

MIX2MORPH: LEARNING SOUND MORPHING FROM NOISY MIXES

Annie Chu^{1,2}, Hugo Flores-García², Oriol Nieto¹, Justin Salamon¹, Bryan Pardo², Prem Seetharaman¹

¹ Adobe Research, San Francisco, USA ² Northwestern University, Evanston, USA
anniechu@u.northwestern.edu

ABSTRACT

We introduce Mix2Morph, a text-to-audio diffusion model finetuned to perform sound morphing without a dedicated dataset of morphs. By finetuning on noisy surrogate mixes at higher diffusion timesteps, Mix2Morph yields stable, perceptually coherent morphs that convincingly integrate qualities of both sources. We specifically target sound infusions, a practically and perceptually motivated subclass of morphing in which one sound acts as the dominant primary source, providing overall temporal and structural behavior, while a secondary sound is infused throughout, enriching its timbral and textural qualities. Objective evaluations and listening tests show that Mix2Morph outperforms prior baselines and produces high-quality sound infusions across diverse categories, representing a step toward more controllable and concept-driven tools for sound design.

Index Terms— sound design, text-to-audio, morphing

1. INTRODUCTION

Sound design encompasses the creative process of shaping and combining audio to achieve expressive narrative goals. A common approach is *sound morphing*: the fusion of two sounds into a single, coherent sound resembling both at once [1]. For instance, to design a menacing alien voice, a sound designer may combine animal vocalizations with human speech, producing a unified sound that is at once recognizable and alien.

A central challenge in morphing is the generation of perceptually coherent midpoints: hybrids that convincingly sound like both sources at once. Human perception of sound is shaped primarily by two dimensions; temporal structure and spectral content, which provide natural anchors for coherence [2, 3, 4, 5]. When these anchors are not preserved, intermediate morphs become perceptually fragile: they may lose the identity of one or both sources, sound more like an additive mix than a fused sound, or drift into unnatural artifacts [6, 7]. This issue is important in practice, as sound designers often aim not to fully transform one sound into another but to preserve the recognizable identity of a primary source while enriching it with qualities of a secondary source. To capture this practically motivated and perceptually distinctive case, we introduce the term *sound infusion*. We define a sound infusion as a special case of asymmetric (or unbalanced [6]) static morphing [8, 9], in which one sound acts as the primary source (e.g., a human voice) while another serves as the secondary (e.g., a lion’s roar), contributing supplemental features throughout (e.g., a rough, roar-like texture woven into its timbre).

In this work, we focus on producing robust, high-quality, perceptually-coherent sound infusions where the temporal dynamics are rooted in the primary source while timbral qualities are jointly blended from both, as in [10]. While sound infusion offers a practically motivated framing of the midpoint problem, prior work has focused on other morphing forms (e.g., dynamic

or repetitive [11, 8, 12]). Classic DSP methods, such as feature interpolation guided by perceptual descriptors [1, 11, 13] and cross-synthesis [14, 15] generally produce good morphs when applied to pitched, harmonic sounds (e.g., instruments, vocals). However, they struggle with unpitched textures like ambient noise, environmental recordings, and everyday sound effects (SFX)—categories central to creative sound design.

To broaden morphing beyond these limits, recent work has turned to deep learning [16]. Systems like MorphFader [12] and SoundMorpher [7] extend the expressive scope of morphing, leveraging the pre-existing abilities of AudioLDM2 [17], a text-to-audio (TTA) model. However, they frequently suffer from midpoint collapse: endpoint generations remain stable, but intermediate morphs are perceptually fragile and drift toward incoherent or additive-sounding mixes.

A complementary line of work decomposes sounds into their dominant perceptual dimensions [2, 3, 4, 5], temporal structure and spectral content, and blends them independently. Dixit et al. [6] targeted temporal envelopes by interpolating amplitude contours via an autoencoder, while Tokui and Baker [10], building on granular resynthesis [18, 19], used neural audio codecs to preserve a primary source’s temporal structure while infusing spectral content from a secondary source. This latter approach is especially relevant to sound infusions, but it remains fragile across diverse sound categories and often introduces audible, noise-like artifacts.

