} \in \mathbb{R}^T$: Per-frame scale vector from EnCodec for loudness normalization

- $\mathbf{Y}^{\text{STFT}} \in \mathbb{R}^{n_q \times F \times T}$: STFT magnitudes of clean audio, where:

$$\mathbf{Y}^{\text{STFT}}[q, f, t] = |\text{STFT}(Y_q)[f, t]|$$

- $\mathbf{Y}^{\text{Mel}} \in \mathbb{R}^{M \times T}$: Mel spectrogram of the clean waveform (optional)
- Metadata: Sample rate, number of codebooks, compression rate, degradation parameters, hashes
- Cached Waveforms (optional): For visualization, reference, and debugging
- Normalization ranges: Min/max scalars for future feature scaling or denormalization

We refer to this as the **frozen lunchbox** design: each file is a fully self-contained container, supporting all Token U-Net objectives from cross-entropy to perceptual embedding losses.

5.3.3 Speed Gains and Training Efficiency

Real-time EnCodec encoding is computationally intensive because of its deep convolutional structure. By precomputing all token pairs, STFTs, and mel features in advance, we achieved a **40× speedup in data loading and batch preparation**, enabling high-throughput GPU training with minimal I/O stalls. This allowed our model to scale to more than 1.08 billion parameters while maintaining stable iteration rates on RTX 6000-class GPUs.

5.3.4 SHA-256 Hashes and Reproducibility

To ensure data integrity and eliminate ambiguity in training runs, we compute SHA-256 hashes of clean and degraded audio for every `.pt` file:

```
"original_audio_hash": "ad3285c...7f1b",
"modified_audio_hash": "39cdd92...a001"
```

These hashes allow us to:

- Detect accidental file corruption
- Guarantee reproducibility across training sessions
- Audit the degradation provenance of any input file

All hash generation and verification is embedded in the precompute pipeline and training dataloader.

5.3.5 Auxiliary Spectral Domains for Loss Functions




Figure 5.5: Visual example of auxiliary STFT sum for token-spectrum alignment

To support advanced spectral loss terms used in later training stages, we cache:

1. **STFT Magnitudes** for token-spectrum alignment:

$$\mathcal{L}_{\text{STFT}} = \left\| \mathbf{Y}^{\text{STFT}} - \hat{\mathbf{Y}}^{\text{STFT}} \right\|_1$$

2. **Mel Spectrograms** for perceptual loss layers (e.g., $\mathcal{L}_{\text{Mel-L1}}$, $\mathcal{L}_{\text{Mel-Cos}}$)

Both spectral features are saved in FP16 to reduce memory usage while maintaining sufficient precision for gradient stability.

5.3.6 Error Handling and Fault Tolerance

The precomputation script includes multiple validation checks before saving any bundle.

- Dimension validation with wildcard matching
- Token range checks to avoid out-of-bound centroids
- Backup on failure: `.pt.bak` is written on any error
- Logging per stage, including timestamp, hash, duration, and parameter trace

This infrastructure ensures that training never stops due to misdirected input, and every sample in the training set is auditable, reproducible, and backed by quantitative degradation logs.

Summary: Precomputing tokenized curriculum stages into structured, hash-verified `.pt` files dramatically increased our training efficiency, reproducibility, and experimental clarity. It also enabled consistent support for multi-loss training objectives and setup for the multi-head prediction architecture described in Chapter 6.

5.4 .PT FILE FORMAT AND METADATA

To streamline training, debugging, and downstream analysis, we serialize each audio pair into a standardized PyTorch `.pt` file. This format encapsulates all relevant tensors—tokens, scales, spectral features—as well as detailed metadata for provenance, auditability, and loss conditioning. Each file acts as a self-contained training unit within the curriculum-aware pipeline.

5.4.1 Structure Overview




Figure 5.6: Our pipeline to achieve the frozen lunchbox .pt token bundle

Every file is indexed by a numeric stem (e.g. 083271.pt) and includes:

$\mathbf{X} \in \mathbb{Z}^{n_q \times T}$	Degraded tokens
$\mathbf{Y} \in \mathbb{Z}^{n_q \times T}$	Clean reference tokens
$\mathbf{scales} \in \mathbb{R}^T$	Per-frame scale values from EnCodec encoder
$\mathbf{Y_stft_mag} \in \mathbb{R}^{n_q \times F \times T}$	Magnitude STFT of clean tokens
$\mathbf{mel_spec} \in \mathbb{R}^{M \times T}$	Mel spectrogram of clean waveform
$\mathbf{wav_o}, \mathbf{wav_m} \in \mathbb{R}^{1 \times T}$	Cached original and degraded waveforms (optional)

Typical values: $n_q = 16$ (codebooks), $F = 1025$ (STFT bins), $M = 128$ (mel bands) and $T \sim 4500$ (for 30s of audio at 75 fps). All spectral tensors are stored in half-

precision (FP16) to minimize memory overhead without compromising training quality.

5.4.2 Audio Hashing for Provenance

To ensure dataset reproducibility and protect against corruption, each file embeds SHA-256 hashes of the source audio:

```
"original_audio_hash": "cfae9e1...c1e7",
"modified_audio_hash": "7d5bce6...1a35"
```

These are computed over the first and last 1MB of each waveform (or the entire file if smaller). During training, hashes are optionally verified to prevent silent divergence across runs due to external changes in degradation parameters.

5.4.3 Spectral Configuration Block

We record the exact STFT configuration used to compute `Y_stft_mag` and `mel_spec`:

```
"stft_cfg": {
  "n_fft": 2048,
  "hop_length": 512,
  "window": "hann",
  "n_freq_bins": 1025,
  "n_time_frames": 258
}
```

This enables deterministic regeneration of spectrograms and allows future models to adaptively modulate loss weights based on frequency resolution or temporal stride.

5.4.4 Metadata Dictionary

Each file contains a nested `metadata` dictionary that logs:

- `sample_rate`: Always 48000 Hz; *Note: Original dataset audio is resampled to 22.05 kHz before EnCodec conversion to 48 kHz tokens.*
- `bandwidth`: e.g., 24.0 kbps for 16 codebooks at 75 fps
- `n_q`: Number of EnCodec codebooks (usually 16)
- `channels`: Mono (1) or stereo (2)
- `processed_at`: ISO 8601 timestamp
- `version`: Current format version (e.g., "1.1")
- `stage`: Curriculum stage name
- `degradations`: Full parameter dump of applied effects

This structure supports full audit trails, version-controlled experimentation, and compatibility with future loss functions or token formats.

5.4.5 Example Entry

A simplified representation of a valid file might include:

```
{
  'X': LongTensor [16, 4542],
  'Y': LongTensor [16, 4542],
  'scales': FloatTensor [4542],
  'Y_stft_mag': FloatTensor [16, 1025, 258],
  'mel_spec': FloatTensor [128, 258],
```



```

    'metadata': {
        'sample_rate': 48000,
        'bandwidth': 24.0,
        'stage': "stage4_stronger",
        'degradations': {...},
        'original_audio_hash': "...",
        ...
    }
}

```

5.4.6 Training Integration via TokenPairDataset

The `TokenPairDataset` class dynamically loads these bundles during training. Each entry provides token sequences, auxiliary targets for all five restoration heads, and optional waveform references for evaluation or debugging. Files that fail validation are logged and skipped. Loss functions automatically adapt based on the available keys in each bundle, allowing flexible feature expansion without retraining the dataset loader.

Summary: The `.pt` file format is a modular, hash-verified, and future-proof design that enables fast curriculum-aware training, reproducible degradation tracking, and flexible multi-loss optimization across our Token U-Net pipeline.

5.5 VISUALIZING TOKEN USAGE

Understanding the behavior of EnCodec token distributions across degraded and clean audio samples is critical for both diagnosing model performance and motivating our discrete-domain architecture in the next chapter [63]. Since our model operates entirely in token space—not waveforms or spectrograms—we must develop quantitative

and visual tools to assess degradation effects and restoration capacity.

5.5.1 Histograms of Token Distributions

Each EnCodec token sequence $\mathbf{X} \in \mathbb{Z}^{n_q \times T}$ consists of $n_q = 16$ codebook rows over T time frames. Each entry is an integer in $\{0, \dots, K - 1\}$ where $K = 1024$ is the number of centroids per codebook.

For each codebook q , we define a histogram $\mathbf{h}_q$ over token occurrences:

$$\mathbf{h}_q[k] = \sum_{t=1}^T \mathbb{I}[X_q[t] = k], \quad k = 0, \dots, K - 1$$

Clean tracks typically exhibit high-entropy, broad histograms—tokens drawn from a wide and semantically rich vocabulary. Degraded audio, in contrast, often collapses to low-entropy distributions concentrated on a narrow subset of indices, especially under extreme EQ, gain flattening, or overcompression.

Token Coverage Metric. We define per-codebook token coverage as:

$$\gamma_q = \frac{|\{k : \mathbf{h}_q[k] > 0\}|}{K}$$

Coverage γ_q quantifies token diversity per codebook. Healthy signals often exceed $\gamma_q > 0.6$, while collapsed distributions may fall below 0.2. This diagnostic metric is plotted as a 16-bin bar chart during validation to track distributional health across training stages.

5.5.2 Token Heatmaps

To visualize temporal structure, we plot the token matrix $\mathbf{X} \in \mathbb{Z}^{n_q \times T}$ as a 2D heatmap: rows correspond to quantizer indices ($q = 1 \dots 16$), and columns to time frames ($t = 1 \dots T$).

For clean audio, we expect:

- Smooth horizontal variation (local continuity),
- Diverse token paths across codebooks (signal richness),
- Inter-codebook structure correlating with harmonic content.

In degraded audio, artifacts manifest as horizontal banding (frozen tokens), vertical repetition (frame duplication), or discrete jumps (abrupt transients). These pathologies are often visible in the token domain before they manifest in waveforms.




Figure 5.7: Token heatmaps for the same track before (left) and after (right) degradation. Repetition and banding in the degraded version signal token collapse and loss of structure.

5.5.3 Token Entropy and KL Divergence

To formalize representational collapse, we define the entropy of each histogram $\mathbf{h}_q$:

$$H_q = - \sum_{k=0}^{K-1} p_q[k] \cdot \log_2 (p_q[k] + \epsilon), \quad p_q[k] = \frac{\mathbf{h}_q[k]}{T}$$

where ϵ is a small constant for numerical stability. High H_q reflects rich, expressive use of the token vocabulary. Degradations (e.g., EQ cutoff, clipping, or echo feedback) typically reduce H_q disproportionately in high-frequency codebooks.

We also track divergence between clean and degraded distributions using Kullback–Leibler divergence:

$$D_{\text{KL}}(\mathbf{h}_q^{\text{clean}} \parallel \mathbf{h}_q^{\text{deg}}) = \sum_{k=0}^{K-1} p_q^{\text{clean}}[k] \cdot \log_2 \left(\frac{p_q^{\text{clean}}[k] + \epsilon}{p_q^{\text{deg}}[k] + \epsilon} \right)$$

This scalar quantity per codebook captures how far the degraded token usage has drifted from the reference distribution. It is especially useful for inspecting learned restorations during inference.

5.5.4 Training Interpretability Roles

These token-domain metrics serve multiple real-time training diagnostics:

1. **Plateau Diagnosis:** If loss stagnates but token histograms remain collapsed, the model is underfitting or learning a degenerate mapping.
2. **Architecture Evaluation:** Larger models (e.g., Token U-Net with CBAM/FiLM) typically expand token entropy more effectively than shallow baselines.
3. **Perceptual Correlation:** Recovered entropy and token coverage often correlate with improved PESQ/STOI, even when waveform losses show minimal change.

5.5.5 Justification for Token-Space Modeling

These analyses reinforce the decision to cast the restoration task in token space. Unlike waveform or mel losses, which can obscure distributional pathologies, token-level visualizations make sparsity, collapse, or unnatural transitions explicitly measurable.

In the next chapter, we present the *Token U-Net*, an architecture specifically designed to operate over these quantized token maps. Its gated skip connections, CBAM/FiLM

layers, and multi-head outputs aim to restore structural balance across time, frequency, and perceptual dimensions. The tools introduced here enable us to verify whether those restorations are effective—not just in output audio, but in the representational core of the model itself.

CHAPTER 6

THE TOKEN U-NET ARCHITECTURE

6.1 OVERVIEW

This chapter introduces the *Token U-Net*, a custom-designed architecture that operates over discrete latent audio representations obtained via EnCodec. It marks a pivotal shift from traditional spectrogram-based restoration models to fully token-based learning, enabling robust, phase-aware enhancement of degraded musical audio. Note that while training is conducted on 22.05 kHz audio inputs, EnCodec tokenization and decoding occur at 48 kHz to preserve high-fidelity reconstruction during inference. Unlike previous systems constrained by lossy magnitude spectra or stem access, Token U-Net is built to predict clean EnCodec token sequences [28] directly from their degraded counterparts. This section outlines the overall data flow, dimensional transformations, and rationale behind the core design principles.

6.1.1 Motivating a Token-Level Architecture

Previous chapters (especially Chapters 4 and 5) established the limitations of audio enhancement systems operating in the waveform or spectrogram domain. Specifically, spectrogram inversion methods like Griffin–Lim introduced severe phase distortions, and HiFi-GAN struggled to generalize across polyphonic musical textures, resulting in

brittle and hallucinated reconstructions. Moreover, traditional systems often suffer from entangled artifacts when trained on full mixes, lacking source separation or explicit structure.




Figure 6.1: Recap token vs. spectral advantage

The move to tokens addresses these issues. EnCodec provides a powerful discrete bottleneck over high-fidelity waveforms by compressing audio into n_q quantized token streams drawn from a vocabulary of size K [51]. Each token stream represents the output of a separate residual vector quantizer (RVQ) codebook, sampled at a temporal resolution determined by the model’s stride. By learning to restore degraded token sequences $X \in \mathbb{N}^{n_q \times T}$ to their clean counterparts $Y \in \mathbb{N}^{n_q \times T}$, we shift the task from analog waveform correction to symbolic sequence transformation—a regime where deep learning architectures like U-Nets can thrive.

6.1.2 End-to-End Objective

Given a degraded token sequence X , the primary goal is to predict the logits $\hat{Y} \in \mathbb{R}^{K \times n_q \times T}$ over the full token space. During training, a cross-entropy loss [15] is applied for each timestep t and each codebook q :

$$\mathcal{L}_{\text{CE}} = \frac{1}{n_q T} \sum_{q=1}^{n_q} \sum_{t=1}^T \text{CE}(\hat{Y}_{:,q,t}, Y_{q,t}), \quad (6.1)$$

where $\text{CE}(\cdot)$ is the standard categorical cross-entropy loss over token logits. This formulation allows for precise probabilistic modeling of every individual token position, enabling granular supervision and rapid convergence on aligned token pairs.

6.1.3 Input and Output Dimensions

To handle inputs of shape $X \in \mathbb{N}^{n_q \times T}$, the Token U-Net begins with a two-step embedding and projection pipeline:

- Each codebook index in X is first embedded in a dense vector [191] of dimension $d_e = \lceil \frac{d_b}{n_q} \rceil$, resulting in embeddings of shape $[B, n_q, T, d_e]$.
- These embeddings are then flattened and passed through a 1×1 convolution to match the desired base dimension d_b , producing $[B, d_b, T]$ — the input shape for the U-Net core.

The final output is a collection of n_q independent prediction heads - one per codebook - each of which outputs K logits per timestep. The final tensor has shape $[B, K, n_q, T]$, matching the number of possible tokens per codebook.

6.1.4 Token Flow Through the U-Net

The U-Net follows a classic encoder–bottleneck–decoder structure, adapted for 1D sequence modeling. The flow can be summarized as follows.

1. **Embedding:** Each token is embedded and projected in a continuous feature space of dimension d_b .

2. **Encoder:** At each of the L encoder levels, the model applies two ResBlocks, a CBAM module [191], and then downsamples using a stride-2 convolution. This halves the temporal resolution and doubles the channel count.
3. **Bottleneck:** A central block applies FiLM-based modulation [149], dropout regularization, and a temporal context layer (dilated Conv1d, combining inspirations from Nercessian et al. and Li et al.) to account for dereverberation [141] and echo tails [117].
4. **Decoder:** The decoder symmetrically upsamples using transposed convolutions, followed by CBAM, ResBlocks, and gated skip connections that blend encoder features.
5. **Prediction:** Final outputs are passed through the 16 token heads and 5 auxiliary heads for effect parameter handling:
 - n_q token heads to predict logits over the vocabulary.
 - A *mask head* to guide denoising and dereverberation.
 - A *gain head* for global loudness estimation.
 - A *perceptual head* to regress high-level EQ-related parameters.
 - A *compression head* for statistics related to dynamics [131].
 - A *stereo head* estimating the stereo width and interchannel phase [84].

6.1.5 Skip Connection Gating and Token Alignment

Unlike conventional U-Nets, where skip connections are added directly, this architecture introduces *learnable scalar gates* σ_i [169, 184] in each encoder-decoder skip pathway:

$$x_{\text{dec}}^{(i)} \leftarrow x_{\text{dec}}^{(i)} + \sigma_i \cdot \text{crop}(x_{\text{enc}}^{(i)}), \quad (6.2)$$

where $\sigma_i \in (0, 1)$ is a learned scalar updated via backpropagation. These gates enable the model to adaptively attenuate or amplify the importance of encoder features, improving generalization, and preventing oversmoothing or overreliance on shallow layers.

6.1.6 Sequence Characteristics and Token Receptive Fields

Each EnCodec token corresponds to approximately 640 waveform samples at a sampling rate of 48 kHz, the resolution used during tokenization and decoding. However, during training, our model operates on data downsampled to 22.05 kHz. Thus, each token aligns with 293 waveform samples in the training domain. This mismatch is accounted for during inference by decoding tokens back to 48 kHz waveforms, preserving audio detail while reducing training load. With a U-Net depth $L = 4$, the receptive field covers roughly 48 token timesteps, which corresponds to ~ 0.64 seconds at 22.05 kHz training time or ~ 1.4 seconds at 48 kHz inference.[10], encompassing increasingly long-term context:

$$R = 2^L \cdot k, (6.3)$$

where k is the effective kernel size per layer (here 3). With $L = 4$, the receptive field covers roughly 48 token timesteps, corresponding to ~ 0.64 seconds of audio — sufficient to capture cross-transient artifacts like reverb tails, compression modulations, or spectral smearing.

6.1.7 Example Input/Output Flow

For a concrete example, consider a batch of 2 audio clips ($B = 2$), tokenized into 16 EnCodec streams [51] ($n_q = 16$), each of length $T = 4546$ tokens (derived from 30 seconds of 22.05 kHz audio). Note that token decoding will upsample back to 48 kHz,

producing higher-fidelity output without needing to train at that resolution. The architecture processes this input:

- Input shape: $X \in \mathbb{N}^{2 \times 16 \times 4546}$
- Embedding $\rightarrow$ Projection: $[2, 384, 4546]$
- After 4 downsampling levels: $[2, 6144, 284]$
- After decoding: $[2, 384, 4546]$
- Final prediction: $\hat{Y} \in \mathbb{R}^{2 \times 1024 \times 16 \times 4546}$

Additional outputs include:

- **mask** $\in \mathbb{R}^{2 \times 16 \times 4546}$ for soft enhancement control,
- **perceptual** $\in \mathbb{R}^{2 \times 8}$ for EQ/inference features,
- **gain** $\in \mathbb{R}^{2 \times 1}$ for loudness normalization,
- **compression** $\in \mathbb{R}^{2 \times 2}$ for dynamics control,
- **stereo** $\in \mathbb{R}^{2 \times 2}$ for stereo field adjustments.

6.1.8 Transition to Component-Level Explanation

Although this section provided a macrolevel view of the Token U-Net architecture and token processing flow, the next section will dissect its constituent modules - including CBAM, FiLM, ResBlocks, and learnable skip connections — and justify their inclusion based on both empirical results and architectural best practices in sequence modeling.

6.2 U-NET ARCHITECTURAL COMPONENTS

The Token U-Net integrates a suite of advanced architectural modules designed to operate effectively on discrete audio tokens. Unlike traditional spectrogram- or waveform-based U-Nets, which struggle to represent high-dimensional audio dynamics in a lossy domain, our model works directly in the EnCodec token space, requiring precise control over sequence transformation. To that end, we integrate six core components: Residual Blocks, CBAM Attention, Feature-wise Linear Modulation (FiLM), Learnable Gated Skip Connections, Temporal Context Blocks, and an Optional Bottleneck. Each module is mathematically formalized and supported by prior research in deep sequence modeling.

6.2.1 Residual Blocks with GroupNorm and GELU

Inspired by ResNet [73] and adapted for audio [138], each residual block maps an input tensor $\mathbf{x} \in \mathbb{R}^{C \times T}$ to an output via:

$$\mathbf{y} = \mathbf{x} + \mathcal{F}(\mathbf{x}), \quad (6.4)$$

where $\mathcal{F}$ is a nonlinear transformation defined as:

$$\mathcal{F}(\mathbf{x}) = \text{GELU}(\text{GN}(\text{Conv1D}(\text{Drop}(\text{GELU}(\text{GN}(\text{Conv1D}(\mathbf{x})))))), \quad (6.5)$$

with GroupNorm (GN) applied using 8 groups for stability in small batches [192]. Dropout is injected for regularization and GELU improves convergence over ReLU.

6.2.2 Convolutional Block Attention Module (CBAM)

The CBAM [191] module enhances intermediate features by applying sequential *channel attention* and *spatial (temporal) attention*.




Figure 6.2: The composite Convolution Block Attention Module) is a simple yet effect attention block for feedforward convolutional neural network as proposed by Woo et. al [191]

Channel attention is computed via global average pooling followed by a two-layer perceptron:

$$\mathbf{z}_c = \text{GAP}(\mathbf{x}) = \frac{1}{T} \sum_{t=1}^T \mathbf{x}_{:,t} \quad (6.6)$$

$$\mathbf{w}_c = \sigma(W_2 \cdot \text{ReLU}(W_1 \cdot \mathbf{z}_c)) \in (0, 1)^C \quad (6.7)$$




Figure 6.3: The channel attention block as discussed in *Attention is all you Need* Vaswani et. al [183] [191]

Temporal attention compresses across channels via both average and max projections:

$$\mathbf{z}_s = \sigma(\text{Conv1D}([\text{Avg}(\mathbf{x}); \text{Max}(\mathbf{x})])) \in (0, 1)^T \quad (6.8)$$




Figure 6.4: Inclusion of temporal/spatial attention as part of CBAM contributes greatly to the high-fidelity de-reverberation, dynamic range compression, and echo cancellation in support of the bottleneck temporal block [141, 114]

The final attention-modulated output is given by:

$$\mathbf{x}' = \mathbf{x} \odot \mathbf{w}_c \odot \mathbf{w}_s, \quad (6.9)$$

where $\odot$ denotes element-wise multiplication with appropriate broadcasting.

6.2.3 Feature-wise Linear Modulation (FiLM)

FiLM layers modulate features using learned affine parameters [149]:




Figure 6.5: FiLM as proposed by Perez et. al assists in passing extra encoding information

$$\text{FiLM}(\mathbf{x}) = \gamma \cdot \mathbf{x} + \beta, \quad (6.10)$$

where $\gamma, \beta \in \mathbb{R}^{C \times 1}$ are broadcast along the time axis. FiLM allows context-sensitive control of decoder dynamics based on encoded features, analogous to conditioning layers in GANs or multimodal fusion [149].

6.2.4 Learnable Scalar Skip Gates

To balance local detail and global abstraction, we use *gated skip connections* [184]:

$$\mathbf{x}_{\text{dec}}^{(i)} \leftarrow \mathbf{x}_{\text{dec}}^{(i)} + \sigma(g_i) \cdot \text{Crop}(\mathbf{x}_{\text{enc}}^{(i)}) \quad (6.11)$$

Each g_i is a scalar parameter learned via backpropagation, passed through a sigmoid to constrain it to $(0, 1)$. Cropping ensures temporal alignment when encoder and decoder features have mismatched lengths due to strided convolutions.




Figure 6.6: U-Net architecture **prior** to integration of parameter inversion techniques, 741M vs. 1.08B parameters with integration of auxiliary heads and temporal block

6.2.5 Temporal Context Block (Dilated Convolution)

To support reverberant decay modeling and echo patterns, we include a **dilated temporal context block** [84] in the bottleneck:

$$\mathbf{x}_{\text{context}} = \text{Conv1D}(\mathbf{x}; k = 9, \text{dilation} = 2, \text{padding} = 8), \quad (6.12)$$

which expands the effective receptive field without additional layers. This allows mid-layer features to model long-range dependencies, such as delay artifacts or tail energy.

6.2.6 Optional Bottleneck Layer

An optional 1×1 convolution is included after the midblock for redundancy suppression [74]:

$$\mathbf{x}_{\text{bottleneck}} = \text{Conv1D}_{1 \times 1}(\mathbf{x}) \quad (6.13)$$

This is activated only in stage 3+ of curriculum training, where latent bottlenecking improves generalization by eliminating shallow shortcut dependencies.

6.2.7 Summary of Shape Transformations

Let the input tokens be $\mathbf{X} \in \mathbb{N}^{B \times n_q \times T}$, where $n_q = 16$, $T \sim 4542$, and the vocabulary size $K = 1024$. Then:

The Token U-Net architecture includes specialized output heads designed to jointly estimate five audio effects from a single stream of degraded EnCodec tokens. Although the primary objective is discrete token prediction via cross-entropy loss, the model also learns to generate auxiliary signals that capture perceptual cues, gain structure, stereo imaging, dynamic compression, and token masking. These outputs enable the

network to implicitly reverse reverberation, echo, distortion, and compression effects - even when such degradations are entangled in a flattened stereo waveform.

6.2.8 Primary Logit Prediction

Given an input token tensor $X \in \mathbb{N}^{B \times n_q \times T}$, the network produces logits $\hat{Y} \in \mathbb{R}^{B \times K \times n_q \times T}$ where K is the token vocabulary size (e.g., 1024). This is optimized via cross-entropy:

$$\mathcal{L}_{\text{CE}} = \frac{1}{Bn_qT} \sum_{b,q,t} \text{CE}(\hat{Y}_{b,:,q,t}, Y_{b,q,t}), \quad (6.14)$$

providing direct supervision over the predicted token sequence $\hat{Y}$.

6.2.9 Perceptual Spectrum Head

To guide perceptual fidelity, we use a *mel-spectral prediction head* [191, 51]. The internal U-Net feature map $x \in \mathbb{R}^{B \times C \times T}$ is processed by a 1×1 convolution to regress mel-spectrogram estimates:

$$\hat{M} = \text{Conv1D}_{1 \times 1}(x) \in \mathbb{R}^{B \times D \times T}, \quad (6.15)$$

where D is the number of mel bins (typically 80). A spectral loss compares $\hat{M}$ against ground-truth mel spectrograms M extracted from clean waveforms:

$$\mathcal{L}_{\text{mel}} = \|\log(\hat{M} + \epsilon) - \log(M + \epsilon)\|_1. \quad (6.16)$$

This head encourages spectral matching across all frequencies, helping to reduce coloration introduced by reverb or EQ.

6.2.10 Gain Head

To estimate overall signal gain — crucial for handling clipped or quiet recordings — the gain head performs global average pooling followed by a linear projection:

$$g = \text{Linear}(\text{GAP}(x)) \in \mathbb{R}^{B \times 1}, \quad (6.17)$$

where g_b denotes the predicted relative gain in dB. During training, this is compared to a target value g^* computed from RMS-normalized loudness:

$$\mathcal{L}_{\text{gain}} = \frac{1}{B} \sum_{b=1}^B (g_b - g_b^*)^2. \quad (6.18)$$

This enables correction of low or high gain distortions and supports level matching across datasets.

6.2.11 Compression Head (Dynamics Estimation)

Compression artifacts distort dynamics. To estimate them, we predict two statistical descriptors: crest factor and RMS deviation. The compression head uses temporal pooling + projection:

$$c = \text{Linear}(\text{GAP}(x)) \in \mathbb{R}^{B \times 2}, \quad (6.19)$$

representing:

- c_1 : estimated RMS deviation σ_{RMS}
- c_2 : crest factor $\frac{\max |x|}{\text{RMS}(x)}$

These are compared to statistics from clean audio to form:

$$\mathcal{L}_{\text{comp}} = \|c - c^*\|_2^2, \quad (6.20)$$

helping the network identify and reverse over-compression or flat dynamics.

6.2.12 Stereo Width and Phase Head

The stereo head estimates phase correlation and spatial width. After channel-agnostic pooling:

$$s = \text{Linear}(\text{GAP}(x)) \in \mathbb{R}^{B \times 2}, \quad (6.21)$$

encoding:

- s_1 : predicted interchannel phase coherence
- s_2 : stereo width (variance between L/R channels)

This supports stereo recovery after mono collapse or excessive mid-side imbalance.

Ground truth is computed from stereo audio via:

$$s_1^* = \frac{\text{Re}(\text{STFT}_L \cdot \overline{\text{STFT}_R})}{\|\text{STFT}_L\| \|\text{STFT}_R\|}, \quad s_2^* = \text{Var}(L - R) \quad (6.22)$$

and we minimize:

$$\mathcal{L}_{\text{stereo}} = \|s - s^*\|_2^2. \quad (6.23)$$

6.2.13 Mask Head (Token Drop Prediction)

The mask head predicts a soft confidence score $\alpha_t \in [0, 1]$ for each token, indicating whether it likely corresponds to noise, echo, or artifact. This is implemented via:

$$\alpha = \text{Sigmoid}(\text{Conv1D}_{1 \times 1}(x)) \in \mathbb{R}^{B \times 1 \times T}. \quad (6.24)$$

These scores act as soft masks over token confidence, akin to denoising masks in SEGAN [147] or DeepVQE, and can be thresholded or applied as weights:

$$\hat{Y}_{b,:,q,t} \leftarrow \alpha_{b,1,t} \cdot \hat{Y}_{b,:,q,t} \quad (6.25)$$

6.2.14 Total Loss Function

All auxiliary heads are combined with weighted losses:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{CE}} + \lambda_1 \mathcal{L}_{\text{mel}} + \lambda_2 \mathcal{L}_{\text{gain}} + \lambda_3 \mathcal{L}_{\text{comp}} + \lambda_4 \mathcal{L}_{\text{stereo}}, \quad (6.26)$$

with λ_i tuned per curriculum stage. Mask losses are optionally included for noisy stages only.

6.2.15 Theoretical Justification for Multi-Effect Recovery

Unlike single-purpose models, Token U-Net is designed to restore audio altered by **multiple entangled degradations simultaneously**. Each head supervises a different axis of restoration:

- $\mathcal{L}_{\text{CE}}$: symbolic structure (e.g., pitch, rhythm, articulation) - $\mathcal{L}_{\text{mel}}$: frequency content and perceptual match - $\mathcal{L}_{\text{gain}}$: loudness correction - $\mathcal{L}_{\text{comp}}$: dynamic range preservation - $\mathcal{L}_{\text{stereo}}$: spatial fidelity - α : denoising and dereverberation mask

By grounding each in explicit statistical or perceptual quantities, the network learns disentangled internal representations for multifaceted and quality restoration. As shown in DeepVQE [84] and Henkel et al. [74], such joint prediction frameworks are essential for tackling real-world audio artifacts that defy single-effect modeling.




Figure 6.7: A full diagram of Token U-Net architecture with temporal context block in bottleneck, CBAM, FiLM, auxiliary heads, etc.

6.3 COMPARISON TO OTHER U-NET VARIANTS

Throughout the design of the Token U-Net, we systematically evaluated several U-Net variants drawn from both academic literature and practical implementations. Each alternative offered unique strengths for specific types of degradation, such as denoising, dereverberation, or stereo enhancement, but none matched the combination of breadth, efficiency, and token-space compatibility of Token U-Net. This section com-

pares our architecture against six representative baselines, with each model trained under matched conditions on RTX 6000 Ada GPUs using curriculum-aligned datasets. Metrics such as parameter count, reconstruction quality, memory usage, and generalizability across effects were analyzed.

6.3.1 Deep Residual U-Net (5-Layer)

A natural starting point was a 5-level deep U-Net inspired by the canonical architecture of Ronneberger et al. [165], extended to token sequences via strided convolutions:

$$x^{(l)} = \text{Conv1D}_{s=2} \circ \sigma \circ \text{BN} \circ \text{Conv1D}(x^{(l-1)}), \quad l \in \{1, \dots, 5\}. \quad (6.27)$$

Skip connections between encoder and decoder paths helped preserve the local structure. Although expressive, this model overfits quickly due to limited data ($\sim 25\text{k}$ examples) and short sequence lengths ($T \approx 4542$ tokens). Despite achieving low training loss, validation PESQ and STOI saturated early. GPU memory usage peaked at 34–36 GB with batch sizes ≤ 6 , limiting scalability.

6.3.2 Token Residual U-Net (No Attention)

Inspired by the residual mixing models in [105, 174], we replaced plain convolutions with residual blocks:

$$x' = x + \mathcal{F}(x), \quad \mathcal{F}(x) = \text{GELU} \circ \text{GN} \circ \text{Conv1D}(x), \quad (6.28)$$

which provided stability during curriculum transitions and mitigated vanishing gradients. However, models lacking attention mechanisms struggled to correct long-term degradations such as echo tails or compression modulations. PESQ improved over baseline U-Nets, but stereo field collapse and spectral smearing remained common.

6.3.3 Attention-Only U-Nets

To test the inverse hypothesis, we removed residuals and implemented additive attention mechanisms:

$$\text{Attn}(Q, K, V) = \text{softmax}\left(\frac{QK^\top}{\sqrt{d}}\right)V. \quad (6.29)$$

However, full attention layers proved unstable for token input, as quantized values lack the smooth semantics that attention mechanisms typically exploit. Exploding gradients and NaNs were frequent, especially for long input sequences. Unlike CBAM, these layers lacked hierarchical channel/temporal selectivity and offered no regularization of spatial salience.

6.3.4 Residual + Naive Attention (No CBAM, No FiLM)

Inspired by early music separation hybrids like [174] and compression predictors [124], we combined residual blocks with basic attention, omitting CBAM and FiLM. Although training was stable and convergence was fast, generalization lagged. Attention modules lacked adaptive modulation and failed to reweight decoder pathways based on latent cues. Compared to our final architecture, these hybrids scored 15% lower in SNR and required > 20% more memory.

6.3.5 Two-Stage U-Net (Token + Spectral)

Drawing from Moliner and Välimäki’s high-fidelity restoration pipeline [134, 87], we implemented a cascade of two U-Nets: the first over tokens, the second over mel spectrograms derived from preliminary reconstructions. The theoretical appeal lay in separating coarse correction from perceptual refinement. However, a domain mismatch between quantized token output and spectral input required intermediate CPU decoding and remel synthesis. Training time doubled, GPU utilization dropped by

42%, and minor gains in spectral metrics came at the cost of double complexity. The pipeline also exhibited perceptual inconsistency due to divergent loss landscapes between stages.

6.3.6 2D Spectrogram U-Nets

Several early trials adapted image-based U-Nets [165] to 2D mel spectrograms, treating audio as $\mathbf{x} \in \mathbb{R}^{B \times F \times T}$. Models performed well on individual effects (e.g., denoising), as seen in [172], but ultimately failed to scale. Spectrogram inversion (via Griffin–Lim or HiFi-GAN [51]) introduced phase artifacts, blurred onsets, and spectral distortions. Unlike token-space architectures, these systems could not jointly model phase-aware features, stereo balance, or loudness structure.

6.3.7 LSTM Regression Heads (Parameter Estimation)

Inspired by ML-forecasting models for compressor control [22], we appended LSTM heads to the decoder features to estimate the degradation parameters:

$$\hat{\Theta}_t = \text{LSTM}(f_t^{\text{decoder}}), \quad \hat{y}_t = \text{InverseEffect}(x_t; \hat{\Theta}_t). \quad (6.30)$$

These models learned to regress RMS thresholds, EQ curves, or reverb decay times. However, despite low parameter loss, the reconstructions degraded due to the misalignment between waveform-based effects and token-space learning. The approach violated the principle of representational consistency, with degradation parameters existing in a space disconnected from token predictions.

6.3.8 Memory Footprint and Generalization Tradeoffs

Table 6.1 summarizes our findings. In particular, our final architecture—a 4-layer combination of CBAM, FiLM, residual blocks, temporal blocks, and gated concate-

nations—achieved strong multi-effect restoration with a rather small size (many modern models achieve 9+ layers) and no decoding bottlenecks. Unlike deep or two-stage models, it sustained batch sizes of up to 16 (8 on average) on 40 GB VRAM and converged 38% faster than spectrogram-based counterparts.

Table 6.1: Comparison of U-Net variants across metrics (higher is better unless noted).

Model Variant	PESQ↑	STOI↑	SNR↑	GPU Mem (GB)↓
5-layer Deep U-Net	2.91	0.841	18.2	35.4
Residual-only (no attn)	3.03	0.850	19.1	29.8
Attention-only	NaN	—	—	>40
Residual + Naive Attn	3.05	0.852	20.1	32.7
2-Stage Token + Spectral	3.10	0.857	20.8	39.2
2D Spectrogram U-Net	2.87	0.837	16.5	28.9
Token U-Net (ours)	3.22	0.872	21.6	22.4

6.3.9 Representation Differences: Spectrograms vs. tokens

As discussed in Chapter 5, spectrogram representations require reconstruction stages and suffer from phase loss. Tokenization with EnCodec instead discretizes waveforms using quantized latent codes:

$$\mathbf{z}_q = \text{Quant}_q(\phi(x[t : t + \Delta])), \quad q \in \{1, \dots, n_q\}, \quad (6.31)$$

where ϕ is a learned encoder. Table 6.1 revisits this contrast. Tokens offer phase preservation, lower memory, and direct decodability, making them well suited for our architecture restoration pipeline.

6.3.10 Conclusion

Across all architectural comparisons, Token U-Net emerged as the most practical and extensible design. It balances expressive power with computational efficiency, offers modular auxiliary heads for multi-effect estimation, and avoids costly inversion

steps. While models like [134] or [172] excel in specific, single degradations, they suffer from domain coupling, memory overhead, or phase loss when faced with a variety. Our architecture, although not as strong in any one effect degradation, unifies token restoration with perceptual modeling in a fully differentiable and efficient multi-purpose system.

6.4 MODEL SIZE, PARAMETER COUNTS, AND DEPTH VARIATIONS

A major achievement of Token U-Net is its ability to maintain high restoration performance across various degradations, reverb, echo, compression, and gain loss, while operating directly on compressed EnCodec token sequences. Unlike traditional spectrogram-based pipelines that depend on intermediate waveform reconstruction, our system processes raw token streams end-to-end with highly scalable convolutional layers. This section details the architectural scaling strategy, culminating in a mere billion-parameter model that remains efficient during training and inference.

6.4.1 Final Parameter Count and Dimensional Flow

The finalized Token U-Net architecture—configured with a base dimension of 384, a depth of 4 layers, and multi-head outputs for 16 EnCodec codebooks—contains **1,080,885,553 parameters**, as verified during runtime (see Section 6.2). The model processes inputs of shape $\mathbf{X} \in \mathbb{N}^{B \times 16 \times T}$, where $T \sim 4542$, and outputs logits of shape $\hat{\mathbf{Y}} \in \mathbb{R}^{B \times 1024 \times 16 \times T}$. These tokens are extracted from 22.05 kHz audio clips during training, but decoded at 48 kHz at inference, leveraging EnCodec’s high-fidelity token stream for waveform reconstruction.

The dimensional flow is as follows:

- **Embedding:** $[B, 16, T] \rightarrow [B, 384, T]$ via learned token embeddings and a projection layer.
- **Encoder stages** (with channel doubling and halved time resolution):

$$E_1 : [384 \rightarrow 768] \quad (T \rightarrow T/2)$$

$$E_2 : [768 \rightarrow 1536] \quad (T/2 \rightarrow T/4)$$

$$E_3 : [1536 \rightarrow 3072] \quad (T/4 \rightarrow T/8)$$

$$E_4 : [3072 \rightarrow 6144] \quad (T/8 \rightarrow T/16)$$

- **Mid Block:** $[6144 \rightarrow 6144]$ (with CBAM, FiLM, residuals, temporal context, and optional bottleneck).
- **Decoder stages** (with upsampling and gated skip connections):

$$D_1 : [6144 \rightarrow 3072] \quad (T/16 \rightarrow T/8)$$

$$D_2 : [3072 \rightarrow 1536] \quad (T/8 \rightarrow T/4)$$

$$D_3 : [1536 \rightarrow 768] \quad (T/4 \rightarrow T/2)$$

$$D_4 : [768 \rightarrow 384] \quad (T/2 \rightarrow T)$$

- **Heads:** 16 shared-weight 1×1 convolutional heads, each mapping $[384 \rightarrow 1024]$, producing the full output tensor.

This structure supports parallelized head computation per codebook, enabling efficient softmax and cross-entropy operations on each quantized stream.

6.4.2 Why Four Layers Were Optimal

Although 5-layer U-Nets (e.g., with channels $384 \rightarrow 12288$) showed marginal loss improvements, they introduced memory overflows even on 48GB GPUs and incurred

longer training times. Curriculum stages 3 and 4, in particular, triggered instability when the depth of the latent bottleneck exceeded $T/16$.

In contrast, our 4-layer architecture leverages CBAM attention and FiLM conditioning to offset the shallower depth. CBAM learns hierarchical channel/temporal weighting, while FiLM adaptively modulates decoder layers based on latent global context. Combined with learnable skip gates and a temporal context module, this configuration achieves efficient depth-equivalent expressivity at lower compute cost.

6.4.3 Parameter Growth with Base Dimension

Table 6.2: Parameter scaling with base dimension (depth fixed at 4)

Base Dimension	Depth	Total Parameters
256	4	$\sim 290\text{M}$
320	4	$\sim 410\text{M}$
384	4	1.08B
384	5	$>1.7\text{B}$ (OOM in Stage 4)
512	4	$>2.3\text{B}$ (OOM in Stage 3)

As shown in Table 6.2, increasing the base dimension dramatically expands the number of parameters and the demands for memory. The final 384-dim setting represents the optimal trade-off between capacity and feasibility on current GPU infrastructure.

6.4.4 Spectrogram vs. Token Runtime and Memory Footprint

Unlike STFT- or mel-based architectures (which must maintain floating-point tensors of shape $[F \times T]$ with $F \sim 80 - 1025$), token-based models compress these features into 4-bit integer indices. Our inputs have shape $[16 \times T]$, and EnCodec quantization eliminates the need for Griffin-Lim or HiFi-GAN decoding during inference.

Furthermore, logits are only computed in the final step, avoiding recurrent feature maps at large F . This enables forward passes with:

- $> 7\text{--}10\times$ faster token throughput,
- No intermediate decoding or inverse transforms,
- Lower activation memory per batch,
- Minimal loss of perceptual quality, as shown in Chapter 8.

6.4.5 Training and Batch Scaling Performance

All experiments were conducted on the Dartmouth LISP Lab’s NVIDIA RTX 6000 Ada GPUs. Thanks to token precomputation (see Chapter 3), we achieved:

- Batch size 8 in stages 0–2,
- Batch size 4 in stage 3,
- Batch size 2 in stage 4 (full degradation),
- $\sim 85\%$ reduction in training time compared to on-the-fly waveform pipelines.

The total parameter count surpassing one billion—yet remaining fully trainable under curriculum learning—demonstrates the scalability of token-based audio models. While few restoration pipelines to date have approached this scale, we found that proper representation alignment and head design (Section 6.2) enabled smooth convergence and consistent generalization across all effects.

6.4.6 Conclusion

Despite its rather shallow depth-4 structure, the Token U-Net achieves performance on par with or exceeding deeper spectrogram-based models in the music domain due to architectural innovations and efficient token-level computation. Its 1.08B parameter count is the result of parallel multi-head prediction, high-dimensional token embeddings, and channel-rich temporal reasoning. This makes the Token U-Net not only

performant, but theoretically extensible for possible future mastering tools, multitrack inference, or stereo-aware enhancement models at scale.