A key bottleneck in training a model for morphing is the lack of training data exemplifying morphs, especially ones that both remain perceptually coherent across diverse sound categories and preserve the asymmetric dominance structure valued in creative sound design. In this work, we present Mix2Morph, a text-to-audio model designed for generating high-quality sound infusions. Mix2Morph’s strength lies in a finetuning strategy that enables morph generation without requiring a dedicated morphing dataset. Inspired by the idea of “making good data from bad data” [20], we construct surrogate training data by repurposing noisy additive mixes, allowing the model to learn how to produce perceptually coherent morphs. Our contributions include:

- Mix2Morph, a text-to-audio model for robust midpoint sound infusions, enabled by a finetuning strategy that repurposes pretrained generative models to perform morphing
- An automated method for constructing a surrogate training dataset for morphing, enabling model training without requiring a pre-existing morph corpus
- Extended objective metrics for evaluating morphing quality, including a *directionality* measure and a proxy for morph identification via latent compressibility
- Objective and subjective validation that Mix2Morph outperforms both the prior base model and state-of-the-art deep learning baselines

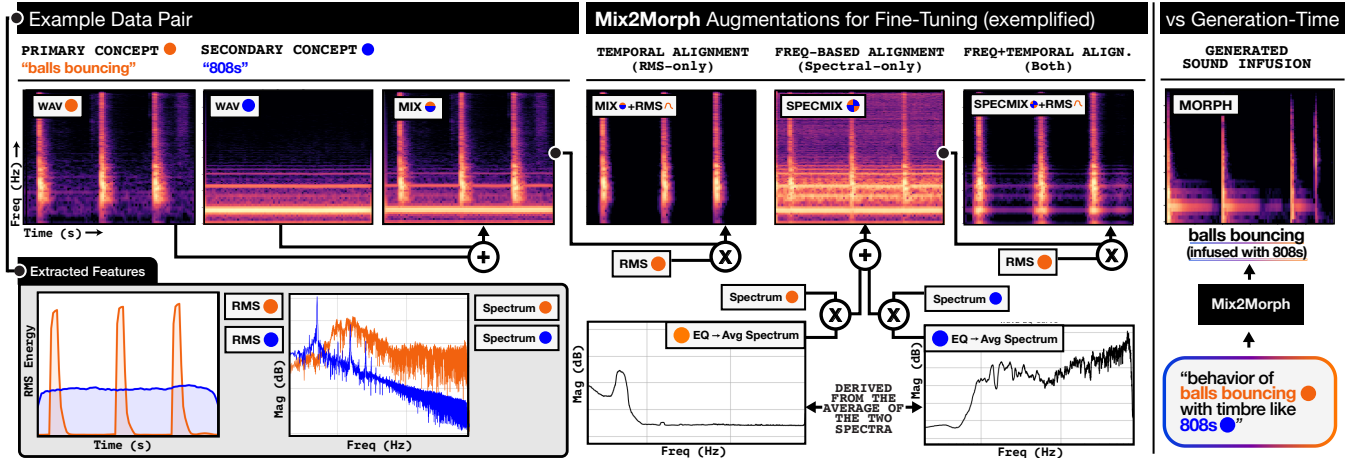


Fig. 1: Mix2Morph pipeline for creating noisy surrogate morphs via augmented additive mixing of sound pairs, used as high-timestep training targets in our diffusion-based TTA model. This example (as well as others) can be listened to at <https://anniechu.github.io/mix2morph/>

2. PROPOSED METHODOLOGY

Training sound morphing models is hindered by the lack of high-quality morph datasets. We address this by constructing noisy surrogates, specially designed *mixes* of two sounds, as training signals.

We adapt a “no-waste” dataset training strategy [20], where low-quality data are not discarded but instead assigned to higher diffusion timesteps. In diffusion training, the denoising objective depends on the noise level (timestep t): at low t , the model emphasizes recovering fine-grained details, while at high t , it focuses on coarse, global structure. We hypothesize that assigning noisy surrogate morphs to higher timesteps encourages the model to capture high-level morphing concepts while suppressing low-level mixing artifacts, thereby leveraging imperfect data without overfitting to artifacts.