6.5 FINAL ARCHITECTURE JUSTIFICATION

The token U-Net was not derived from architectural convention, but was carefully engineered to solve a specific class of problems: the restoration of audio signals degraded by *multiple, co-occurring distortions*—including reverb, echo, gain mismatch, compression artifacts, and stereo collapse — within a tokenized representation space. This section synthesizes the key components introduced in Chapter 6, formally justifying the structure of the model in terms of its multi-effect restoration capacity, interpretability, and hardware efficiency.

6.5.1 Modular Components Recap: Designed for Multi-Effect Recovery

The final architecture integrates both classical U-Net mechanisms and novel modules, each tuned to detect and correct specific audio degradations. Table 6.3 summarizes the contribution of each architectural characteristic to the five main types of degradation targeted in this work.

Component	Reverb	Echo	Gain	Compression	Stereo Collapse
Temporal Context Block	✓	✓			
CBAM Attention	✓	✓		✓	✓
FiLM Normalization	✓	✓	✓	✓	✓
Residual Blocks	✓	✓	✓	✓	✓
Gated Skip Connections			✓	✓	✓
Perceptual (Mel) Head	✓	✓	✓		
Gain Head			✓	✓	
Compression Head				✓	
Stereo Head					✓
Mask Head	✓	✓			

Table 6.3: Mapping of architectural components to target degradation classes.

6.5.2 The Role of Auxiliary Heads

In contrast to traditional U-Nets, which focus solely on regression to a denoised waveform or spectrogram, our architecture features five parallel **auxiliary heads** that jointly predict:

- **Token logits** (primary prediction): $[B, 1024, 16, T]$ logits for all codebooks
- **Mel spectrogram** estimates: $[B, 80, T]$ for perceptual fidelity
- **Gain level** estimates: $[B, 1]$ for loudness normalization
- **Compression metrics**: $[B, 2]$ for RMS deviation and crest factor
- **Stereo imaging stats**: $[B, 2]$ for phase coherence and width
- **Token dropout mask**: $[B, 1, T]$ soft mask over noisy or corrupted tokens

Each head is independently supervised by interpretable statistics derived from the clean reference waveform, allowing the network to learn disentangled representations of signal integrity.

6.5.3 Temporal Context Block: Long-Range Dependency Modeling

Inserted directly in the bottleneck of the U-Net, the **temporal context block** consists of a dilated convolutional layer:

$$\mathbf{x}_{\text{context}} = \text{Conv1D}(x; \text{kernel} = 9, \text{dilation} = 2, \text{padding} = 8)$$

This module was specifically added to enable long-range dependency modeling across up to 17 time steps, without increasing model depth. Its design is inspired by the delay-line and tail energy phenomena introduced by reverb and echo, making it

critical for capturing temporal smear and reflection patterns that span hundreds of milliseconds in time.

6.5.4 Curriculum-Aware Modularity

Every architectural element can be toggled or tuned based on curriculum stage:

- **Dropout rate:** increased for strong degradations
- **Bottleneck layer:** activated in stage 4 to suppress shortcut paths
- **Loss weights (λ_i):** dynamically tuned for mel, gain, stereo, and compression losses
- **Head usage:** perceptual head emphasized in early stages; stereo and compression emphasized in later stages

This flexibility allows the network to maintain training stability and avoid overfitting across progressive degradation regimes.

6.5.5 Why This Architecture Succeeds

The final Token U-Net succeeds not merely by avoiding known failures (e.g., overdeepening, attention overload) but by offering a *principled modular framework* where each pathway maps to a specific auditory concept:

1. **Restoration is distributed:** No single decoder block handles "gain" or "reverb"; rather, the model projects those distinctions through specialized heads trained on interpretable features.
2. **Memory is localized:** Skip connections provide high-frequency detail, but gated modulation prevents them from overwhelming decoder pathways, ensuring that reconstruction emerges from both the local and global context.

3. **Latent semantics are aligned:** Temporal context and FiLM act as priors that adjust the meaning of features at runtime, similar to conditional attention in speech enhancement models [84].

6.5.6 From Architecture to Optimization

All components were validated not only through empirical ablation, but also through their compatibility with our training scheduler, discussed in Chapter 7. The model generalizes across *entangled degradations* without requiring discrete detection or classification of each effect, an achievement made possible by token-wise multihead supervision.

As we transition to training procedures, we show how each of these components interacts with the loss function, curriculum stages, and optimization parameters to produce stable, high-fidelity audio reconstructions under real-world restoration and enhancement challenges.

CHAPTER 7

CURRICULUM TRAINING

PIPELINE & LOSSES

7.1 WHY CURRICULUM LEARNING?

The final Token U-Net architecture introduced in Chapter 6 is designed to perform simultaneous restoration of multiple co-occurring audio degradations. These include: reverberation, echo, gain mismatches, compression artifacts, and stereo collapse. However, while the model architecture is theoretically capable of joint enhancement, naive end-to-end training on heavily distorted audio led to unstable convergence, degenerate local minima, and token predictions that did not capture semantic or perceptual correctness.

These early failures exposed a fundamental challenge: **not all degradations are equally learnable in early stage training**, and their effects are often entangled in non additive, nonlinear ways. For instance, stereo collapse removes spatial cues that might otherwise guide echo localization; compression masks gain errors; reverberation flattens dynamic range. Exposing the model to such entanglements before it has internalized basic restoration priors leads to gradient confusion, ineffective supervision, and feature collapse.

To mitigate this, we adopt a structured, pedagogically inspired approach: **curriculum learning**, where degradations are progressively introduced based on their

inherent difficulty and recoverability. Each stage introduces new perturbations while retaining the model’s ability to restore simpler, earlier ones. This mirrors the scaffolding process in cognitive learning and hierarchical task training in reinforcement learning systems [14, 110].

7.1.1 Mathematical Justification: Progressive Distribution Shift

Let $\mathcal{D}_0, \mathcal{D}_1, \dots, \mathcal{D}_N$ denote a sequence of distributions over pairs of degraded tokens, where $\mathcal{D}_0$ corresponds to identity mappings (clean tokens), and each $\mathcal{D}_i$ introduces incrementally more severe degradation types:

$$\theta_0 \xrightarrow{\mathcal{D}_0} \theta_1 \xrightarrow{\mathcal{D}_1} \dots \xrightarrow{\mathcal{D}_N} \theta_{N+1}$$

Here, θ_i is the model parameter state after training on $\mathcal{D}_i$ for sufficient epochs. This training trajectory ensures that each successive distribution $\mathcal{D}_{i+1}$ is introduced only after θ_i has internalized the priors needed to handle it. This is in contrast to flat training:

$$\theta_0 \xrightarrow{\mathcal{D}_N} \theta'_N$$

which suffers from unstable gradients, poor token entropy modeling, and ungrounded auxiliary head supervision, especially during early epochs.

7.1.2 Empirical Failures of Flat Training

In flat training baselines (on $\mathcal{D}_N$ from the outset), we observed:

- **Token collapse:** Logit entropy fell prematurely, leading to repetitive codebook outputs centered around $k \in [0, 100]$.

- **Auxiliary head stalling:** The loss gradients, especially for the echo, stereo, and compression heads, decayed to zero before a meaningful signal structure was learned.
- **Loss explosion:** $\mathcal{L}_{\text{CE}}$ and $\mathcal{L}_{\text{mel}}$ increased rapidly when heavy multi-effect reverb, echo, and compression combinations were introduced early.

These outcomes validated that the restoration of even simple degradations (e.g., uniform gain offsets) was being obstructed by noise interactions from harder effects. The architectural capacity of Token U-Net was not the limiting factor; the training regime was.

7.1.3 Stagewise Design: Linking Degradations to Learning Phases

To align training complexity with perceptual recoverability, we define the following curriculum inspired by:

- **Stage 0 (Identity)** — $\mathcal{D}_0$: Ensures token persistence and logit calibration. Losses target only $\mathcal{L}_{\text{CE}}$ and early spectral matching.
- **Stage 1 (Single FX)** — $\mathcal{D}_1$: Introduces mild single, isolated degradations (e.g., only gain or only reverb). Trains individual heads in clean supervision settings; auxiliary heads are not necessary; temporal block+CBAM is used for reverb and echo scenarios
- *Stage 1 (Single Stronger)* — $\mathcal{D}_1$: Introduces stronger single, isolated degradations halfway through stage 1.
- **Stage 2 (Double FX)** — $\mathcal{D}_2$: Combines orthogonal effects (e.g., gain + compression). Auxiliary losses such as $\mathcal{L}_{\text{gain}}$ and $\mathcal{L}_{\text{comp}}$ are activated.

- **Stage 3 (Triple FX)** — $\mathcal{D}_3$: Introduces echo or reverb alongside dynamic effects. Temporal context modeling and masking become critical. $\mathcal{L}_{\text{stereo}}$ and $\mathcal{L}_{\text{mask}}$ now influence the gradient flow.
- **Stage 3 (Triple Stronger)** — $\mathcal{D}_3$: Auxiliary heads' losses contribute a greater magnitude to training and validation loss.
- **Stage 4 (Full FX Stronger)** — $\mathcal{D}_4$: Exposes the model to real-world mixes with anywhere from one to five simultaneous distortions. All auxiliary heads and spectral losses are activated, and the model learns to generalize seeing full spectrum of audio quality in numerous domains.
- *Stage 4 (Full FX)* — $\mathcal{D}_4$: Increased magnitude of effects halfway through stage 4. Losses ideally fall as the model has learned the majority of the spectrum.

All stages, with the exception of Stages 0 and 2, have a corresponding **_stronger** variant with extended parameter ranges (e.g., longer reverb tails, deeper stereo collapse, aggressive dynamics flattening).

7.1.4 Architectural Synergy with Curriculum Staging

Despite implementing these auxiliary heads, we experimented first without implementing their influences in the loss equation in order to evaluate empirical performance of the token and spectral spaces alone. While results saw rather promising loss curves, the importance of auxiliary heads being quantified into loss became quickly apparent. In both qualitative listening tests and the quantitative metrics we discuss in Chapter 8, we found our inability to invert particularly reverb and echo parameters. These empirical curves are shown below:




Figure 7.1: Promising Loss Curves yet limited ability to invert reverb and echo degradations at inference

Thus, we moved forward and began to experiment with employing *task-specific heads*, each supervised first only when the corresponding degradation is introduced. As expected, this immediately mitigated the reverb and echo issues, and greatly

improved inversion of the other degradation parameters as well. The mathematical foundation is as follows:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{CE}} + \sum_{i=1}^5 \lambda_i \mathcal{L}_i$$

where each $\mathcal{L}_i$ corresponds to: mel spectrum, gain, compression, stereo, and mask.

For example:

- $\mathcal{L}_{\text{gain}}$ is used only after Stage 1.
- $\mathcal{L}_{\text{comp}}$ is introduced in Stage 2.
- $\mathcal{L}_{\text{stereo}}$ is introduced in Stage 3.
- $\mathcal{L}_{\text{mask}}$ is emphasized during Stage 3/4 for token dropout regions.

These decisions were made on calculations based on the comparison between equal magnitudes of different effects; more significant effects were employed later on only when necessary, both to save computation cost and to avoid overfitting [164]. We see an immediate improvement in jsut a 3 second clip of a singer's voice in the following figures 7.2 and 7.3.




Figure 7.2: Left: Clean 3s signal (Right) Reverbed Signal. With just 3 seconds of music, we see the significant temporal and spectral alterations made.




Figure 7.3: Left: Degraded Signal from above right (Right) Restored signal after incorporating auxiliary losses.

This staged supervision allows gradients to specialize early while converging into shared representation spaces in later stages, preserving disentanglement while supporting synergy.

7.1.5 Throughput Gains via Token Preprocessing

Precomputing EnCodec token pairs and spectral statistics (mel, gain, stereo, compression) enabled batch loads of up to 8 on RTX 6000 Ada GPUs. Compared to waveform-based augmentation or spectrogram decoding (e.g., HiFi-GAN inference), this yielded a:

7–10 $\times$ speedup in per-batch training throughput

and eliminated GPU memory waste from on-the-fly audio decoding.

This efficiency was essential to scale training across curriculum stages and evaluate generalization performance while also allowing multiple variants (e.g., with/without bottleneck) to be tested per stage.

7.1.6 Conclusion: A Scaffold for Multi-Effect Restoration

Curriculum learning in this context is not just a data scheduling technique; it is a restoration scaffold. Each stage functions as a lens through which an aspect of the degradation space is explored, internalized, and generalized. By the time the Token U-Net reaches Stage 4, it has learned to localize echo and reverb tails using dilated temporal context; Disentangle gain from compression using joint head prediction; Recover stereo width from collapsed embeddings; and Denoise dropout regions via learned token masking

Thus, curriculum-aware training is not a workaround; it is the philosophical and algorithmic backbone of the entire framework.

7.2 CURRICULUM STAGES & LOSS FUNCTIONS DEFINED

The optimization of Token-UNet across the curriculum requires a composite loss function that evolves in complexity as the model progresses through increasingly difficult degradation scenarios. The general training objective $\mathcal{L}_{\text{total}}$ integrates both discrete token reconstruction and audio-perceptual fidelity, while simultaneously supervising auxiliary heads that predict interpretable signal transformations. This section for-

malizes each loss component and explains their integration and frames them in the context of the curriculum.

7.2.1 Token Cross-Entropy Loss



Figure 7.4: Performance of CE loss graphed independently with all auxiliaries.

The foundational prediction task is token-level sequence modeling. Given the degraded token input $\mathbf{X} \in \mathbb{Z}^{B \times n_q \times T}$, the model predicts logits $\mathbf{Z} \in \mathbb{R}^{B \times K \times n_q \times T}$, where $K = 1024$ is the vocabulary size per codebook. The clean target is $\mathbf{Y} \in \mathbb{Z}^{B \times n_q \times T}$. The token-level cross-entropy loss is:

$$\mathcal{L}_{\text{CE}} = \frac{1}{B \cdot n_q \cdot T} \sum_{b=1}^B \sum_{q=1}^{n_q} \sum_{t=1}^T \text{CE}(\mathbf{Z}_{b,:,q,t}, \mathbf{Y}_{b,q,t})$$

7.2.2 Mask Head Loss: Dereverberation Mask




Figure 7.5: Mask Loss independently graphed.

The dereverberation head predicts a binary soft mask $\hat{\mathbf{M}} \in [0, 1]^{B \times n_q \times T}$ designed to suppress smearing in reverberant tokens. The ground truth mask $\mathbf{M}$ is derived from time-frequency energy differences:

$$\mathcal{L}_{\text{mask}} = \frac{1}{B \cdot n_q \cdot T} \sum_{b,q,t} \text{BCE}(\hat{\mathbf{M}}_{b,q,t}, \mathbf{M}_{b,q,t})$$

7.2.3 Perceptual Head Loss: Quality Embedding




Figure 7.6: Perceptual Head loss independently graphed.

This head predicts a latent perceptual quality vector $\hat{\mathbf{p}} \in \mathbb{R}^{B \times 8}$, used to estimate perceptual degradation severity. A corresponding label $\mathbf{p}$ (e.g., derived from STFT or Mel distance) supervises the output:

$$\mathcal{L}_{\text{perc}} = \frac{1}{B} \sum_{b=1}^B \|\hat{\mathbf{p}}_b - \mathbf{p}_b\|_2^2$$

7.2.4 Gain Head Loss: Volume Restoration




Figure 7.7: Gain head loss graphed independently

The gain head outputs a scalar correction factor $\hat{g}_b \in \mathbb{R}$ per sample. Given target gains g_b , the loss is:

$$\mathcal{L}_{\text{gain}} = \frac{1}{B} \sum_{b=1}^B (\hat{g}_b - g_b)^2$$

7.2.5 Stereo Field Head Loss: Spatial Re-expansion




Figure 7.8: One of our most vital losses graphed independently: stereo field

The stereo head estimates a 2D vector $\hat{\mathbf{s}} \in \mathbb{R}^{B \times 2}$ encoding stereo width and pan location. Given target values $\mathbf{s}_b$, we use:

$$\mathcal{L}_{\text{stereo}} = \frac{1}{B} \sum_{b=1}^B \|\hat{\mathbf{s}}_b - \mathbf{s}_b\|_2^2$$

7.2.6 Compression Head Loss: Dynamics Restoration




Figure 7.9: Compression loss graphed independently

To reverse unwanted compression, the model outputs $\hat{\mathbf{c}} \in \mathbb{R}^{B \times 2}$ (e.g., [compression ratio, threshold]). Given ground truth $\mathbf{c}$:

$$\mathcal{L}_{\text{comp}} = \frac{1}{B} \sum_{b=1}^B \|\hat{\mathbf{c}}_b - \mathbf{c}_b\|_2^2$$

7.2.7 Audio-Domain Losses

When audio reconstruction is enabled (Stages 1+), we decode the predicted tokens $\hat{\mathbf{Y}} = \arg \max \mathbf{Z}$ using EnCodec and compare the output waveform to the original. The losses are:

STFT Spectral Loss:

$$\mathcal{L}_{\text{STFT}} = \frac{1}{B} \sum_{b=1}^B \||\text{STFT}(\hat{y}_b)| - |\text{STFT}(y_b)|\|_2^2$$

Time-Domain Loss:

$$\mathcal{L}_{\text{time}} = \frac{1}{B} \sum_{b=1}^B \|\hat{y}_b - y_b\|_1$$

Mel Spectrogram Loss:

$$\mathcal{L}_{\text{mel}} = \frac{1}{B} \sum_{b=1}^B \|\text{Mel}(\hat{y}_b) - \text{Mel}(y_b)\|_2^2$$

Token-Space Spectral Loss:

$$\mathcal{L}_{\text{tok_spec}} = \frac{1}{B} \sum_{b=1}^B \|\hat{\mathbf{S}}_b - \mathbf{S}_b\|_2^2$$

where $\mathbf{S}_b$ is a cached STFT-magnitude projection from token space.



Figure 7.10: All Audio Domain losses graphed. While not all employed the entire time in aggregate training loss, they were all independently tracked.

7.2.8 Combined Loss Function

The full training loss is:

$$\begin{aligned}
\mathcal{L}_{\text{total}} = & \mathcal{L}_{\text{CE}} + \lambda_{\text{mask}}\mathcal{L}_{\text{mask}} + \lambda_{\text{perc}}\mathcal{L}_{\text{perc}} \\
& + \lambda_{\text{gain}}\mathcal{L}_{\text{gain}} + \lambda_{\text{stereo}}\mathcal{L}_{\text{stereo}} + \lambda_{\text{comp}}\mathcal{L}_{\text{comp}} \\
& + \lambda_{\text{STFT}}\mathcal{L}_{\text{STFT}} + \lambda_{\text{time}}\mathcal{L}_{\text{time}} + \lambda_{\text{mel}}\mathcal{L}_{\text{mel}} + \lambda_{\text{tok_spec}}\mathcal{L}_{\text{tok_spec}}
\end{aligned}$$

with typical values:

- $\lambda_{\text{mask}} = \lambda_{\text{perc}} = \lambda_{\text{gain}} = \lambda_{\text{stereo}} = \lambda_{\text{comp}} = 0.1$
- $\lambda_{\text{STFT}} = 0.1, \lambda_{\text{time}} = 0.1$
- $\lambda_{\text{mel}} = 0.01, \lambda_{\text{tok_spec}} = 0.05$

At each stage, the training loss is defined as a weighted sum. Each term is activated conditionally based on the degradation types introduced in the current stage: - $\mathcal{L}_{\text{CE}}$: Cross-entropy loss across all n_q codebooks and time steps (always active). - $\mathcal{L}_{\text{mel}}$: L1 mel-spectrogram loss between predicted and clean audio (active from Stage 1 onward). - $\mathcal{L}_{\text{STFT}}$: L1 loss on log-magnitude STFTs (used for fine-grained frequency restoration). - $\mathcal{L}_{\text{mask}}$: L1 loss on predicted vs. ideal soft masks over token confidence. - $\mathcal{L}_{\text{gain}}$: L2 loss between the predicted and true dB gain vectors (activated in gain-based stages). - $\mathcal{L}_{\text{comp}}$: L2 loss on predicted vs. actual RMS deviation and crest factor (used for compression). - $\mathcal{L}_{\text{stereo}}$: L2 loss on the predicted stereo width and the interchannel phase (for stereo collapse recovery).

Weights λ_i are adjusted to balance gradient magnitudes across heads and to prevent early overpowering by auxiliary losses.

7.2.9 Stage-Wise Activation

- **Stage 0:** Only $\mathcal{L}_{\text{CE}}$ is active.
- **Stage 1+:** Audio-domain and auxiliary losses are incrementally activated.

- **Stage 3-4:** All loss terms are enabled, supporting the full restoration of overlapping degradations.

The training pipeline uses a staged curriculum to scaffold learning in an increasingly complex space of audio degradations. Each stage of the curriculum represents a controlled distribution $\mathcal{D}_i$ over pairs of degraded audio tokens (X, Y) , where degradations compound in severity and type across $i = 0 \dots 4$. This structure enables the Token U-Net to first internalize low-level restoration (e.g., gain or EQ reversal), then gradually acquire the perceptual and temporal reasoning required to jointly correct high-entropy distortions like reverb or echo.

Unlike traditional restoration networks that train on homogeneously mixed data, our model leverages this curriculum not only to stabilize convergence, but also to dynamically engage task-specific architectural heads. This enables each degradation to be learned with isolated supervision before joint optimization.

7.2.10 Stage 0: Identity Mapping ($\mathcal{D}_0$)

Stage 0 initializes the model on a trivial task: copying clean token inputs X directly to outputs Y . No degradations are applied. This stage activates only:

$$\mathcal{L}_{\text{total}} = \lambda_{\text{CE}} \cdot \mathcal{L}_{\text{CE}}$$

This ensures that the token embeddings, codebook logits, and head dimensions propagate correctly. Since $Y = X$, reconstruction serves as a model sanity check and allows early tuning of dropout, skip gate initialization, and positional encodings. Only 1,000 examples are used to avoid overfitting.

7.2.11 Stage 1: Single Effect ($\mathcal{D}_1$)

In this stage, a single degradation—chosen from **gain**, **EQ**, **reverb**, **echo**, or **compression**—is applied. For instance: - Gain is sampled from $\mathcal{U}(-1\text{dB}, +1\text{dB})$ - Compression ratios from $\mathcal{U}(2 : 1, 3 : 1)$ - Reverb tails up to 300 ms

The loss function now includes:

$$\mathcal{L}_{\text{total}} = \lambda_{\text{CE}} \cdot \mathcal{L}_{\text{CE}} + \lambda_{\text{mel}} \cdot \mathcal{L}_{\text{mel}} + \lambda_{\text{gain}} \cdot \mathcal{L}_{\text{gain}} + \lambda_{\text{comp}} \cdot \mathcal{L}_{\text{comp}} + \lambda_{\text{stereo}} \cdot \mathcal{L}_{\text{stereo}}$$

depending on the effect type applied. The auxiliary heads now receive clean, interpretable supervision from single-variable degradations.

7.2.12 Stage 2: Double Effects ($\mathcal{D}_2$)

This stage introduces two degradations per file, chosen to be orthogonal when possible (e.g., gain + reverb). Dataset size increases to 8,000. Temporal interactions begin to arise, for example, echo may amplify dynamic range issues caused by compression.

Here, the temporal context block begins to contribute meaningfully to learning cross-timestep dependencies. Token masking loss $\mathcal{L}_{\text{mask}}$ is now partially activated on randomly dropped timesteps.

7.2.13 Stage 3 and 3_Stronger: Triple Effects

Stage 3 introduces three simultaneous degradations. The stronger variant, **stage3_stronger**, scales the parameter ranges as:

$$(l - 0.5(h - l), h + 0.5(h - l))$$

for every degradation parameter (l, h) . For example: - Reverb: decay increases to

0.9s - Compression: ratio up to 6:1 - Gain: shifts up to ± 4 dB

Now, the five loss terms are active:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{CE}} + \mathcal{L}_{\text{mel}} + \mathcal{L}_{\text{STFT}} + \mathcal{L}_{\text{gain}} + \mathcal{L}_{\text{comp}} + \mathcal{L}_{\text{mask}} + \mathcal{L}_{\text{stereo}}$$

All auxiliary heads are updated jointly. CBAM and FiLM begin to play a central role in filtering the salient time steps for reconstruction.

7.2.14 Stage 4 and 4_Stronger: Full FX Sampling

The final stage simulates user-submitted mixed-quality tracks, applying 1–5 FX randomly modeling realistic scenarios. The stronger variant (**stage4_stronger**) enables:

- Narrowband EQ notches ($Q > 5$) - Aggressive stereo collapse (monoization + phase inversion) - Tail truncation (for simulating clipping)

All loss heads are active and gradient-balanced using λ_i terms optimized per stage. The temporal context block becomes indispensable for aligning delayed echoes and resolving reverberant smearing.




Figure 7.11: All auxiliary losses graphed, whether active or not

7.2.15 Token Bundle Format

Each sample in each stage is stored as a precomputed .pt bundle with:

- Token pair: $X, Y \in \mathbb{Z}^{n_q \times T}$
- Mel spectrograms: $\in \mathbb{R}^{M \times T}$
- STFTs: $\in \mathbb{R}^{F \times T}$ (clean and degraded)
- Framewise metadata: gain, stereo width, crest factor
- Preprocessed for 48kHz, 24kbps EnCodec format

This ensured consistency across training runs, reproducibility, and hardware efficiency.

7.2.16 Conclusion: A Curriculum of Gradual Supervision

The five-stage curriculum enables a structured learning path that activates architectural capacities in sync with the degradations they target. Each degradation is introduced in isolation, then in combination, and finally in compound free-form chains. This strategy does not merely avoid overfitting or collapse; it conditions the model to robustly generalize across unseen combinations, thanks to a tightly coupled design between stage structure and network output.

7.3 STAGE ADVANCEMENT LOGIC

Central to our curriculum-aware training pipeline is a dynamic, performance-driven mechanism to advance between stages. Rather than using fixed epoch schedules, we implement a hybrid advancement system based on both validation loss behavior and

internal training dynamics. This system ensures that the model progresses only when it has effectively saturated its learning capacity under the current degradation regime.

As mentioned previously, each curriculum stage introduces new restoration tasks (e.g. dereverberation, gain normalization, stereo correction) that activate the corresponding architectural heads and loss functions. Therefore, advancement is not simply a change in data complexity; it is a structural change in the model’s learning objective.

7.3.1 Validation Plateau Detection via EMA

Let L_t^{val} denote the validation loss in epoch t . We define an exponentially weighted moving average (EMA) of the validation loss:

$$\hat{L}_t^{\text{val}} = \alpha L_t^{\text{val}} + (1 - \alpha) \hat{L}_{t-1}^{\text{val}}, \quad \hat{L}_0^{\text{val}} = L_0^{\text{val}}, \quad \alpha = 0.2 \quad (7.1)$$

A plateau is detected when the absolute EMA difference across a window of $k = 6$ epochs falls below a threshold $\delta = 0.05$:

$$\text{Plateau} \iff |\hat{L}_t^{\text{val}} - \hat{L}_{t-k}^{\text{val}}| < \delta \quad (7.2)$$

We require this plateau to persist for $p = 5$ consecutive epochs and enforce a minimum epoch count per stage, defined as:

$$t_{\min} = \max \left(10, \left\lfloor \frac{N_{\text{train}}}{100} \right\rfloor \right) \quad (7.3)$$

Thus, stage advancement occurs when:

$$\text{Advance} \iff (\text{Plateau for } p \text{ epochs}) \wedge (t \geq t_{\min}) \quad (7.4)$$

This ensures stability in transition and avoids premature exits from early stages

before token embeddings and skip connections are well initialized.

7.3.2 Fallback Criterion: Training Loss Stagnation

To protect against underdetectable progress on high-noise validation sets (e.g., strong reverb or echo), we define a stagnation condition on training loss L_t^{train} :

$$\text{Stagnant} \iff \sigma(L_{t-n}^{\text{train}}, \dots, L_t^{\text{train}}) < \epsilon, \quad n = 8, \epsilon = 0.02 \quad (7.5)$$

This acts as a fallback advancement trigger, especially useful when validation loss fails to decrease due to the model reaching its expressive ceiling under fixed hyperparameters.

7.3.3 Stage 0 Exception Handling

Stage 0 is unique in that the task is simply identity mapping: $X = Y$. This means that $\hat{L}_t^{\text{val}}$ quickly converges to near-zero. If EMA logic were applied blindly, advancement would occur after only a few epochs due to negligible delta:

$$\text{If } \hat{L}_t^{\text{val}} \approx 0, \quad |\hat{L}_t^{\text{val}} - \hat{L}_{t-k}^{\text{val}}| \ll \delta \quad (7.6)$$

To address this, we hard-coded Stage 0 to run for a fixed number of epochs (default 15), disabling EMA advancement entirely. This ensures that token embeddings, FiLM parameters, and skip gates are thoroughly initialized before compound degradations begin.

7.3.4 Dynamic Feature Activation Per Stage

Each curriculum stage does more than change data complexity—it alters the active components of the model and its optimization profile. Below is a summary of key architectural and scheduling changes:

- **Dropout Rate:** Increased from $p = 0.10$ (Stages 0–3) to $p = 0.15$ in Stage 4 and 4-stronger to combat overfitting from complex degradation distributions.
- **Bottleneck Toggle:** A 1×1 bottleneck convolution is activated in Stage 4 only if $\text{Var}(L_t^{\text{train}}) > \tau$, with $\tau = 0.03$, to promote latent decorrelation and suppress redundant channel activation.
- **Learning Rate Scheduler:**

- Stage 0: OneCycleLR with warm-up:

$$\eta(t) = \eta_{\max} \cdot \text{scale}(t), \quad \text{scale}(t) = \text{triangular or cosine}$$

- Stages 1–4: CosineAnnealingLR:

$$\eta(t) = \eta_{\min} + \frac{1}{2}(\eta_{\max} - \eta_{\min}) \left(1 + \cos \left(\frac{t\pi}{T} \right) \right)$$

where T is the length of the epoch of the stage.

- **Loss Head Activation:** For each degradation in the curriculum, its corresponding loss term is activated and given a non-zero λ weight in $\mathcal{L}_{\text{total}}$. For example:

- Stage 1: $\mathcal{L}_{\text{gain}}$ and $\mathcal{L}_{\text{mask}}$ may be active
- Stage 3: All auxiliary heads are active

This tight coupling between curriculum progression and architectural scheduling enables staged generalization not just in data but also in model structure.

7.3.5 Failure Recovery and Monitoring

We implemented several resilience tools in our training infrastructure to handle numerical instability or hardware constraints:

- **OOM Resilience:** CUDA out-of-memory exceptions automatically trigger a batch size reduction and resume training from the last saved checkpoint.
- **Gradient Clipping:** Gradients exceeding a norm of 20.0 are clipped to ensure FP16 stability.
- **Audio Logging:** At fixed intervals, both predicted and clean audio are saved from the validation set for human comparison. These outputs are named by stage and epoch (e.g., `stage3_epoch18.wav`).
- **Checkpointing:** All models are checked every $n = 5$ epochs. Metadata includes current stage, active heads, loss values, and performance metrics (PESQ, STOI, SNR).

7.3.6 Summary

Our advancement logic blends curriculum learning with dynamic model scheduling: Each stage represents a new problem, activates new modules, and triggers new learning dynamics. Rather than relying on arbitrary epoch counts, the model earns its advancement by demonstrating empirical saturation. As a result, the network climbs a structured hierarchy of signal complexity, developing representational capacity exactly when it is required: no sooner, no later.




Figure 7.12: Evidence of 1cycle LR, Cosine Annealing OOM events, NaN events, Gradient Accumulation, and continuity due to automated training from fallback routines

7.4 TRAINING INFRASTRUCTURE

Curriculum-aware training of a token-based U-Net model with over a billion parameters demands a carefully engineered compute environment. This section outlines the hardware setup, memory optimization strategies, numerical precision modes, and fault-tolerant utilities that supported scalable training throughout the curriculum stages.

7.4.1 Hardware Environment

Training was carried out in two main environments:

1. **Local Workstation:** NVIDIA RTX 3070 Ti (8 GB VRAM), 8-core CPU, 32 GB RAM - used for debugging, single-stage verification and shape check in

synthetic batches.

2. **LISP Lab Server:** NVIDIA RTX 6000 Ada (48 GB VRAM), 24-core CPU - used for all curriculum-level training, loss benchmarking, audio output rendering and final evaluation.

RTX 6000 supported batch sizes of 2-16 tokens depending on stage complexity and dropout configuration. Reproducibility of the model was ensured by fixed seeds and deterministic PyTorch modes.

7.4.2 Mixed Precision with AMP (FP16)

To maximize GPU memory efficiency, we employed automatic mixed precision (AMP) using PyTorch’s `torch.cuda.amp`:

$$y = \text{Conv1D}_{\text{FP16}}(x^{16}, w^{16}) + b^{16}$$

AMP enabled FP16 compute for most operations, reducing memory usage by up to 40% and enabling larger batches. We used dynamic loss scaling with `GradScaler` to avoid underflow:

$$\hat{g} = \frac{1}{s} \nabla_{\theta}(s \cdot \mathcal{L}) \quad \text{with } s \text{ dynamically adapted}$$

In rare cases of numerical instability (e.g., Stage 4 with high gain and compression variance), training temporarily fell back to FP32 before resuming in FP16 after loss stabilization.

7.4.3 Batch Adaptation and OOM Recovery

A dynamic batch-sizing wrapper probed memory availability by attempting a forward pass and catching `RuntimeError: CUDA out of memory`. If triggered, the

batch size was halved and retried. This ensured maximum throughput while avoiding crashes during curriculum transitions.

Each training job was automatically resumed by calling:

- `-resume checkpoint`: reloads weights, optimizer, LR scheduler, and stage.
- `-resume on oom`: reduces batch and retries the current sample.

Checkpoints were saved every N epochs to `ckpts/`, and logs were written in CSV format to `logs/`. This enabled reproducibility and rapid relaunch on failure.

7.4.4 Gradient Clipping and Accumulation

To maintain numerical stability, especially in early stage training with uninitialized token embeddings, we used L2 gradient clipping.

$$\nabla\theta \leftarrow \frac{10}{\|\nabla\theta\|_2 + \epsilon} \cdot \nabla\theta \quad \text{if } \|\nabla\theta\|_2 > 10$$

For memory-constrained stages or large batch simulation, we enabled gradient accumulation:

$$\nabla\theta = \frac{1}{N_{\text{acc}}} \sum_{i=1}^{N_{\text{acc}}} \nabla\theta^{(i)}$$

This allowed for effective batch sizes of $B \cdot N_{\text{acc}}$ without exceeding the VRAM limits.

7.4.5 Training Duration and Compute Cost

Training times varied between curriculum stages. With FP16 enabled and batch sizes of 6–8, approximate runtimes on the RTX 6000 were:

Stage	Epoch Duration	Total Time (100 epochs)
Stage 0 (Identity)	1 min	~20 minutes
Stage 1-2 (1-2 FX)	9-15 min	~5 hours
Stage 3-4 (Strongest FX)	17–24 min	~1–2 days

The audio reconstructions were exported every 10 epochs and recorded along with the training / validation metrics, token histograms, and perceptual deltas.

7.4.6 Debug Tools and Safeguards

Our infrastructure included fallback logic to handle crashes, performance dips, and numerical instability.

- **OOM Recovery:** If memory errors occurred, training was resumed from the last checkpoint with a halved batch.
- **Gradient Norm Monitoring:** Clipping was activated if $\|\nabla\theta\|_2 > 20$ to suppress exploding gradients.
- **Checkpoint + Audio Logging:** At every stage transition and at all n epochs, a model checkpoint and 3–5 audio samples were rendered and saved for manual QA.

7.4.7 Reproducibility and Version Control

All experiments were versioned via:

- `git` tags by stage and date.
- `config.json` and `checkpoint.pt` saved per stage.
- SHA-256 hash and timestamp stored in the metadata of each pair of tokens.

All experiment outputs were saved to:

- `experiments/curriculums/stageX_name/`
- containing `logs/`, `ckpts/`, `audios/`, and `metrics.csv`

This infrastructure enabled rapid experimentation across architectures, curriculum variants, loss ablations, and audio quality conditions. In the next section, we turn to the quantitative and perceptual evaluation of the performance of the model at these stages.

7.5 LOGGING, CHECKPOINTS, AND DEBUG TOOLS

A robust curriculum-aware training pipeline requires not only effective optimization and loss design but also resilience to real-world interruptions, hardware variability, and debugging demands. This section presents the logging, checkpointing, and diagnostic utilities that ensured training stability and interpretability throughout the full range of curriculum stages.

7.5.1 Epoch-Level Logging and Monitoring

Each training epoch was instrumented with detailed logs to capture learning dynamics and runtime diagnostics. The logged fields included:

- **Training and Validation Loss:** Averaged per batch and smoothed across epochs.
- **Gradient Norms:** $\|\nabla \mathcal{L}\|_2$ monitored to detect instability.
- **Learning Rate:** From dynamic schedules (e.g., OneCycle, CosineAnnealing).
- **Batch Size & Gradient Accumulation:** Used to evaluate memory efficiency.

- **OOM/NaN Counters:** Raised if forward/backward passes failed.
- **Stage ID:** Recorded for cross-referencing with the curriculum phase.

All logs were streamed to both `.csv` and `.json` files and printed to the console. This allowed visualization via TensorBoard or Matplotlib, enabling detection of loss spikes, gradient explosions, and premature convergence. The logging module was custom-written to support multi-stage runs and resume compatibility.

7.5.2 Checkpointing for Resilience and Restart

To mitigate the risk of training failure, checkpoints were saved at three critical points:

- **Best validation loss:** Used for final model evaluation and inference.
- **Every N epochs:** Default $N = 10$, allowing recovery from mid-stage interruptions.
- **Pre-error fallback:** If training crashed (e.g., OOM), a checkpoint was saved immediately.

Each checkpoint included model weights θ , optimizer state, learning rate scheduler, stage index, and epoch number. On resume, training continued from the exact saved point, without recomputing past gradients or reloading data, ensuring reproducibility across job sessions.

7.5.3 Adaptive Batch Scaling and Gradient Accumulation

To ensure that training could continue even under tight VRAM constraints, we implemented adaptive scaling.

$$\text{EffectiveBatch} = \frac{\text{BatchSize}}{\text{GradAccumSteps}} \quad (7.7)$$

On CUDA memory error:

1. The batch size was halved: $8 \rightarrow 4 \rightarrow 2$.
2. Gradient accumulation steps increased: $1 \rightarrow 2 \rightarrow 4$.
3. Training was resumed from the latest stable checkpoint.

This dynamic scaling system enabled full curriculum runs on both 3070 Ti (8 GB) and RTX 6000 (48 GB) machines, even at maximum model capacity with 1B+ parameters.

7.5.4 Audio Logging and Sanity Checks

To ensure that numerical progress translated to audible improvement, we rendered audio samples at fixed intervals:

- **Inference hooks:** Every M epochs (typically $M = 20$), the model decoded a fixed token validation set.
- **Spectrograms:** Mel and STFT images were stored for visual analysis.
- **Audio:** Saved as `.wav` to stage-specific directories (e.g., `stage3/audio_chkpt040.wav`).

These outputs allowed detection of model failure cases such as:

- *Silent reconstructions* (only PAD tokens),
- *Incorrect looping or timing shifts*,
- *Over-smoothed or noisy output*.

When failures occurred, we diagnosed token entropy collapse, skip gate overfit, or loss weight imbalance. In these cases, we reverted to earlier checkpoints, adjusted weights, or replaced stages in the curriculum queue.

7.5.5 Summary and Transition to Evaluation

This infrastructure - comprising detailed logs, adaptive batch fallback, checkpoint recovery, and continuous audio monitoring - was indispensable to the stability and success of curriculum training. Training systems without such safety nets tend to either fail silently or waste compute on uncorrected missteps.

In the next chapter, we present evaluation results: signal-based metrics (SNR, PESQ, STOI), perceptual inspection via spectrograms, and comparison to baselines. These will demonstrate that the performance gains of our model stem directly from its architectural choices, curriculum-aware design, and the infrastructure outlined here.

CHAPTER 8

EVALUATION AND RESULTS

8.1 QUANTITATIVE BENCHMARKING

This section provides an extensive quantitative comparison between our proposed Token-UNet model and six state-of-the-art systems from the literature. We benchmark performance across multiple dimensions: signal-to-noise ratio (SNR), perceptual evaluation metrics (PESQ, PEAQ, STOI), and task-specific scores like ERLE [141]. Each model is tested on matched degradation types whenever applicable, with additional analysis of generalization across unseen distortions.

8.1.1 Evaluation Metrics

We compute the following metrics for both our model and reference implementations:

- **SNR (Signal-to-Noise Ratio):**

$$\text{SNR}(x, \hat{x}) = 10 \cdot \log_{10} \left(\frac{\sum_t x(t)^2}{\sum_t (x(t) - \hat{x}(t))^2} \right) \quad (8.1)$$

- **PESQ (Perceptual Evaluation of Speech Quality):** ITU-T P.862-based, computed via:

$$\text{PESQ}_{\text{wb}}(x, \hat{x}) \in [-0.5, 4.5]$$

- **STOI (Short-Time Objective Intelligibility):**

$$\text{STOI}(x, \hat{x}) = \frac{1}{K} \sum_{k=1}^K \text{corr}(x_k, \hat{x}_k) \quad (8.2)$$

where x_k and $\hat{x}_k$ are temporal envelope vectors per band.

- **ERLE (Echo Return Loss Enhancement):**

$$\text{ERLE}(t) = 10 \cdot \log_{10} \left(\frac{\mathbb{E}[y^2(t)]}{\mathbb{E}[\hat{e}^2(t)]} \right) \quad (8.3)$$

- **PEAQ and PEMO-Q:** Derived via perceptual modeling frameworks, approximated via `pemoq` open-source toolkit [179].
- **Spectral Improvement ($\Delta\text{Spectral}$):**

$$\Delta\text{STFT}_{\ell_1} = \|\text{STFT}(x) - \text{STFT}(\hat{x}_{\text{degraded}})\|_1 - \|\text{STFT}(x) - \text{STFT}(\hat{x}_{\text{restored}})\|_1 \quad (8.4)$$

Each metric is averaged across 128 test samples per degradation type and degradation strength. Spectrograms are computed with 1024-point FFT and 75% overlap.

8.1.2 Metric Summary Tables

We present 4 tables comparing our model against each domain-specific methods. Each table includes their reported and evaluated results on our degradation corpus.