In our context, the low-quality data are *simple additive mixes* of training pairs, which we treat as “bad” morphs. To strengthen their role as surrogate morphs, we apply temporal and spectral augmentations that encourage overlap and alignment, effectively bringing both sources into a shared structural and/or spectral space (Fig. 1). These noisy surrogates are used to train a model only as high-timestep training targets, never as low-timestep targets. At inference, we rely on refinement learned in pretraining to construct fine detail in low- t timesteps. Crucially, this lets us construct a scalable training set without requiring high-quality morph examples, making this the first work to leverage noisy surrogate morphs for training at scale.

2.1. Augmentation Techniques

Temporal Alignment (RMS Anchoring): We begin with an additive mix of a primary and secondary waveform, combined at 0 dB SNR (signals normalized to equal power). To ensure temporal overlap, the secondary sound is truncated or looped to match the primary’s duration. We then extract the RMS envelope of the primary and apply it to the mix, anchoring the composite to follow the macroscopic temporal “behavior” of the primary source. This produces a more realistic surrogate morph: rather than two sounds overlapping arbitrarily, the mix exhibits the temporal structure of the primary with added timbral detail from the secondary.

Frequency-Domain Alignment (Spectral Interpolation): Inspired by classic DSP interpolation, we employ a simple frequency-domain alignment trick to align the spectral content of two sounds. Given

length-matched waveforms, we compute their FFT magnitudes and construct an averaged target spectrum: $|Y_{\text{target}}| = 0.5 \cdot |Y_1| + 0.5 \cdot |Y_2|$. This serves as a shared spectral envelope. For each sound, we derive a frequency-dependent gain mask (EQ curve) as the ratio between the target and original spectrum, specifying how much to boost or attenuate each band. These EQ curves are smoothed with a moving window, applied to the signals, and summed to yield the spectral composite. Perceptually, this produces a surrogate morph in which both sounds are projected into a common timbral space.

2.2. Augmentation Modes

Diverse Training Targets: To diversify training, we define four augmentation modes (Fig. 1): (i) RMS-only, (ii) Spectral-only, (iii) Both (Spectral+RMS), or (iv) None (unaugmented mix). During training, one mode is randomly assigned to each audio pair.

Caption Conditioning: Each augmentation mode is paired with a corresponding caption: (i) RMS-only, where captions indicate primary behavior with textures from both (“*The behavior of X with textures from X and Y*”), (ii) Spectral-only, highlighting timbral blending (“*A spectral blend of X and Y*”), (iii) Both, describing primary behavior with blended timbres (“*The behavior of X with a spectral blend of X and Y*”), and (iv) None (“*A mix of X and Y*”).

3. MORPHING EVALUATION METRICS

A generative model can produce a morph, a mix, or collapse into a single-concept generation. To our knowledge, there is no standardized metric for distinguishing morphs from mixes or single-concept generations. To explore potential proxies, we tested the *compressibility* of various low-level features (MFCCs, DAC [21] latents, or log-melspectrograms) via dimensionality reduction, hypothesizing that true morphs, and especially single-concept generations, form more coherent, lower-dimensional representations, while mixes retain more independent variability. To validate candidate proxies, we conducted a human study on 250 clips spanning mixes, single-concept generations, and morphs, annotated by a professional sound designer and the first author. We then compared features against these annotations using Spearman correlations and ROC-AUC scores. The most reliable indicator was the cumulative variance explained by the first two principal components of the DAC latent

space, which we term the **Latent Compressibility Score (LCS)**. LCS correlated strongly with humans labeling audio as a morph (Spearman’s $\rho=0.9$, $p=0.037$), suggesting it as a practical proxy for a system’s tendency to produce morphs.

While LCS measures whether an output is perceived as a morph at all, the following metrics assess the semantic composition of that morph. To measure perceptual-relevance, we follow Caetano & Osaka’s [9] morphing criteria of *correspondence* and *intermediateness*. We adopt an embedding-based approach, using FLAM [22], a shared text–audio model, to compute clip-wise cosine similarities between the morphed audio and the text embeddings of primary and secondary source concepts, sim_X and sim_Y respectively. **Correspondence** is defined as the joint semantic presence of both source concepts in the morph. As this requires both similarities to be high, we operationalize this as the harmonic mean of sim_X and sim_Y (analogous to F1 scoring): $\frac{2 \cdot sim_X \cdot sim_Y}{sim_X + sim_Y}$.