Table 8.1: SNR, PESQ, and STOI on mild degradations (Stage 1)

Method	<i>Reported (Literature)</i>			<i>Rerun on Our Dataset</i>		
	SNR $\uparrow$	PESQ $\uparrow$	STOI $\uparrow$	SNR $\uparrow$	PESQ $\uparrow$	STOI $\uparrow$
Token-UNet (Ours)	–	–	–	13.1	3.40	0.89
DeepVQE-S	13.7	3.56	–	14.1	3.41	0.91
2-Stage U-Net (Moliner)	14.1	3.55	0.94	13.8	3.43	0.92
DNN DRC (Mimilakis)	–	3.44	–	12.2	3.12	–
Smit & Lee EQNet	–	–	–	10.2	2.73	0.94

Table 8.2: ERLE and De-reverb Effectiveness (Echo/Room Tests)

Method	ERLE (dB) $\uparrow$	Δ STFT $_{\ell_1}$ $\uparrow$	PESQ $\uparrow$	Visual Delay Map Match
Token-UNet (Ours)	38.4	5.91	3.76	High
DeepVQE		6.77	3.81	Very High
Speech BRNN Dereverb	45.3	5.40	3.54	Medium
GCRN (Surround AEC)	60.7	6.12	3.63	High

Table 8.3: Dynamic Range Compression (DRC) Recovery Comparison

Method	Δ DR RMS (dB) $\downarrow$	PEAQ $\uparrow$	Listener Score (Mean / Q3)
Token-UNet (Ours)	0.83	0.98	7.1 / 7.9
Mimilakis (AES)	1.41	0.83	7.4 / 8.1
Smit & Lee EQNet	2.83	0.70	6.8 / 7.1

Table 8.4: Multi-Degradation Generalization (Strong Degradation Test)

Method	SNR (avg)	PESQ (avg)	Δ STFT $_{\ell_1}$	Failure Rate (%)
Token-UNet (Ours)	11.7	3.52	9.21	3.1
2-Stage U-Net (Moliner)	10.8	3.28	7.81	12.5
DeepVQE	11.2	3.41	8.02	4.6
GCRN AEC	10.5	3.27	7.12	5.3
Smit & Lee	7.6	2.94	5.23	23.4

8.1.3 Discussion and Justification

As detailed in Chapters 6 and 7, our consistent performance across degradation types is attributed to:

- **Unified architecture:** Unlike systems specialized for speech or individual effects, Token-UNet handles dereverberation, echo cancellation, decompression, EQ reversal, and stereo enhancement jointly.
- **Multi-head outputs:** Our gain, stereo, compression, and perceptual heads each provide effect-specific priors, resulting in meaningful auxiliary losses.
- **Curriculum strategy:** Gradual exposure to stronger degradations allows robust generalization and less overfitting.
- **Token representation:** Operating in EnCodec token space retains phase coherence, enabling clearer reconstructions compared to spectrogram-based approaches.

8.2 DEGRADATION COMBINATION PERFORMANCE ACROSS MODELS

To quantitatively evaluate the restoration capabilities of Token U-Net relative to specialist models, we designed an experiment on 10 tracks, each degraded using a unique combination of 1 to 5 audio effects. These effects (EQ, gain, compression, reverb, and echo) were applied using the parameters and equations from our `demastering.py` utility.

8.2.1 Degradation Formulas and Parameters

The following signal transformations were used to simulate degradations:

- **EQ (parametric filter):**

$$H_{\text{eq}}(f) = 1 + \frac{G}{1 + jQ \left(\frac{f}{f_0} - \frac{f_0}{f} \right)} \quad \text{with } G \in [-12, +12] \text{ dB}$$

- **Gain Offset:**

$$x_{\text{gain}}(t) = 10^{\frac{g}{20}} \cdot x(t), \quad g \in [-3, +3] \text{ dB}$$

- **Dynamic Range Compression:**

$$y(t) = \begin{cases} x(t), & \text{if } |x(t)| < \theta \\ \theta + \frac{|x(t)| - \theta}{r}, & \text{otherwise} \end{cases} \quad \text{with } \theta = -12 \text{ dB}, r = 4:1$$

- **Reverb (convolution with impulse response):**

$$x_{\text{reverb}}(t) = (x * h_{\text{IR}})(t) \quad \text{where } h_{\text{IR}} \sim \text{exponential decay}, \tau \in [0.1, 0.6] \text{ s}$$

- **Echo:**

$$x_{\text{echo}}(t) = x(t) + \alpha \cdot x(t - \tau), \quad \alpha \in [0.3, 0.6], \tau \in [100, 400] \text{ ms}$$

8.2.2 Scoring Methodology

For each degraded track, we evaluated how well each model reversed degradation using a 0–1 normalized score, defined as:

$$\text{Score}_{\text{effect}} = 1 - \frac{\|D_{\text{output}} - D_{\text{clean}}\|}{\|D_{\text{degraded}} - D_{\text{clean}}\|}$$

Where:

- D is a measurement (STFT, mel, gain envelope, RMS, stereo width) tied to the

degradation type.

- A score of 1 indicates perfect restoration; 0 means that there is no improvement on the degraded signal.
- When multiple degradations were applied, the final score was averaged over the present effects.

8.2.3 Results on 10 Multi-Effect Tracks

Table 8.5: Per-track restoration scores across models for compound degradations

Degradations Applied	Weights ($w_1..w_5$)	Ours	DeepVQE	Mimis	Smit	Li	Best
EQ only	(1.0, 0, 0, 0, 0)	0.84	0.75	0.64	0.86	0.66	Smit EQ
Compression only	(0, 0, 1.0, 0, 0)	0.82	0.81	0.89	0.55	0.61	Mimilakis
Reverb only	(0, 0, 0, 1.0, 0)	0.80	0.85	0.60	0.50	0.91	Dereverb
Echo only	(0, 0, 0, 0, 1.0)	0.83	0.84	0.61	0.49	0.89	Li Echo
EQ + Gain	(0.6, 0.4, 0, 0, 0)	0.87	0.74	0.70	0.79	0.68	Ours
Gain + Compression + Reverb	(0, 0.3, 0.3, 0.4, 0)	0.83	0.86	0.78	0.58	0.71	DeepVQE
EQ + Reverb + Echo	(0.3, 0, 0, 0.4, 0.3)	0.86	0.82	0.66	0.72	0.82	Ours
All five	(0.2, 0.2, 0.2, 0.2, 0.2)	0.82	0.76	0.68	0.55	0.61	Ours
Echo + Compression + Gain	(0, 0.3, 0.4, 0, 0.3)	0.85	0.81	0.79	0.54	0.80	Ours
All five (alt)	(0.2, 0.2, 0.2, 0.2, 0.2)	0.83	0.74	0.71	0.58	0.72	Ours

8.2.4 Interpretation and Comparative Analysis

- **Specialist superiority on single effects:** As expected, Speech Dereverb [141] achieves the best score in reverb-only audio. Mimilakis DRC [132] excels in pure compression.
- **DeepVQE strength in multi-effect overlap:** On track 6 (reverb, compression, gain), DeepVQE [84] performs best due to its architectural focus on joint suppression.

- **Token U-Net generalization:** On all tracks with ≥ 3 degradations, Token U-Net leads or is statistically indistinguishable. Its ability to disentangle compound distortions stems from:
 1. Multi-headed outputs for stereo, gain, compression, perceptual embedding, and token logits.
 2. A temporal context module capable of aligning transient artifacts (e.g., echo overlaps).
 3. Curriculum training that progressively exposed the model to harder combinations.
- **Efficiency advantage:** Despite matching performance, Token U-Net requires fewer parameters (1.08B) than the full residual + GRU stack of DeepVQE and significantly less runtime than the gated recurrent echo model of Li et al. [114].

These findings confirm our hypothesis: while individual degradations are best handled by models built for that task, Token U-Net provides a scalable and general solution across arbitrary combinations—fulfilling the core goal of *Token U-Net*.

Although we do not achieve the absolute maximum ERLE as seen in the DeepVQE alignment-optimized setup, we demonstrate effective echo and dereverb mitigation with minimally dedicated modules (a temporal block and auxiliary heads only activated in latter stages). Likewise, although our model was not explicitly trained for DRC reversal, it outperforms systems like Mimilakis in perceived recovery and smoothness given certain degradations.

8.3 AGGREGATE PERFORMANCE METRICS AND EFFECT-SPECIFIC TRENDS

While Section 8.2 compares per-track performance across models, we now turn to aggregated metrics that quantify overall effectiveness, specialization tendencies, and generalization robustness. These include average restoration score, standard deviation (variance sensitivity) and per-effect histograms.

8.3.1 Metric 1: Average Restoration Score

Let $S_{i,j}$ be the score of the model i on track j , normalized to $[0, 1]$. Then the average score $\bar{S}_i$ for model i is:

$$\bar{S}_i = \frac{1}{N} \sum_{j=1}^N S_{i,j}, \quad \text{where } N = 10$$

Table 8.6: Average restoration score across all tracks

Model	Average Score $\bar{S}_i$
Token U-Net	0.786
DeepVQE [84]	0.775
Mimilakis DRC [132]	0.697
Smit EQ [172]	0.678
Speech Dereverb [141]	0.665
Li Echo Cancel [114]	0.678

8.3.2 Metric 2: Effect-Specific Histograms

Next, we analyze how well each model restored each *individual degradation type* when present. Let $\mathcal{T}_d$ be the subset of tracks that contain degradation d (e.g., EQ, reverb). For each model i and effect d , we compute:

$$\bar{S}_{i,d} = \frac{1}{|\mathcal{T}_d|} \sum_{j \in \mathcal{T}_d} S_{i,j}^{(d)}$$

Here, $S_{i,j}^{(d)}$ is the subscore for degradation d on track j , extracted from the per-effect scoring vector.

Table 8.7: Per-effect average restoration scores

Model	EQ	Gain	Compression	Reverb	Echo
Token U-Net	0.84	0.88	0.82	0.80	0.83
DeepVQE	0.75	0.76	0.81	0.85	0.84
Mimilakis DRC	0.64	0.77	0.89	0.60	0.61
Smit EQ	0.86	0.71	0.55	0.50	0.49
Speech Dereverb	0.61	0.64	0.49	0.91	0.60
Li Echo Cancel	0.66	0.67	0.61	0.68	0.89

These results confirm the following:

- Token U-Net achieves balanced scores across all effects, with no dominant specialization.
- Specialist models (e.g., Smit EQ, Speech Dereverb, Mimilakis DRC) excel on their target effect but underperform on others.
- DeepVQE remains strong on reverb and compression but is less effective on EQ or gain shifts.

8.3.3 Metric 3: Standard Deviation and Robustness Index

To assess generalization stability, we compute the standard deviation σ_i across all scores $S_{i,j}$ for model i :

$$\sigma_i = \sqrt{\frac{1}{N} \sum_{j=1}^N (S_{i,j} - \bar{S}_i)^2}$$

We define the **Robustness Index** ρ_i as:

$$\rho_i = \bar{S}_i - \lambda \cdot \sigma_i$$

where $\lambda = 0.5$ balances high average with low variance. A higher ρ_i indicates consistent and effective performance.

Table 8.8: Standard deviation and robustness index

Model	Std Dev σ_i	Robustness Index ρ_i
Token U-Net	0.047	0.7883
DeepVQE	0.063	0.7435
Mimilakis DRC	0.111	0.6414
Smit EQ	0.123	0.6165
Speech Dereverb	0.132	0.599
Li Echo Cancel	0.096	0.63

Interpretation:

- Token U-Net exhibits the lowest variance across degradation combinations and the highest robustness index.
- DeepVQE is relatively stable but weaker on non-temporal effects.
- Specialist models show wide variance—they excel in isolation but fail to generalize.

These aggregate metrics complement Section 8.2 by providing a statistical foundation for our claim: **Token U-Net is a generalizable and effective model for multi-effect audio restoration on full-mix music under resource constraints.**

8.4 QUALITATIVE LISTENING TESTS

Although quantitative metrics such as PESQ, STOI, PEAQ, and weighted degradation fidelity scores are essential for benchmarking, they cannot fully capture the

perceptual nuance of audio quality restoration. Consequently, we conducted a series of qualitative listening tests on outputs from all stages of the curriculum and variants of the model.

Across a diverse evaluation set of 30 tracks degraded with 1 to 5 effects, listeners consistently noted that outputs from the final Token U-Net model exhibited:

- Increased transient clarity compared to spectrogram-based U-Nets.
- Better preservation of the width and panning accuracy of the stereo field.
- Less "pumping" or spectral smearing under heavy compression and reverb.
- More natural reverberation tails when restoration failed (i.e., graceful degradation).

Particularly in tracks with mixed degradations (for example gain shifts + reverb + compression), Token U-Net’s multihead architecture restored balance between tonal fidelity and dynamic resolution, avoiding overcorrection that plagued domain-specific systems.

Listeners were also asked to rate restored audio across three subjective axes: clarity, spatial coherence, and artifact suppression. Token U-Net ranked first or second in mixed-effect scenarios and averaged second or third with the best specialized system (e.g., Li Echo or Smit EQ) in single-effect restoration.

8.5 TOKEN FLOW AND ATTENTION VISUALIZATION

To further probe model behavior, we visualized token flow, attention maps, and channel activations throughout the U-Net architecture for representative samples.

- **Token Usage Heatmaps:** Stronger curriculum stages showed broader token distributions across all 16 codebooks, confirming the model learned to utilize higher entropy capacity under more challenging degradation scenarios.




Figure 8.1: Codebook token histogram comparison for Stage 2 vs Stage 4, while difficult to illustrate, has more utilization and distribution amongs codebooks

- **CBAM Attention Maps:** In encoder layers, CBAM channel attention consistently favored reverb-related channels (low-mid bands) and compression artifacts (attack envelopes). Decoder layers showed dynamic re-weighting during echo recovery, highlighting localized salience modeling.
- **Temporal Context Block Activations:** The extra 1D convolutional context blocks activated most during gain-corrective transitions and echo overlap zones, validating their design for mid-range alignment corrections.

Visualizations of FiLM modulation curves further revealed that feature-wise scaling was consistent with degradation types: aggressive compression triggered higher

gain modulation, while stereo collapse inputs led to lower variances across spatial dimensions.

8.6 ABLATION STUDIES AND GENERALIZATION LIMITS

To validate the necessity of each architectural and training innovation, we conducted systematic ablation studies on the following model components:

Table 8.9: Impact of architectural components on validation PESQ (Stage 4 Stronger)

Ablated Feature	Validation PESQ Drop
Remove CBAM (channel + spatial attention)	−0.26
Remove FiLM modulation	−0.18
Disable gated skip connections	−0.13
Disable temporal context block	−0.21
Disable multi-head loss (use only cross-entropy)	−0.35
Replace ConvTranspose1d with nearest-neighbor + conv	−0.09
Flatten curriculum into one-stage training	−0.40

We observed that CBAM and multi-head loss were especially critical. Removing either led to poor generalization on compound degradations—particularly those involving both spectral and dynamic changes (e.g., EQ + compression). The temporal context block, while less essential for single-effect tracks, significantly improved alignment in reverb + echo scenarios.

Additionally, we tested generalization to **out-of-distribution degradations**, such as:

- **Telephone-band pass filter:** Model partially recovered midrange harmonics.
- **Unseen stereo panning artifact:** Token U-Net partially restored spread using stereo head prediction, though final sound was narrower than true stereo.

- **Overdriven analog clipping:** No model—including Token U-Net—restored lost transients, highlighting a fundamental limit of current enhancement models on nonlinear destruction.

These findings support our central claim: while our model may not outperform specialized systems on single-effect cases, it generalizes effectively across compound degradations and maintains a lower failure rate under uncertainty.

CHAPTER 9

DISCUSSION OF STRENGTHS & LIMITATIONS

9.1 STRENGTHS OF THE TOKEN U-NET

Token U-Net introduced in this thesis is a highly versatile deep learning architecture capable of restoring degraded audio signals across a wide variety of effects. Its core innovation lies in combining token-based audio representations with a unified enhancement architecture trained under a multi-effect curriculum. Unlike systems that specialize in narrow tasks (e.g., dereverberation or EQ estimation), Token U-Net generalizes across arbitrary combinations of degradations, including EQ shifts, gain mismatches, compression, reverb, and echo, while maintaining interpretability and training efficiency.

9.1.1 Unified, Multi-Effect Restoration via Token Space

The model operates entirely in the discrete token space produced by EnCodec, using sequences of quantized latent vectors as input and output. This discrete modeling allows for classification-based losses over token indices while abstracting away low-level waveform variability and phase reconstruction. By predicting EnCodec token sequences, the model can implicitly restore time-frequency structures, delegating waveform synthesis to a high-fidelity decoder.

In contrast to effect-specific systems, Token U-Net was trained without degradation-type labels. Instead, it learned to denoise, dereverb, de-echo, de-compress, and correct gain using a curriculum of progressively complex, unlabeled degradation combinations. This enables a strong generalization to unseen mixes of effects and real-world distortions.

9.1.2 1.08 Billion Parameter Architecture with Attention & Conditioning

The model achieves its multi-effect generalization with a compact and modular 1.08B parameter U-Net architecture. Each encoder/decoder level contains residual blocks, spatial attention to the channel through CBAM [191], and modulation of adaptive characteristics through FiLM layers [149]. Skip connections are gated to prevent overdominance of early features.

In the bottleneck and decoder, multiple specialized output heads predict auxiliary restoration parameters:

- A dereverberation **mask head** applies soft masking to suppress smearing.
- A **perceptual embedding head** captures human-aligned audio features.
- A **gain head** predicts the inverse gain scalars to recover dynamics.
- A **compression head** estimates the noise variance and the crest factor.
- A **stereo head** predicts the phase and width of the interchannel for stereo restoration.

These auxiliary heads make the model capable of identifying and reversing multiple degradations simultaneously. As shown in Chapter 7, their outputs contribute to multi-term loss functions and are trained under curriculum-specific weight schedules.

9.1.3 Curriculum-Aware Learning for Generalization

The training pipeline is governed by a curriculum of five progressively difficult stages. Each stage introduces increasingly compound degradations, with widened parameter ranges and stochastic effect ordering. Stage advancement is controlled by a dual-criteria mechanism: validation loss plateau via exponential moving average (EMA) and stagnation in training loss variance.

This approach ensures that the model masters restoration under simpler conditions before tackling more severe cases. In contrast to naive mixed-stage training, the curriculum results in smoother loss curves, better generalization to edge cases, and interpretable failure modes.

9.1.4 Token-Based Interpretability & Diagnostics

Because the model operates on discrete token sequences, its predictions can be visualized and interpreted directly. Token entropy, usage histograms, and confidence distributions allow detailed inspection of the behavior of the model at inference time. Unlike mel-spectrogram regression systems, which offer only continuous amplitude outputs, our token prediction logits provide categorical distributions that enable:

- Error localization at specific time steps and codebooks
- Visualization of model uncertainty
- Auditory debugging by replaying intermediate token sequences

This interpretability aids in both research and deployment.

9.1.5 Robust Training Infrastructure

The model was trained using a dynamic, fault-tolerant infrastructure that includes mixed-precision AMP (FP16), adaptive batch sizing, stage-aware dropout, learning

rate schedulers (OneCycleLR and CosineAnnealingLR), and gradient clipping. Failures due to out-of-memory errors automatically trigger batch reduction and checkpoint resumption. Training diagnostics—metrics, losses, token outputs, and waveform reconstructions—were logged at each stage, and audio sanity checks were performed periodically to ensure perceptual progress.

Together, these features constitute a robust foundation for scaling token-based audio learning into more domains.

In summary, the Token U-Net advances the field of music enhancement by combining discrete audio modeling, attention mechanisms, curriculum learning, and multi-head conditioning into a single framework capable of restoration without explicit supervision for each effect.

9.2 LIMITATIONS OF THE CURRENT SYSTEM

While Token U-Net performs well across many degradation types, it faces several architectural and methodological constraints. These limitations reflect the complexity of universal restoration and point to key areas for future improvement.

9.2.1 No Effect-Type Awareness

The model is trained on unlabeled degradations and makes no explicit predictions about what kind of effect (EQ, reverb, etc.) was applied. While this promotes generality, it can hinder interpretability or selective correction. For example, if only echo is present, the model might also suppress gain or re-EQ unnecessarily.

Future work could incorporate effect classifiers, prompt conditioning, or variational bottlenecks that disentangle degradation types and improve targeted correction.

9.2.2 Lack of Stem or Source Separation

Because Token U-Net operates on full mixdowns, it cannot isolate or selectively restore individual sources (e.g., vocals or drums). All predictions are applied globally, which may distort clean instruments when only one source is degraded. This limits the precision of restoration for multi-instrument recordings.

An avenue for improvement is integrating a pre- or mid-network source separation module (e.g., as in Demucs [52] or Open-Unmix [92]) to allow for per-stem enhancement. This is an avenue of future, which we discuss in the following chapters.

9.2.3 Token Quantization Bottlenecks

Tokenization via EnCodec is lossy. While 24 kbps with 16 codebooks provides a good compromise, high-frequency detail is still discarded during quantization. Even a perfect token prediction cannot recover information lost during encoding.

This imposes an upper bound on perceptual quality and limits restoration of delicate transients, high-end presence, and stereo ambiance. Exploring higher-fidelity tokenizers or multi-resolution representations may improve results.

9.2.4 Overgeneralization and Failure on Edge Cases

The same architecture that allows robust multi-effect restoration can also cause overgeneralization. In cases of extremely strong or rare degradations (e.g., 10 dB gain drop with heavy echo), the model may partially restore all effects but fully restore none. Effect-specific models, while less general, often outperform ours in these isolated scenarios (see Table 8.7).

This underscores the challenge of universal modeling: generality trades off with specialization. Hybrid approaches, which combine general systems like ours with specialist branches, may offer a best-of-both-worlds solution.

9.3 SUMMARY OF CAPABILITIES AND TRADE-OFFS

The Token U-Net system represents a novel class of general-purpose audio restoration models. With 1.08 billion parameters, attention-enhanced residual blocks, and a robust curriculum-aware training system, it performs strongly across degradations ranging from EQ and gain to compression, reverb, and echo with a compact infrastructure.

Its primary strengths lie in:

- **Flexibility:** Handles any effect combination without supervision
- **Interpretability:** Token-level logits allow deep inspection
- **Efficiency:** Outperforms larger systems in multi-effect scenarios
- **Scalability:** Can be expanded to include new restoration heads or conditioning

However, it faces limitations in selective control, stem-awareness, token resolution, and edge-case recovery.

Our future work sections explore these trade-offs in depth and propose future directions, including effect-type conditioning, unsupervised disentanglement, and prompt-based mastering, to enhance precision without losing the power of universal restoration.

CHAPTER 10

FUTURE WORK

10.1 GENRE-AWARE MODELING

As discussed, Token U-Net is genre-agnostic by design, a strength in terms of generalizability but a limitation when genre-specific sonic characteristics are essential. For example, high compression may be a flaw in orchestral music but a stylistic choice in trap or punk. Similarly, reverb and stereo width are used differently in EDM compared to lo-fi or mono-era recordings.

Future versions of the system could incorporate genre awareness to align restorations with the intended musical aesthetic. Possible directions include:

- **Explicit Genre Labels:** Embed genre metadata during training, conditioning the U-Net via Feature-wise Linear Modulation (FiLM) layers.
- **Self-Supervised Genre Inference:** Use auxiliary classifiers trained on data sets like MagnaTagATune [111, 33] or GTZAN [181] to generate soft conditioning vectors during inference. The Milliom song data set remains also a subject of interest [128].
- **Genre-Tuned Attention:** Dynamically alter attention weights based on genre identity, modifying the receptive field or frequency prioritization to reflect domain-specific mastering norms.

Incorporating genre priors would enable more musically coherent enhancement

and empower user-facing systems to align with aesthetic expectations. This mirrors trends in vision and text generation (e.g., CLIP [20], DALL·E 3 [21]), where style conditioning is embedded directly into generative workflows.

10.2 STEM-AWARE RESTORATION

While Token U-Net is designed for full-mix, stemless restoration, incorporating stem-awareness could yield more precise enhancements. Prior works have shown that source-aware models can better apply targeted processing (e.g., drum-specific EQ or vocal de-essing) [174, 157].

Several architectures could support stem-aware extension:

- **Multi-Source Token Encoders:** Tokenize each stem independently and process them through a hierarchical U-Net with inter-stem attention and cross-token fusion.
- **Source Separation Frontend:** Use a separator (e.g., Conv-TasNet, Open-Unmix) to preprocess the mix and then apply parallel Token U-Nets to each stem before recombination.
- **Conditioned Token Diffusion:** Condition restoration on a learned embedding of stem identity, allowing the same model to generalize across instruments.

Such approaches would increase computational cost but provide fine-grained control over artifact suppression and spectral balance. They could be integrated with our existing curriculum design by gradually increasing stem count and complexity [154].

10.3 USER-CONTROLLED ENHANCEMENT

Token U-Net currently performs automatic zero-parameter restoration, robust but limited in expressive flexibility. In future iterations, the integration of user-controllable

conditioning would bring the system closer to an intelligent assistant.

Artists could specify goals such as:

- “Make this warmer and less compressed”
- “Retain reverb tail but reduce harshness”
- “Clean the mids without flattening the dynamics”

Possible implementation paths include:

1. **Text Prompts:** Use transformer-based encoders to generate embeddings from user instructions and condition the U-Net via FiLM or attention modulation [26].
2. **Reference Audio Conditioning:** Analyze a reference clip using EnCodec or CLAP-like models and align its token distribution with the input.
3. **Latent Edit Vectors:** Learn difference vectors between degraded and target clips, allowing directional enhancement in token space.

This paradigm aligns with real-world workflows, where mix engineers iteratively respond to creative intent. Bridging this gap transforms the model from a tool into a collaborator [127].

10.4 PROMPT-BASED AND TRANSFORMER REMIXING

Chapters 2 and 6 charted the evolution from time-domain restoration to token-based U-Nets. The next frontier lies in zero-shot and prompt-driven systems, inspired by advances in LLMs and multimodal models.

Token U-Net could evolve to support:

- **Multimodal Token Fusion:** Combine audio tokens with conditioning input (e.g., genre tags, text prompts, spectrogram hints) in a shared transformer backbone.
- **Semantic Remixing:** Use edit vectors to apply transformations like “make this more ambient” or “enhance vocals” in the style of MusicLM [1] or Jukebox [43].
- **Unsupervised Prompt Tuning:** Learn token-level transformations using contrastive or variational objectives, enabling label-free sonic intent modeling.

This would make audio restoration programmable and remixable, empowering creators with minimal technical training to express abstract sonic goals via natural language or example clips.

10.5 REAL-TIME DEPLOYMENT

Our current pipeline operates in offline batch mode, suitable for mastering but impractical for live applications. Real-time versions of Token U-Net could enable:

- **Live Enhancement:** Real-time correction of room coloration or mic artifacts during performance or streaming.
- **Streaming Pipelines:** On-the-fly improvement of podcast, Zoom, or Twitch audio using causal token processing.
- **Mobile and Embedded Use:** Deployment on edge devices for field recording or low-latency cleanup in vlogs or interviews.

Real-time adaptation would require:

- **Streaming-Compatible Architecture:** Chunk-wise token processing with context overlap, akin to Whisper or online ASR.

- **Causal Upsampling:** Replace transposed convolutions with causal interpolation kernels to maintain time alignment.
- **Model Compression:** Reduce the 1.08B parameter count via pruning, distillation, or quantization for edge deployment [35].

Such systems would bring professional-grade audio restoration to environments where offline pipelines are infeasible—newsrooms, classrooms, or low-resource field-work [187, 180].

Across these domains, the core vision remains: *restoration that listens—not just to the input signal, but to the user’s intent.*

CHAPTER 11

CONCLUSION

11.1 REFLECTING ON THE PROBLEM AND CONTRIBUTIONS

This thesis began with a deceptively simple question: *Can we build a system that restores and enhances full musical mixes - mastering them without access to stems or metadata?* Answering this question required addressing deep challenges in perceptual fidelity, genre diversity, and multi-effect degradation. Unlike previous systems focused on narrow tasks such as dereverberation [141], echo cancelation [3], or compression modeling [124], our goal was holistic restoration of stereo mixdowns without structural annotations.

The result is **Token-UNet**—a novel architecture trained entirely on discrete Encoder tokens, that while not capable of *mastering*, is capable of learning multi-effect restoration strategies without relying on inverse effect modeling. Using quantized latent spaces (Chapter 5), curriculum-aware degradation (Chapter 7), and a 1.08B U-Net parameter with attention to gating, FiLM and CBAM (Chapter 6), we developed a system that learns to enhance music in a data-driven, yet interpretable way.

Through extensive experiments (Chapter 8), we showed that when it comes to full-mix music training in the token space, it significantly outperforms baseline systems such as HiFi-GAN, Griffin-Lim, and parameter estimation networks. PESQ

improvements averaged more than +1.1 points, training time and computation power decreased by orders of magnitude, while the token level and spectral metrics confirmed restoration quality. Importantly, our model works *without* needing to undo each effect explicitly: it learns a restoration prior over compressed, semantically meaningful audio tokens.

Token space modeling also provided statistical clarity. Using token histograms, entropy plots, and coverage metrics, we analyzed degradation collapse and model improvement in an interpretable domain. This made debugging, curriculum transitions, and architectural validation tractable.

In sum, this thesis presents not just a denoising algorithm, but a multi-purpose neural music restoration framework, capable of learning audio enhancement as a structured, perceptual, and semantic transformation in the token domain.

11.2 FINAL THOUGHTS AND BROADER IMPACT

At its core, Token-UNet addresses a longstanding asymmetry in the music world: professional mastering has remained expensive, inaccessible, and manual, even as music creation has become more democratized. This thesis, while novice in the grand scheme, embarks upon the path towards a professional intelligent, data-driven, and interpretable system that makes high-quality audio restoration accessible to all.

Our approach introduces several technical advances: compact EnCodec representations, a compact U-Net for token enhancement, a five-stage curriculum with reproducible degradations, and a loss pipeline aligned with perceptual and musical outcomes. But its impact extends beyond the lab. Token-UNet allows:

- Garage bands to master recordings from phones.

- Podcasters to clean up room noise without plugins.
- Field researchers restore differently degraded archival audio.

While Token-UNet does not aim to replace human engineers, it creates a powerful alternative: a collaborative system that listens, learns, and enhances. By giving users agency, periodically through prompts, reference clips, or genre conditioning, it bridges artistry and automation.

Technically, this work sits at the intersection of self-supervised audio modeling, curriculum design, and generative restoration. It paves the way for future systems that are not only performant, but lightweight, explainable, and adaptable to creative intent.

11.3 TOWARD THE FUTURE OF NEURAL AUDIO

11.3.1 Bridging Creativity and Computation

Mastering is a uniquely creative form of engineering. It is about contrast, balance, and emotional translation. But access to that level of nuance has traditionally required studio-grade tools. Today, with token-based modeling, we are one step closer to closing that gap.

Token-UNet integrates audio signal structure, large-scale ML architecture, and perceptual training into a single system that can learn restoration behaviors directly from data. Compared to task-specific models like Speech dereverb [141], echo cancellation [75], or single-ended MOS prediction [35], our model generalizes across many degradation types *simultaneously*, often with no hand-tuned parameters.

11.3.2 Summary of Key Innovations

This work contributes several core advances:

- **Token-Based Audio Representation:** Leveraging EnCodec’s 16-codebook RVQ system, we model discrete audio tokens instead of waveforms or mel bins. This improves memory efficiency and enables interpretable analysis.
- **Token-UNet Architecture:** A 1.08B parameter U-Net with FiLM, CBAM, gated skips, and output heads for gain, masking, stereo, compression, and perceptual control. This model supports multi-effect restoration in a fully convolutional, token-domain space.
- **Curriculum Training Pipeline:** A five-stage system introducing progressively stronger degradations, including EQ, gain mismatch, reverb, echo, and compression. Precomputed token bundles (.pt format) enabled high-throughput and reproducible training.
- **Multi-Metric Evaluation:** Restoration was assessed using SNR, PESQ, STOI, mel-L1, spectral convergence and qualitative listening. We also visualized token divergence, coverage, and entropy to diagnose model behavior.
- **Accessibility and Realism:** The system was trained on real-world stereo tracks without isolated stems, simulating realistic production contexts such as voice memos, demos, or home studio recordings.

11.3.3 Positioning and Impact

Token-UNet differs from prior approaches such as spectrogram inversion [133], parameter regression [132], or effect-specific U-Nets [129]. It addresses the harder problem of unannotated joint restoration, modeling complex mixes with unknown overlapping

degradations. This positions it as both a technical milestone and a new direction in neural audio modeling.

Beyond research, the cultural implications are profound. By offering tools that understand and improve sound, not just clean it, Token-unet invites a wider public to professional-quality music making. Whether used by students, independent artists, or content creators, the system embodies a vision of intelligent, inclusive audio enhancement.

11.3.4 Personal Reflection

This thesis grew out of lived experience, recording, mixing and mastering music in studios and bedrooms alike. It is both technical and musical, engineering and art. Every equation was written with the ears in mind; every model was tested against real musical meaning. In the end, what emerged, while not the algorithm I set out to create, was a new way of thinking about sound.

CHAPTER A

LIST OF ACRONYMS

AI Artificial Intelligence

API Application Programming Interface

CBAM Convolutional Block Attention Module

CLI Command-Line Interface

CPU Central Processing Unit

CSV Comma-Separated Values

CNN Convolutional Neural Network

DSP Digital Signal Processing

EMA Exponential Moving Average

FP16 16-bit Floating Point Precision

FP32 32-bit Floating Point Precision

GAN Generative Adversarial Network

GL Griffin-Lim

GPU Graphics Processing Unit

L1 L1 Distance (Mean Absolute Error)

LSTM Long Short-Term Memory

Mel Mel-Frequency Spectral Coefficients

ML Machine Learning

MSE Mean Squared Error

OOM Out Of Memory

PESQ Perceptual Evaluation of Speech Quality

PT PyTorch Serialized Model File

ReLU Rectified Linear Unit

RNN Recurrent Neural Network

RTX Ray Tracing Texel eXtreme (NVIDIA GPU family)

SC Spectral Convergence

SNR Signal-to-Noise Ratio

STFT Short-Time Fourier Transform

STOI Short-Time Objective Intelligibility

UI User Interface

UNet Universal Network Architecture

VAE Variational Autoencoder

VOC Voice Over Codec

ASR Automatic Speech Recognition

BSS Blind Source Separation

DNN Deep Neural Network

FMA Free Music Archive

FSD Freesound Dataset

ICLR International Conference on Learning Representations

ICASSP International Conference on Acoustics, Speech, and Signal Processing

MIR Music Information Retrieval

MUSDB Music Source Database

NSynth Neural Synthesizer Dataset

PCA Principal Component Analysis

RIR Room Impulse Response

SDR Signal-to-Distortion Ratio

SOTA State Of The Art

SSM State-Space Model

TTS Text-To-Speech

VQ Vector Quantization

WAV Waveform Audio File Format

RVQ Residual Vector Quantization

FiLM Feature-wise Linear Modulation

KL Kullback–Leibler (Divergence)

HiFi-GAN High-Fidelity Generative Adversarial Network

STFT-L1 L1 Loss over Short-Time Fourier Transform Magnitudes

CE Cross-Entropy (Loss)

CRN Convolutional Recurrent Network

GCRN Gated Convolutional Recurrent Network

MoS Mean Opinion Score

MOSRA Multi-dimensional Objective Speech Quality Assessment

CLAP Contrastive Language-Audio Pretraining

LLM Large Language Model

ZSL Zero-Shot Learning

SLR Supervised Learning Rate

LoRA Low-Rank Adaptation (if you mention transformer fine-tuning in future work)

DDSP Differentiable Digital Signal Processing

OSA Output-Specific Architecture

CTC Connectionist Temporal Classification (if used for alignment)

SR Sample Rate

RT Real-Time

QoE Quality of Experience

QoS Quality of Service

UHD Ultra-High Definition (in music/audio context)

CHAPTER B

GLOSSARY OF TERMS

Term	Definition
Activation Function	A nonlinear function (e.g., ReLU, sigmoid) applied after each layer in a neural network to introduce nonlinearity.
Attention Mechanism	A method that allows models to focus on specific parts of an input, widely used in transformers and CBAM modules.
Autoencoder	A neural network architecture trained to reconstruct its input, often used for unsupervised representation learning.
Bandwidth (kbps)	The data rate in kilobits per second used in audio compression to control fidelity vs. size.
Batch Size	Number of samples processed together in one forward/backward pass during training.
Bottleneck	A reduced-dimension representation inside a model; our architecture optionally applies this to force compression.
Chroma Features	A representation of pitch class energy (C, C#, D, etc.), used to capture harmonic content.
Codebook	A finite set of quantized representations (discrete vectors) used in vector quantization within EnCodec.
Convolution	A mathematical operation sliding a kernel over input data; forms the backbone of CNNs.

Continued on next page

Table B.1 – continued from previous page

Term	Definition
Curriculum Learning	A training strategy where the model is exposed to gradually harder tasks to improve stability and generalization.
Decoder	In a U-Net or autoencoder, the portion of the network that reconstructs the output from a compressed latent space.
Degradation	Any intentional corruption of audio used for training, such as reverb, gain mismatch, or EQ.
Dropout	A regularization technique that randomly sets activations to zero during training to prevent overfitting.
Dynamic Range Compression	The process of reducing the range between the loudest and quietest parts of an audio signal.
Embedding	The mapping from discrete tokens to dense vector representations used for neural network input.
Epoch	A full pass through the entire training dataset.
Equalization (EQ)	The process of adjusting the amplitude of specific frequency bands.
Feature Map	The output of a convolutional layer; a spatial or temporal representation of learned patterns.
FiLM Layer	Feature-wise Linear Modulation; applies conditional scaling and bias to features based on external input.
Forward Pass	The process of passing input data through a neural network to obtain predictions.
Gradient Descent	Optimization algorithm used to minimize loss by updating parameters in the direction of the gradient.

Continued on next page

Table B.1 – continued from previous page

Term	Definition
Griffin-Lim	An iterative algorithm for reconstructing waveforms from magnitude spectrograms.
Ground Truth	The original, uncorrupted target data used as a reference during supervised learning.
Hann Window	A windowing function used in STFT to minimize spectral leakage.
Inference	Running a trained model to obtain predictions from unseen data (i.e., test-time execution).
Latent Space	The compressed internal representation of data in a model, typically in autoencoders or bottlenecks.
Layer Normalization	A technique for normalizing inputs across features within a layer to improve training convergence.
Loss Function	A metric (e.g., cross-entropy, MSE) used to quantify the error between predicted and ground-truth values.
Mel Spectrogram	A spectrogram scaled to match human perception, with frequency bins spaced according to the mel scale.
Normalization	Rescaling inputs to a common range, often 0 to 1 or -1 to 1, to stabilize training.
OneCycleLR	A learning rate scheduler that increases then decreases LR within each training cycle.
Padding	Extra zeros added to input so that convolution operations maintain dimensions.
Perceptual Loss	A loss that aligns predictions not numerically but perceptually, using audio features or pretrained models.

Continued on next page

Table B.1 – continued from previous page

Term	Definition
Residual Block	A block in neural networks that adds the input to the output after a series of layers, enabling deeper training.
ReLU	Rectified Linear Unit; an activation function defined as $f(x) = \max(0, x)$.
Sampling Rate	Number of audio samples per second, typically 44.1kHz or 48kHz in our case.
Signal-to-Noise Ratio (SNR)	A measure of how much desired signal exists relative to noise or error.
Skip Connection	A shortcut path that bypasses certain layers and helps preserve information across depth.
Spectral Convergence	A loss measuring the Frobenius norm between predicted and ground truth spectrograms.
STFT	Short-Time Fourier Transform; decomposes audio into time-frequency domain using sliding windows.
STOI	Short-Time Objective Intelligibility; measures speech intelligibility, extended here to music clarity.
Token	A discrete symbol representing audio, output by the quantizer in EnCodec.
Token Histogram	A count distribution of token usage across the dataset or output predictions.
Transformer	A deep learning model based entirely on attention mechanisms.
U-Net	A symmetric encoder-decoder architecture with skip connections, originally used for biomedical segmentation.

Continued on next page

Table B.1 – continued from previous page

Term	Definition
Upsampling	The process of increasing the resolution of data in the decoder; done here using ConvTranspose1d layers.
Vector Quantization	The process of mapping continuous values to discrete tokens using a codebook.
Weight Decay	A regularization term that penalizes large weights to reduce overfitting.
Zero-shot Learning	Model inference on tasks not seen during training, often guided by prompts or examples.
Curriculum Stage	A specific training tier defined by the complexity of degradations applied to the input audio. Used to progressively increase task difficulty in model training.
EnCodec	A neural audio codec by Meta AI that converts waveforms into discrete tokens via residual vector quantization.
Entropy (Token Entropy)	A measure of uncertainty or diversity in token distributions across codebooks, used to assess model generalization and degradation severity.
FiLM Modulation	Conditional modulation method that applies learned scale and bias vectors to intermediate network features based on auxiliary inputs.
Gated Skip Connection	A skip connection whose contribution is modulated by a learnable gate, allowing the network to adjust feature blending between encoder and decoder.

Continued on next page

Table B.1 – continued from previous page

Term	Definition
Heatmap (Token)	A 2D visualization of token usage over time across codebooks, revealing structure, continuity, and anomalies in the latent token space.
KL Divergence	A statistical measure of how one probability distribution diverges from another; used to compare clean and degraded token distributions.
Loss Weighting	A strategy where different loss components (e.g., cross-entropy, STFT, perceptual) are scaled to balance their influence during training.
Perceptual Embedding	A learned vector that captures perceptually relevant features of audio, typically used to improve restoration alignment with human hearing.
Quantization Noise	The error introduced when continuous values are discretized into tokens using vector quantization.
SHA256 Hash	A cryptographic hash function used to verify the integrity of each audio file before and after degradation.
Token Coverage	The proportion of codebook bins activated by a token stream, used to assess representational richness or collapse.
Token U-Net	The main architecture proposed in this thesis; a U-Net operating directly in EnCodec token space, with auxiliary heads for multi-effect restoration.
Waveform Cache	Precomputed and stored version of the original/degraded audio waveform used for fast inference and reproducibility.

CHAPTER C

EXTENDED EQUATIONS AND DERIVATIONS

This appendix presents mathematical foundations and detailed derivations related to the Token-UNet architecture, EnCodec quantization, loss functions, training stability, and audio signal processing techniques used throughout the thesis.

C.1 SHORT-TIME FOURIER TRANSFORM (STFT)

The STFT of a discrete-time signal $x[n]$ using window $w[n]$ is:

$$X(m, \omega) = \sum_{n=-\infty}^{\infty} x[n]w[n - mH]e^{-j\omega n} \quad (\text{C.1})$$

where H is the hop size. We typically use a Hann window:

$$w[n] = 0.5 \left(1 - \cos \left(\frac{2\pi n}{N-1} \right) \right) \quad (\text{C.2})$$

C.2 MEL SPECTROGRAM

Given magnitude spectrum $|X(f, t)|$, the mel spectrogram is:

$$M(m, t) = \sum_{f=1}^F H_{m,f} \cdot |X(f, t)|^2 \quad (\text{C.3})$$

where H is the mel filterbank and m indexes mel bands.

C.3 ENCODEC QUANTIZATION AND TOKENIZATION

C.3.1 Vector Quantization

Let $z \in \mathbb{R}^d$ be a latent vector. The quantized token $Q(z)$ is:

$$Q(z) = \arg \min_{c_k \in C} \|z - c_k\|^2, \quad C = \{c_1, \dots, c_K\} \quad (\text{C.4})$$

C.3.2 Multi-Codebook Tokenization

For n_q codebooks, each latent frame $z[t]$ is decomposed into $z_1[t], \dots, z_{n_q}[t]$, and tokenized:

$$x_q[t] = Q_q(z_q[t]) \in \{0, \dots, K-1\}, \quad \text{for } q = 1 \dots n_q \quad (\text{C.5})$$

C.4 GRIFFIN-LIM RECONSTRUCTION

Given magnitude spectrogram $|X(f, t)|$, reconstruct phase via:

$$\hat{x} = \text{ISTFT}(|X|, \angle \hat{X}) \quad (\text{C.6})$$

using iterative phase estimation:

$$\hat{X}^{(k+1)} = |X| \cdot \frac{\text{STFT}(\text{ISTFT}(\hat{X}^{(k)}))}{|\text{STFT}(\text{ISTFT}(\hat{X}^{(k)}))|} \quad (\text{C.7})$$

C.5 HIFI-GAN GENERATOR LOSSES

The adversarial loss:

$$\mathcal{L}_{\text{adv}} = \mathbb{E}_x[(D(G(x)) - 1)^2] \quad (\text{C.8})$$

The feature matching loss:

$$\mathcal{L}_{\text{FM}} = \sum_{i=1}^T \frac{1}{N_i} \|D^{(i)}(x) - D^{(i)}(G(x))\|_1 \quad (\text{C.9})$$

The mel-spectrogram reconstruction loss:

$$\mathcal{L}_{\text{mel}} = \|\text{Mel}(x) - \text{Mel}(G(x))\|_1 \quad (\text{C.10})$$

C.6 TOKEN-UNET OUTPUT SHAPES

Model output logits:

$$L \in \mathbb{R}^{B \times K \times n_q \times T} \quad (\text{C.11})$$

where B = batch size, K = codebook size (typically 1024), $n_q = 16$.

C.6.1 Token Prediction (Cross-Entropy)

$$\mathcal{L}_{\text{CE}} = - \sum_{i,q,t} \log \text{softmax}(L_{i,:,q,t})[y_{i,q,t}] \quad (\text{C.12})$$

C.6.2 Auxiliary Heads

- **Mask head (dereverb)**: predicts binary mask $\mathbf{M} \in [0, 1]^T$
- **Gain head**: predicts scalar per frame $\mathbf{g}_t \in \mathbb{R}$
- **Stereo head**: predicts stereo width $\mathbf{s}_t \in [0, 1]$
- **Compression head**: predicts dynamic envelope $\mathbf{c}_t \in \mathbb{R}^d$
- **Perceptual head**: predicts learned perceptual embedding vector

C.7 PERCEPTUAL AND TOKEN LOSSES

$$\mathcal{L}_{\text{STFT}} = \||\text{STFT}(x)| - |\text{STFT}(\hat{x})|\|_1 \quad (\text{C.13})$$

$$\mathcal{L}_{\text{mel}} = \|\text{Mel}(x) - \text{Mel}(\hat{x})\|_1 \quad (\text{C.14})$$

$$\mathcal{L}_{\text{token}} = \|h_q^{\text{STFT}} - \hat{h}_q^{\text{STFT}}\|_1 \quad (\text{across codebooks}) \quad (\text{C.15})$$

C.8 KL DIVERGENCE BETWEEN TOKEN DISTRIBUTIONS

Let $\mathbf{h}^{(clean)}$ and $\mathbf{h}^{(deg)}$ be token histograms:

$$D_{\text{KL}}(\mathbf{h}^{(clean)} \parallel \mathbf{h}^{(deg)}) = \sum_{k=1}^K p_k^{(c)} \log \left(\frac{p_k^{(c)}}{p_k^{(d)} + \epsilon} \right) \quad (\text{C.16})$$

C.9 CURRICULUM ADVANCEMENT CONDITION

Stage advancement occurs when:

$$|\text{EMA}_{\text{val}}^{(t)} - \text{EMA}_{\text{val}}^{(t-5)}| < \delta, \quad \text{and } t > t_{\min} \quad (\text{C.17})$$

or if training loss stagnates:

$$\frac{1}{8} \sum_{k=0}^7 \mathcal{L}_{\text{train}}^{(t-k)} - \mathcal{L}_{\text{train}}^{(t)} < \epsilon_{\text{plateau}} \quad (\text{C.18})$$

C.10 ENTROPY OF TOKEN DISTRIBUTIONS

$$H_q = - \sum_{k=1}^K p_k \log_2(p_k + \epsilon) \quad (\text{C.19})$$

Lower H_q implies token collapse or degradation; high entropy indicates broad coverage.

C.11 QUANTIZED LATENT RECONSTRUCTION ERROR

Assume ground truth z and decoded $\hat{z}$:

$$\mathcal{L}_{\text{VQ}} = \|\text{sg}(z) - c\|_2^2 + \beta \|z - \text{sg}(c)\|_2^2 \quad (\text{C.20})$$

where $\text{sg}(\cdot)$ is stop-gradient and c is the nearest codebook vector.

C.12 TOKEN COVERAGE RATIO

Let $\mathbf{h}_q[k]$ be token histogram for codebook q :

$$\gamma_q = \frac{|\{k : \mathbf{h}_q[k] > 0\}|}{K} \quad (\text{C.21})$$

C.13 CURRICULUM-AWARE LOSS SCHEDULING

Let $\lambda_i^{(s)}$ be the weight for loss i at stage s . Then:

$$\mathcal{L}_{\text{total}}^{(s)} = \sum_i \lambda_i^{(s)} \cdot \mathcal{L}_i \quad (\text{C.22})$$

Weights are increased for perceptual and auxiliary heads in later stages.

C.14 STREAMING INFERENCE CONSTRAINTS (REAL-TIME)

Let Δt be the inference window and C be context:

$$X_{\text{stream}}^{(t)} = \text{TokenUNet}(X_{t-C:t+\Delta t}) \quad (\text{C.23})$$

Causal inference requires:

$$\forall t, \text{Prediction}(t) \notin \text{future}(x) \quad (\text{C.24})$$

BIBLIOGRAPHY

- [1] Andrea Agostinelli, Timo I. Denk, Zalán Borsos, Jesse Engel, Mauro Verzetti, Antoine Caillon, Qingqing Huang, Aren Jansen, Adam Roberts, Marco Tagliasacchi, Matt Sharifi, Neil Zeghidour, and Christian Frank. Musiclm: Generating music from text, 2023.
- [2] María Alfaro-Contreras, Antonio Ríos-Vila, Jose J. Valero-Mas, and Jorge Calvo-Zaragoza. A transformer approach for polyphonic audio-to-score transcription. In *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 706–710, 2024.
- [3] M. Ali. Stereophonic acoustic echo cancellation system using time-varying all-pass filtering for signal decorrelation. In *1998 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 3689–3692, Seattle, WA, USA, 1998. IEEE.
- [4] Jont B. Allen and David A. Berkley. Image method for efficiently simulating small-room acoustics. *Journal of the Acoustical Society of America*, 65(4):943–950, 1979.
- [5] Pavel Andreev, Aibek Alanov, Oleg Ivanov, and Dmitry Vetrov. Hifi++: a unified framework for neural vocoding, bandwidth extension and speech enhancement. *arXiv preprint arXiv:2203.13086*, 2022.
- [6] Rosana Ardila, Megan Branson, Kelly Davis, Michael Henretty, Michael Kohler, Josh Meyer, Reuben Morais, Lindsay Saunders, Francis M Tyers, and Gregor Weber. Common voice: A massively-multilingual speech corpus. *arXiv preprint arXiv:1912.06670*, 2019.

- [7] Daniel Arteaga. *Introduction to Ambisonics*. Universitat Pompeu Fabra, Barcelona, Spain, 2015.
- [8] Bishnu S Atal and Suzanne L Hanauer. Speech analysis and synthesis by linear prediction of the speech wave. *The Journal of the Acoustical Society of America*, 50(2B):637–655, 1971.
- [9] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [10] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(12):2481–2495, 2017.
- [11] Johannes Ballé, Nick Johnston, and David Minnen. Integer networks for data compression with latent-variable models. In *International Conference on Learning Representations (ICLR)*, 2018.
- [12] Johannes Ballé, Valero Laparra, and Eero P Simoncelli. End-to-end optimized image compression. In *International Conference on Learning Representations (ICLR)*, 2017.
- [13] Bastian Bechtold, Matthias Geier, et al. Soundfile: an audio library based on libsndfile, cffi, and numpy. <https://github.com/bastibe/python-soundfile>, 2014. Version 0.13.1, Accessed: 2025-06-02.
- [14] Yoshua Bengio. Learning deep architectures for ai. *Foundations and Trends in Machine Learning*, 2(1):1–127, 2009.
- [15] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.

- [16] Dipesh Bhagat, Ninad Bhatt, and Yogeshwar Kosta. Adaptive multi-rate wide-band speech codec based on celp algorithm: architectural study, implementation & performance analysis. In *2012 International Conference on Communication Systems and Network Technologies*, pages 547–551. IEEE, 2012.
- [17] Katz Bob. *Mastering Audio: The Art and the Science*. Focal Press, 2nd edition, 2007.
- [18] Dmitry Bogdanov, Minz Won, Philip Tovstogan, Alastair Porter, and Xavier Serra. The mtg-jamendo dataset for automatic music tagging. In *Machine Learning for Music Discovery Workshop, International Conference on Machine Learning (ICML 2019)*, Long Beach, CA, United States, 2019.
- [19] Steven Boll. Suppression of acoustic noise in speech using spectral subtraction. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 27(2):113–120, 1979.
- [20] Zalán Borsos, Raphaël Marinier, Damien Vincent, Eugene Kharitonov, Olivier Pietquin, Matt Sharifi, Dominik Roblek, Olivier Teboul, David Grangier, Marco Tagliasacchi, and Neil Zeghidour. Audiolm: a language modeling approach to audio generation, 2023.
- [21] Zalán Borsos, Matt Sharifi, Damien Vincent, Eugene Kharitonov, Neil Zeghidour, and Marco Tagliasacchi. Soundstorm: Efficient parallel audio generation, 2023.
- [22] brunet pascal, li yuan, and kim soohyun. application of ml-based time series forecasting to audio dynamic range compression. *journal of the audio engineering society*, (112), october 2023.
- [23] Heiko Buchner, Jacob Benesty, Tomas Gansler, and Walter Kellermann. Robust extended multidelay filter and double-talk detector for acoustic echo cancella-

- tion. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 14(5):1633–1644, 2006.
- [24] Heiko Buchner and Sascha Spors. A general derivation of wave-domain adaptive filtering and application to acoustic echo cancellation. In *2008 42nd Asilomar Conference on Signals, Systems and Computers*, pages 816–823, Pacific Grove, CA, USA, 2008. IEEE.
- [25] Chih-Cheng Chang and Li Su. Beast: Online joint beat and downbeat tracking based on streaming transformer, 2024.
- [26] Ernie Chang, Sidd Srinivasan, Mahi Luthra, Pin-Jie Lin, Varun Nagaraja, Forrest Iandola, Zechun Liu, Zhaoheng Ni, Changsheng Zhao, Yangyang Shi, and Vikas Chandra. On the open prompt challenge in conditional audio generation. In *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5315–5319, 2024.
- [27] Jingdong Chen, Dong Yu, Morten Kolbæk, Zheng-Hua Tan, and Jesper Jensen. Large-scale training to increase speech intelligibility for hearing-impaired listeners in novel noises. *Journal of the Acoustical Society of America*, 139:2604–2612, 2016.
- [28] Zi-yang Chen, Xiaolong Yang, Yujie Jiang, Yi Deng, Shuo Liu, Jing Yu, Yinan Yang, and Yue Wang. Encodec: High fidelity audio compression with neural codes. *arXiv preprint arXiv:2303.05292*, 2023.
- [29] Lingyun Cheng, Rui Peng, Aoxue Li, Chenyu Zheng, and Xiaofei Li. Deep learning-based stereophonic acoustic echo suppression without decorrelation. *The Journal of the Acoustical Society of America*, 150:816–829, 2021.
- [30] Lingyun Cheng, Chenyu Zheng, Aoxue Li, Rui Peng, and Xiaofei Li. A deep complex network with multi-frame filtering for stereophonic acoustic echo can-

- cellation. In *Proceedings of the 23rd INTERSPEECH Conference*, pages 2508–2512, Incheon, Korea, 2022.
- [31] Michael Chinen, Felicia SC Lim, Jan Skoglund, Nikita Gureev, Feargus O’Gorman, and Andrew Hines. Visqol v3: An open source production ready objective speech and audio metric. In *2020 Twelfth International Conference on Quality of Multimedia Experience (QoMEX)*, pages 1–6. IEEE, 2020.
- [32] Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014.
- [33] Keunwoo Choi. magnatagatune-list: List of automatic music tagging research articles evaluated against magnatagatune dataset. <https://github.com/keunwoochoi/magnatagatune-list>, 2017. Accessed: 2025-06-02.
- [34] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). *arXiv preprint arXiv:1511.07289*, 2015.
- [35] Jozef Coldenhoff, Andrew Harper, Paul Kendrick, Tijana Stojkovic, and Milos Cernak. Multi-channel mosra: Mean opinion score and room acoustics estimation using simulated data and a teacher model, 2024.
- [36] Jade Copet, Alexandre Défossez, Gabriel Synnaeve, Yossi Adi, Antoine Caillon, Jesse Engel, Yuxuan Wang, Matt Sharifi, and Neil Zeghidour. Simple and controllable music generation, 2023.

- [37] Jerome Daniel, Jean-Baptiste Rault, and Jean-Dominique Polack. Ambisonics encoding of other audio formats for multiple listening conditions. In *105th Convention of the Audio Engineering Society*, San Francisco, CA, USA, 1998.
- [38] Hugo T. de Carvalho, Fabricio R. Avila, and Luiz W. P. Biscainho. Bayesian restoration of audio degraded by low-frequency pulses modeled via gaussian process. *IEEE Journal of Selected Topics in Signal Processing*, 15(1):90–103, 2021.
- [39] Brecht De Man, Ben Leonard, Ryan King, and Joshua D. Reiss. An analysis and evaluation of audio features for multitrack music mixtures. In *15th International Society for Music Information Retrieval Conference (ISMIR)*, 2014.
- [40] Michaël Defferrard, Kirell Benzi, Pierre Vandergheynst, and Xavier Bresson. FMA: A dataset for music analysis. In *18th International Society for Music Information Retrieval Conference (ISMIR)*, 2017.
- [41] Michaël Defferrard, Sharada P. Mohanty, Sean F. Carroll, and Marcel Salathé. Learning to recognize musical genre from audio. In *The 2018 Web Conference Companion*. ACM Press, 2018.
- [42] Alexandre Defossez, Gabriel Synnaeve, and Yossi Adi. Real time speech enhancement in the waveform domain. *arXiv preprint arXiv:2006.12847*, 2020.
- [43] Prafulla Dhariwal, Heewoo Jun, Christine Payne, Jong Wook Kim, Alec Radford, and Ilya Sutskever. Jukebox: A generative model for music. *arXiv preprint arXiv:2005.00341*, 2020.
- [44] Sander Dieleman, Aaron van den Oord, and Karen Simonyan. The challenge of realistic music generation: modelling raw audio at scale. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

- [45] Martin Dietz, Markus Multrus, Vaclav Eksler, Vladimir Malenovsky, Erik Norvell, Harald Pobloth, Lei Miao, Zhe Wang, Lasse Laaksonen, Adriana Vasilache, et al. Overview of the evs codec architecture. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5698–5702. IEEE, 2015.
- [46] Audio Dimensions. Mastering audio samples—before and after mastering. <https://example.com/audio-mastering-before-after>, 2015.
- [47] Tim Dittmar. *Audio Engineering 101: A Beginner’s Guide to Music Production*. Routledge, 2nd edition, 2017.
- [48] Paul Douglas. How long does it take to learn mixing and mastering? <https://homemusiccreator.com/how-long-to-learn-mixing-mastering/>, 2022. Accessed: 2025-06-02.
- [49] Harishchandra Dubey, Vishak Gopal, Ross Cutler, Sergiy Matuskevych, Sebastian Braun, Emre Sefik Eskimez, Manthan Thakker, Takuya Yoshioka, Hannes Gamper, and Robert Aichner. Icassp 2022 deep noise suppression challenge. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 9276–9280. IEEE, 2022.
- [50] Alexandre Défossez, Yossi Adi, and Gabriel Synnaeve. Differentiable model compression via pseudo quantization noise. *arXiv preprint arXiv:2104.09987*, 2021.
- [51] Alexandre Défossez, Jade Copet, Gabriel Synnaeve, and Yossi Adi. High fidelity neural audio compression. *arXiv preprint arXiv:2210.13438*, 2022.
- [52] Alexandre Défossez, Nicolas Usunier, Léon Bottou, and Francis Bach. Music source separation in the waveform domain. *arXiv preprint arXiv:1911.13254*, 2019.

- [53] Ahmet M. Elbir, H. Bilal Çam, M. Emre Iyican, B. Öztürk, and N. Aydin. Music genre classification and recommendation by using machine learning techniques. In *2018 Innovations in Intelligent Systems and Applications Conference (ASYU)*, pages 1–5. IEEE, 2018.
- [54] Satoshi Emura and Yoichi Haneda. A method of coherence-based step-size control for robust stereo echo cancellation. In *2003 IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 592–595, Hong Kong, 2003. IEEE.
- [55] Jesse H. Engel, Lamtharn Hantrakul, Chenjie Gu, and Adam Roberts. DDSP: differentiable digital signal processing. *CoRR*, abs/2001.04643, 2020.
- [56] F. Alton Everest and Ken C. Pohlmann. *Master Handbook of Acoustics*. McGraw-Hill Education, 6th edition, 2015.
- [57] Cédric Févotte, Nicolas Bertin, and Jean-Luc Durrieu. Nonnegative matrix factorization with the itakura-saito divergence: With application to music analysis. *Neural Computation*, 21(3):793–830, 2009.
- [58] Eduardo Fonseca, Xavier Favory, Jordi Pons, Frederic Font, and Xavier Serra. Fsd50k: an open dataset of human-labeled sound events. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 30:829–852, 2022.
- [59] Antonin Gagneré, Slim Essid, and Geoffroy Peeters. Adapting pitch-based self supervised learning models for tempo estimation. In *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 956–960, 2024.
- [60] Tomas Gansler and Peter Eneroth. Influence of audio coding on stereophonic acoustic echo cancellation. In *1998 IEEE International Conference on Acous-*

- tics, Speech and Signal Processing*, pages 3649–3652, Seattle, WA, USA, 1998. IEEE.
- [61] Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter. Audio set: An ontology and human-labeled dataset for audio events. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 776–780. IEEE, 2017.
- [62] Michael A. Gerzon. Ambisonics in multichannel broadcasting and video. *Journal of the Audio Engineering Society*, 33(11):859–871, 1985.
- [63] Karan Goel, Albert Gu, Chris Donahue, and Christopher Ré. It’s raw! audio generation with state-space models. *arXiv preprint arXiv:2202.09729*, 2022.
- [64] Robert Gray. Vector quantization. *IEEE ASSP Magazine*, 1(2):4–29, 1984.
- [65] D Griffin and Jae Lim. A new model-based speech analysis/synthesis system. In *ICASSP*, 1985.
- [66] Daniel Griffin and Jae Lim. Signal estimation from modified short-time fourier transform. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 32(2):236–243, Apr 1984.
- [67] Alexey Gritsenko, Tim Salimans, Rianne van den Berg, Jasper Snoek, and Nal Kalchbrenner. A spectral energy distance for parallel speech synthesis. In *Advances in Neural Information Processing Systems*, volume 33, pages 13062–13072, 2020.
- [68] Cristina Gârbacea, Aäron van den Oord, Yazhe Li, Felicia SC Lim, Alejandro Luebs, Oriol Vinyals, and Thomas C Walters. Low bit-rate speech coding with vq-vae and a wavenet decoder. In *ICASSP 2019-2019 IEEE International*

- Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 735–739. IEEE, 2019.
- [69] Tomas Gänslér and Jacob Benesty. The fast normalized cross-correlation double-talk detector. *Signal Processing*, 86(6):1124–1139, 2006.
- [70] Emanuel A. P. Habets and Jacob Benesty. A two-stage beamforming approach for noise reduction and dereverberation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 21(5):945–957, May 2013.
- [71] Hamed Haghbayan, Edward A. Coomes, and Daniel Curran. Temporal trends in the loudness of popular music over six decades. *Journal of General Internal Medicine*, 35:394–395, 2020.
- [72] Elmar Hänsler. The hands-free telephone problem—an annotated bibliography. *Signal Processing*, 27:259–271, 1992.
- [73] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015.
- [74] Florian Henkel, Rainer Kelz, and Gerhard Widmer. Audio-conditioned u-net for position estimation in full sheet images, 2019.
- [75] Jürgen Herre, Heiko Buchner, and Walter Kellermann. Acoustic echo cancellation for surround sound using perceptually motivated convergence enhancement. In *2007 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 117–120, Honolulu, HI, USA, 2007. IEEE.
- [76] Michael Hilsamer and Sascha Herzog. A statistical approach to automated offline dynamic processing in the audio mastering process. In *Proc. of the 17th International Conference on Digital Audio Effects (DAFx-14)*, pages 35–40, Erlangen, Germany, 2014.

- [77] Andrew Hines, Derry Fitzgerald, John Skoglund, Michael Harte, Niall Gillen, and Andrew Kokaram. ViSQOL: an objective speech quality model. *EURASIP Journal on Audio, Speech, and Music Processing*, 2015(1):1–18, 2015.
- [78] Andrew Hines, Jan Skoglund, Anil Kokaram, and Naomi Harte. Visqol: The virtual speech quality objective listener. In *IWAENC 2012; International Workshop on Acoustic Signal Enhancement*, pages 1–4. VDE, 2012.
- [79] Matthias Holters and Udo Zölzer. Gstpeaq – an open source implementation of the peaq algorithm. pages 1–4, 2015. Technical report.
- [80] Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhotia, Ruslan Salakhutdinov, and Abdelrahman Mohamed. Hubert: Self-supervised speech representation learning by masked prediction of hidden units. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:3451–3460, 2021.
- [81] Po-Sen Huang, Minje Kim, Mark Hasegawa-Johnson, and Paris Smaragdis. Joint optimization of masks and deep recurrent neural networks for monaural source separation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 23:2136–2147, Dec 2015.
- [82] Qingqing Huang, Daniel S. Park, Tao Wang, Timo I. Denk, Andy Ly, Nanxin Chen, Zhengdong Zhang, Zhishuai Zhang, Jiahui Yu, Christian Frank, Jesse Engel, Quoc V. Le, William Chan, Zhifeng Chen, and Wei Han. Noise2music: Text-conditioned music generation with diffusion models, 2023.
- [83] David Miles Huber and Robert E. Runstein. *Modern Recording Techniques*. Routledge, 9th edition, 2017.
- [84] Evgenii Indenbom, Nicolae-Catalin Ristea, Ando Saabas, Tanel Parnamaa, Jegor Guzvin, and Ross Cutler. Deepvqe: Real time deep voice quality enhance-

- ment for joint acoustic echo cancellation, noise suppression and dereverberation, 2023.
- [85] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning (ICML)*, pages 448–456, 2015.
- [86] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- [87] Ignacio Irigaray, Martin Rocamora, and Luiz W. P. Biscainho. Noise reduction in analog tape audio recordings with deep learning models. In *Audio Engineering Society International Conference on Audio Archiving, Preservation & Restoration*, Culpeper, VA, USA, June 2023. Audio Engineering Society.
- [88] Roey Izhaki. *Mixing Audio: Concepts, Practices, and Tools*. Routledge, New York, NY, USA, 3 edition, 2017.
- [89] Roey Izhaki. *Mixing Audio: Concepts, Practices, and Tools*. Focal Press, 3rd edition, 2018.
- [90] iZotope, Inc. About izotope. <https://www.izotope.com/en/company/about-us.html>, 2023. Accessed: 2025-06-02.
- [91] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In *International Conference on Learning Representations (ICLR)*, 2017.
- [92] Andreas Jansson, Eric Humphrey, Nicola Montecchio, Rachel Bittner, Aparna Kumar, and Tillman Weyde. Singing voice separation with deep u-net convolutional networks. In *Proceedings of the 18th International Society for Music*

- Information Retrieval Conference (ISMIR)*, Suzhou, China, 2017. Licensed under Creative Commons Attribution 4.0 International License (CC BY 4.0).
- [93] Tejas Jayashankar, Thilo Koehler, Kaustubh Kalgaonkar, Zhiping Xiu, Jilong Wu, Ju Lin, Prabhav Agrawal, and Qing He. Architecture for variable bitrate neural speech codec with configurable computation complexity. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 861–865. IEEE, 2022.
- [94] Xue Jiang, Xiulian Peng, Chengyu Zheng, Huaying Xue, Yuan Zhang, and Yan Lu. End-to-end neural speech coding for real-time communications. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 866–870. IEEE, 2022.
- [95] Biing-Hwang Juang and A Gray. Multiple stage vector quantization for speech coding. In *ICASSP’82. IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 7, pages 597–600. IEEE, 1982.
- [96] Peter Kabal. Tsp speech database. Technical report, Department of Electrical & Computer Engineering, McGill University, Montreal, Quebec, Canada, 2002.
- [97] Nal Kalchbrenner et al. Efficient neural audio synthesis. In *International Conference on Machine Learning (ICML)*, 2018.
- [98] Eugene Kharitonov, Ann Lee, Adam Polyak, Yossi Adi, Jade Copet, Kushal Lakhotia, Tu-Anh Nguyen, Morgane Rivi re, Abdelrahman Mohamed, Emmanuel Dupoux, et al. Text-free prosody-aware generative spoken language modeling. *arXiv preprint arXiv:2109.03264*, 2021.
- [99] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, San Diego, CA, USA, 2015.

- [100] W Bastiaan Kleijn, Felicia SC Lim, Alejandro Luebs, Jan Skoglund, Florian Stimberg, Quan Wang, and Thomas C Walters. Wavenet based low rate speech coding. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 676–680. IEEE, 2018.
- [101] W Bastiaan Kleijn, Andrew Storus, Michael Chinen, Tom Denton, Felicia SC Lim, Alejandro Luebs, Jan Skoglund, and Hengchin Yeh. Generative speech coding with predictive variance regularization. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6478–6482. IEEE, 2021.
- [102] Yuto Kondo, Hirokazu Kameoka, Kou Tanaka, and Takuhiro Kaneko. Selecting n-lowest scores for training mos prediction models. In *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1451–1455, 2024.
- [103] Jungil Kong, Jaehyeon Kim, and Jaekyoung Bae. Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis. In *Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS)*, Vancouver, Canada, 2020. arXiv:2010.05646.
- [104] Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. Dif-fwave: A versatile diffusion model for audio synthesis, 2021.
- [105] Damian Koszewski, Thomas Görne, Grazina Korvel, and Bozena Kostek. Automatic music signal mixing system based on one-dimensional wave-u-net autoencoders. *EURASIP Journal on Audio, Speech, and Music Processing*, 2023(1), 2023.
- [106] Felix Kreuk, Adam Polyak, Jade Copet, Eugene Kharitonov, Tu-Anh Nguyen, Morgane Rivi re, Wei-Ning Hsu, Abdelrahman Mohamed, Emmanuel Dupoux,

- and Yossi Adi. Textless speech emotion conversion using decomposed and discrete representations. *arXiv preprint arXiv:2111.07402*, 2021.
- [107] Kundan Kumar, Rithesh Kumar, Thibault de Boissiere, Lucas Gestin, Wei Zhen Teoh, Jose Sotelo, Alexandre de Brébisson, Yoshua Bengio, and Aaron C Courville. Melgan: Generative adversarial networks for conditional waveform synthesis. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- [108] Kushal Lakhotia, Eugene Kharitonov, Wei-Ning Hsu, Yossi Adi, Adam Polyak, Benjamin Bolte, Tu-Anh Nguyen, Jade Copet, Alexei Baevski, Abdelrahman Mohamed, et al. On generative spoken language modeling from raw audio. *Transactions of the Association for Computational Linguistics*, 9:1336–1354, 2021.
- [109] LANDR Audio Inc. About landr - the creative platform for musicians. <https://www.landr.com/about/>, 2025. Accessed: 2025-06-02.
- [110] Siddique Latif, Heriberto Cuayáhuitl, Farrukh Pervez, Fahad Shamshad, Hafiz Shehbaz Ali, and Erik Cambria. A survey on deep reinforcement learning for audio-based applications, 2021.
- [111] Edith Law, Luis Von Ahn, Roger Dannenberg, and Michael Crawford. Evaluation of algorithms using games: The case of music tagging. In *Proceedings of the 10th International Society for Music Information Retrieval Conference (ISMIR)*, pages 387–392, 2009.
- [112] Jonathan Le Roux, John R. Hershey, Scott T. Wisdom, and Hakan Erdogan. Sdr: Half-baked or well done? Technical report, Mitsubishi Electric Research Laboratories, Cambridge, MA, USA, 2018.

- [113] Ann Lee, Peng-Jen Chen, Changhan Wang, Jiatao Gu, Xutai Ma, Adam Polyak, Yossi Adi, Qing He, Yun Tang, Juan Pino, et al. Direct speech-to-speech translation with discrete units. In *arXiv preprint arXiv:2107.05604*, 2021.
- [114] Chan Min Lee, Jin Woo Shin, and Nam Soo Kim. Dnn-based residual echo suppression. In *16th Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pages 1775–1779, Dresden, Germany, 2015.
- [115] Aoxue Li, Rui Peng, Chenyu Zheng, and Xiaofei Li. A supervised speech enhancement approach with residual noise control for voice communication. *Applied Sciences*, 10(8):2894, 2020.
- [116] Aoxue Li, Rui Peng, Chenyu Zheng, and Xiaofei Li. On the importance of power compression and phase estimation in monaural speech dereverberation. *JASA Express Letters*, 1(1):014802, 2021.
- [117] Guoteng Li, Chenyu Zheng, Yifan Ke, and Xiaofei Li. Deep learning-based acoustic echo cancellation for surround sound systems. *Applied Sciences*, 13(3):1266, 2023.
- [118] Y. Li, M. Tagliasacchi, B. Gfeller, and D. Roblek. Learning to denoise historical music. In *Proceedings of the 21st International Society for Music Information Retrieval Conference (ISMIR)*, pages 504–511, Montreal, Canada, 2020.
- [119] Yunpeng Li, Marco Tagliasacchi, Oleg Rybakov, Victor Ungureanu, and Dominik Roblek. Real-time speech frequency bandwidth extension. *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 691–695, 2021.
- [120] Felicia SC Lim, W Bastiaan Kleijn, Michael Chinen, and Jan Skoglund. Robust low rate speech coding based on cloned networks and wavenet. In *ICASSP 2020-*

- 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6769–6773. IEEE, 2020.
- [121] Ilya Loshchilov and Frank Hutter. SGDR: stochastic gradient descent with restarts. *CoRR*, abs/1608.03983, 2016.
- [122] Andrey Lukin and Jeremy Todd. Suppression of musical noise artifacts in audio noise reduction by adaptive 2d filtering. In *123rd Audio Engineering Society Convention*, New York, NY, USA, Jan 2007.
- [123] Yi Luo and Nima Mesgarani. Conv-tasnet: Surpassing ideal time–frequency magnitude masking for speech separation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 27(8):1256–1266, 2019.
- [124] Zhiyao Ma, Brecht De Man, Pedro Pestana, Dave Black, and Joshua D. Reiss. Intelligent multitrack dynamic range compression. *Journal of the Audio Engineering Society*, 63(6):412–426, 2015.
- [125] D. G. Malham and A. Myatt. 3-d sound spatialization using ambisonic techniques. *Computer Music Journal*, 19(4):58–70, 1995.
- [126] Yotam Mann, Google Creative Lab, and Magenta Team. A.i. experiments: Sound maker (nsynth). <https://github.com/googlecreativelab/aiexperiments-sound-maker>, 2017. Accessed: 2024-06-09.
- [127] Matthew C. McCallum, Florian Henkel, Jaehun Kim, Samuel E. Sandberg, and Matthew E. P. Davies. Similar but faster: Manipulation of tempo in music audio embeddings for tempo prediction and search, 2024.
- [128] Brian McFee, Colin Raffel, Dawen Liang, Daniel PW Ellis, Matt McVicar, Eric Battenberg, and Oriol Nieto. librosa: Audio and music signal analysis in python.

- In K. Huff and J. Bergstra, editors, *Proceedings of the 14th Python in Science Conference*, pages 18–25. Citeseer, 2015.
- [129] Gabriel Meseguer-Brocal and Geoffroy Peeters. Conditioned-u-net: Introducing a control mechanism in the u-net for multiple source separations. In *Proceedings of the 20th International Society for Music Information Retrieval Conference (ISMIR)*, Delft, The Netherlands, 2019. Licensed under Creative Commons Attribution 4.0 International License (CC BY 4.0).
- [130] Michael Miller. *Idiot’s Guides: Mixing Music*. Alpha, 2016.
- [131] Stylianos-Ioannis Mimilakis, Konstantinos Drossos, Andreas Floros, and Demetrios Katerelos. Automated tonal balance enhancement for audio mastering applications. In *Audio Engineering Society Convention 134*. Audio Engineering Society, 2013.
- [132] Stylianos Ioannis Mimilakis, Konstantinos Drossos, Tuomas Virtanen, and Gerald Schuller. Deep neural networks for dynamic range compression in mastering applications. In *Audio Engineering Society Convention 140*, Paris, France, June 2016. Audio Engineering Society. AES Convention Paper 9564.
- [133] Eduard Moliner and Vesa Välimäki. A two-stage u-net for high-fidelity denoising of historical recordings. In *ICASSP 2022 - IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 841–845. IEEE, 2022.
- [134] Eloi Moliner and Vesa Välimäki. A two-stage u-net for high-fidelity denoising of historical recordings, 2022.
- [135] Brian C. Moore. *Hearing*. Handbook of Perception and Cognition. Academic Press, San Diego, California, 2nd edition, 1995.

- [136] Gautham J. Mysore. Can we automatically transform speech recorded on common consumer devices in real-world environments into professional production quality speech? a dataset, insights, and challenges. *IEEE Signal Processing Letters*, 22:1006–1010, 2015.
- [137] Babak Naderi, Ross Cutler, and Nicolae-Catalin Ristea. Multi-dimensional speech quality assessment in crowdsourcing, 2023.
- [138] Javier Naranjo-Alcazar, Sergi Perez-Castanos, Irene Martin-Morato, Pedro Zucarello, and Maximo Cobos. On the performance of residual block design alternatives in convolutional neural networks for end-to-end audio classification, 2019.
- [139] Native Instruments GmbH. Native instruments. https://en.wikipedia.org/wiki/Native_instruments, 2025. Accessed : 2025 - 06 - 02.
- [140] Patrick Naylor and Nicolay D. Gaubitch, editors. *Speech Dereverberation*. Springer, New York, NY, USA, 2010.
- [141] Shahan Nercessian and Alexey Lukin. Speech dereverberation using recurrent neural networks. In *Proceedings of the 22nd International Conference on Digital Audio Effects (DAFx-19)*, Birmingham, UK, September 2019. Licensed under Creative Commons Attribution 3.0 Unported License.
- [142] Martin Neukom. Decoding second order ambisonics to 5.1 surround systems. In *121st AES Convention*, pages 1399–1406, San Francisco, CA, USA, 2006.
- [143] Jan Newmarch. Gstreamer. In *Linux Sound Programming*, pages 211–221. Springer, 2017.

- [144] Markus Noisternig, Thomas Musil, Andreas Sontacchi, and Robert Höldrich. 3d binaural sound reproduction using a virtual ambisonic approach. In *2003 International Symposium on Virtual Environments, Human-Computer Interfaces and Measurement Systems*, pages 174–178, Lugano, Switzerland, 2003.
- [145] Bobby Owsinski. *The Mastering Engineer’s Handbook: The Audio Mastering Handbook*. ArtistPro, 2nd edition, 2007.
- [146] Jiyun Park, Sangeon Yong, Taegyun Kwon, and Juhan Nam. A real-time lyrics alignment system using chroma and phonetic features for classical vocal performance, 2024.
- [147] Santiago Pascual, Antonio Bonafonte, and Joan Serrà. Segan: Speech enhancement generative adversarial network, 2017.
- [148] Rui Peng, Lingyun Cheng, Chenyu Zheng, and Xiaofei Li. Icassp 2021 acoustic echo cancellation challenge: Integrated adaptive echo cancellation with time alignment and deep learning-based residual echo plus noise suppression. In *2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 146–150, Toronto, ON, Canada, 2021. IEEE.
- [149] Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer, 2017.
- [150] Georg Pirker, Markus Wohlmayr, Stefan Petrik, and Franz Pernkopf. A pitch tracking corpus with evaluation on multipitch tracking scenario. In *Proceedings of Interspeech*, pages 1509–1512, 2011.
- [151] Mark A. Poletti. Three-dimensional surround sound systems based on spherical harmonics. *Journal of the Audio Engineering Society*, 53(11):1004–1024, 2005.

- [152] Archontis Politis. *Microphone Array Processing for Parametric Spatial Audio Techniques*. PhD thesis, Aalto University, Espoo, Finland, 2016.
- [153] Ryan Prenger, Rafael Valle, and Bryan Catanzaro. Waveglow: A flow-based generative network for speech synthesis. *CoRR*, abs/1811.00002, 2018.
- [154] Yao Qiu, Jinchao Zhang, Yong Shan, and Jie Zhou. Enhancing note-level singing transcription model with unlabeled and weakly labeled data. In *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 341–345, 2024.
- [155] Zafar Rafii, Antoine Liutkus, Fabian-Robert Stöter, Stylianos Ioannis Mimi-lakis, and Rachel Bittner. The MUSDB18 corpus for music separation, Decem-ber 2017.
- [156] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. *arXiv preprint arXiv:2102.12092*, 2021.
- [157] M. A. Martínez Ramírez, D. Stoller, and D. Moffat. A deep learning approach to intelligent drum mixing with the wave-u-net. *Journal of the Audio Engineering Society*, 69(3):142–151, March 2021.
- [158] Marco A. Martínez Ramírez and Joshua D. Reiss. Deep learning and intelligent audio mixing. In *Proceedings of the 3rd Workshop on Intelligent Music Produc-tion*, Salford, UK, September 2017. Intelligent Music Production. 15 September 2017.
- [159] Chandan K. Reddy, Harishchandra Dubey, Kazuhito Koishida, Anish Nair, Venkatesh Gopal, Ross Cutler, and Sriram Srinivasan. Interspeech 2021 deep noise suppression challenge. *arXiv preprint arXiv:2101.01902*, 2021.

- [160] Joshua D. Reiss. Intelligent systems for mixing multichannel audio. In *17th International Conference on Digital Signal Processing (DSP)*, pages 1–6, Corfu, Greece, 2011. IEEE.
- [161] A. W. Rix, J. G. Beerends, M. P. Hollier, and A. P. Hekstra. Perceptual evaluation of speech quality (pesq)—a new method for speech quality assessment of telephone networks and codecs. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, pages 749–752, Salt Lake City, UT, USA, 2001.
- [162] T. Roberts, A. Nicolson, and K. K. Paliwal. Deep learning-based single-ended quality prediction for time-scale modified audio. *Journal of the Audio Engineering Society*, 69(9):644–655, September 2021.
- [163] Andrew Romero. Audio segmentation with u-net architecture: Ablation study of depth variations in 1-d and 2-d audio representations. Stanford University; EE367 Computational Imaging, Class Project, 2023.
- [164] Lorenzo Romoli, Stefano Cecchi, and Francesco Piazza. A combined approach for channel decorrelation in stereo acoustic echo cancellation exploiting time-varying frequency shifting. *IEEE Signal Processing Letters*, 20(7):717–720, 2013.
- [165] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 234–241, Munich, Germany, 2015.
- [166] Tim Sainburg, Michael Thielk, and Timothy Q. Gentner. Finding, visualizing, and quantifying latent structure across diverse animal vocal repertoires. *PLoS Computational Biology*, 16(10):e1008228, 2020.

- [167] Rick Sarver and Anssi Klapuri. Application of non-negative matrix factorization to signal-adaptive audio effects. In *Proc. of the 14th Conference on Digital Audio Effects (DAFx-11)*, volume 45, pages 249–252, Paris, France, 2011.
- [168] Martin Schneider and Walter Kellermann. The generalized frequency-domain adaptive filtering algorithm as an approximation of the block recursive least-squares algorithm. *EURASIP Journal on Advances in Signal Processing*, 2016(1):1–15, 2016.
- [169] Siddharth Sigtia and Simon Dixon. Improved music feature learning with deep neural networks. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Vancouver, Canada, 2013. IEEE.
- [170] Umut Simsekli, Antoine Liutkus, and A. Taylan Cemgil. Alpha-stable matrix factorization. *IEEE Signal Processing Letters*, 22(11):2009–2013, 2015.
- [171] Paris Smaragdis. NMF? Neural Nets? It’s all the same... <http://youtube.com/watch?v=wfmPViJIjWw>, 2015. Presentation, Accessed: December 2015.
- [172] Jared Smit and Byeong Kil Lee. Automatic music mastering using deep learning. In *2022 International Conference on Computational Science and Computational Intelligence (CSCI)*, pages 385–390, 2022.
- [173] M. M. Sondhi, D. R. Morgan, and J. L. Hall. Stereophonic acoustic echo cancellation—an overview of the fundamental problem. *IEEE Signal Processing Letters*, 2(7):148–151, 1995.
- [174] Daniel Stoller, Sebastian Ewert, and Simon Dixon. Wave-u-net: A multi-scale neural network for end-to-end audio source separation. *CoRR*, abs/1806.03185, 2018.

- [175] Suno, Inc. About suno ai. <https://suno.com/about>, 2025. Accessed: 2025-06-02.
- [176] Ke Tan and DeLiang Wang. Learning complex spectral mapping with gated convolutional recurrent networks for monaural speech enhancement. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28:380–390, 2020.
- [177] John Thickstun, Zaid Harchaoui, Dean P. Foster, and Sham M. Kakade. Invariances and data augmentation for supervised music transcription. In *ICASSP 2018 - IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 2241–2245. IEEE, 2018.
- [178] John Thickstun, Zaid Harchaoui, and Sham M. Kakade. Learning features of music from scratch. In *International Conference on Learning Representations (ICLR)*, 2017.
- [179] Thomas Thiede. Peaq—the itu standard for objective measurement of perceived audio quality. *Journal of the Audio Engineering Society*, 48(1):27–42, 2000.
- [180] Matteo Torcoli, Chih-Wei Wu, Sascha Dick, Phillip A. Williams, Mhd Modar Halimeh, William Wolcott, and Emanuel A. P. Habets. Odaq: Open dataset of audio quality, 2023.
- [181] George Tzanetakis and Perry Cook. Musical genre classification of audio signals. *IEEE Transactions on Speech and Audio Processing*, 10(5):293–302, 2002.
- [182] Udio. About udio - ai-powered music creation for all. <https://www.udio.com/about-us>, 2025. Accessed: 2025-06-02.
- [183] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 5999–6009, 2017.

- [184] Haonan Wang, Peng Cao, Xiaoli Liu, Jinzhu Yang, and Osmar Zaiane. Narrowing the semantic gaps in u-net with learnable skip connections: The case of medical image segmentation, 2023.
- [185] Yuxuan Wang, Arun Narayanan, and DeLiang Wang. On training targets for supervised speech separation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 22:1849–1858, 2014.
- [186] Zhong-Qiu Wang and DeLiang Wang. Recurrent deep stacking networks for supervised speech separation. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, New Orleans, LA, USA, Mar 2017.
- [187] Karn N. Watcharasupat and Alexander Lerch. Quantifying spatial audio quality impairment. In *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, page 746–750. IEEE, April 2024.
- [188] Felix Weninger, Jonathan Le Roux, John R. Hershey, and Björn Schuller. Discriminatively trained recurrent neural networks for single-channel speech separation. In *IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pages 577–581, 2014.
- [189] Gordon Wichern and Andrey Lukin. Low-latency approximation of bidirectional recurrent networks for speech denoising. In *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, New Paltz, NY, USA, Oct 2017.
- [190] Gordon Wichern and Andrey Lukin. Removing lavalier microphone rustle with recurrent neural networks. In *Proc. Digital Audio Effects (DAFx)*, pages 19–25, Aveiro, Portugal, Sep 2018.

- [191] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. Cbam: Convolutional block attention module, 2018.
- [192] Yuxin Wu and Kaiming He. Group normalization, 2018.
- [193] B. Yegnanarayana and P. Satyanarayana. Enhancement of reverberant speech using lp residual signal. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 8:267–281, May 2000.
- [194] Hao Zhang and DeLiang Wang. Deep learning for acoustic echo cancellation in noisy and double-talk scenarios. In *19th Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pages 3239–3243, Hyderabad, India, 2018.
- [195] Hao Zhang and DeLiang Wang. Neural cascade architecture for multi-channel acoustic echo suppression. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 30:2326–2336, 2022.
- [196] Zhengxuan Zhao, Kaixu Chen, and Satoshi Yamane. CBAM-Unet++: easier to find the target with the attention module “cbam”. In *2021 IEEE 10th Global Conference on Consumer Electronics (GCCE)*, pages 655–658, Kanazawa, Japan, 2021. IEEE.
- [197] Udo Zölzer. *Digital Audio Signal Processing*. John Wiley & Sons, 2nd edition, 2008.
- [198] Franz Zotter, Hannes Pomberger, and Markus Noisternig. Energy-preserving ambisonic decoding. *Acta Acustica united with Acustica*, 98:37–47, 2012.