Intermediateness captures balance between concepts; higher when the morph is equally similar to both concepts and lower when one dominates. We compute this as the normalized similarity difference, $1 - \frac{|sim_X - sim_Y|}{\max(sim_X, sim_Y)}$.

Due to the asymmetric nature of sound infusions, we extend evaluation with a softmax-based **directionality** score to assess a morph’s alignment with the intended prompt direction (“*behavior like X, timbre like Y*”) versus its reverse (“*behavior like Y, timbre like X*”). We compute cosine similarities between the audio embedding and the intended (s_{int}) and reversed (s_{rev}) prompt embeddings and apply a softmax ($T = 0.05$) rather than raw difference to emphasize relative directional preference, resolving small but consistent directional biases detectable under semantic overlap. The resulting probability p is mapped to $[-1, 1]$ via $2p - 1$, where +1 indicates bias towards intended prompt and -1 towards the reverse. We also report FAD [23] as an audio quality metric with reference embeddings computed from a proprietary dataset of 40k high-quality SFX with a sound-class distribution matched to the training data.

4. EXPERIMENTS & RESULTS

4.1. Curating Concept Pairs for Evaluation

We curated a set of 50 concept pairs of distinct sound classes designed to test morphing performance, informed by sound-design practices [24, 25, 26] and taxonomies¹. The set aims for categorical coverage across source types (e.g., impact/temporally localized events, continuous textures) and infusion types, with an emphasis on *inter-class* blends that are more difficult in practice to achieve (e.g., blending a TNT explosion with a choir swell) rather than *intra-class* blends (e.g., a soft whoosh with a sparky, metallic whoosh). Each pair is used in both directions (“behavior of X , timbre like Y ” & vice versa), yielding 100 sound infusion prompts total.

4.2. Model Implementation Details

Our base model is a large text-to-audio (TTA) latent diffusion transformer pretrained to perform single sound generation, similar to [27, 28]. The architecture consists of a VAE that compresses 48 kHz stereo audio into 256-dimensional latent sequences at 40 Hz. Generated latent sequences are decoded back to waveform via the VAE decoder. The training data for both pretraining and fine-tuning is a large set of proprietary, licensed SFX datasets and publicly available CC-licensed general audio corpora. After pretraining, we fine-

¹<https://universalcategorysystem.com/>

| Model | t start | t end | RMS | Spectral | Both | None | LCS \uparrow | Correspond. \uparrow | Intermediat. \uparrow | Direct. \uparrow | FAD \downarrow |
|----------------------|-----------|---------|-----|----------|------|--------------|----------------|------------------------|-------------------------|--------------------|------------------|
| base | - | - | - | - | - | - | 0.136 | 0.678 | 0.611 | 0.525 | 1.219 |
| +Timestep Allocation | 0 | 1 | ✓ | × | × | × | 0.128 | 0.699 | 0.646 | 0.173 | 1.230 |
| | 0.25 | 1 | ✓ | × | × | × | <u>0.143</u> | 0.705 | <u>0.658</u> | 0.278 | 1.235 |
| | 0.5 | 1 | ✓ | × | × | × | <u>0.141</u> | <u>0.721</u> | 0.672 | 0.296 | <u>1.221</u> |
| +Augmentation Mode | 0.75 | 1 | ✓ | × | × | × | 0.134 | 0.717 | 0.653 | <u>0.364</u> | 1.225 |
| | 0.5 | 1 | ✓ | × | ✓ | × | 0.135 | 0.700 | 0.623 | 0.363 | 1.226 |
| | 0.5 | 1 | ✓ | ✓ | ✓ | × | 0.150 | 0.725 | 0.648 | 0.436 | 1.220 |
| 0.5 | 1 | ✓ | ✓ | ✓ | ✓ | <u>0.143</u> | 0.712 | 0.650 | 0.349 | 1.222 | |
| Simple Mixing | - | - | - | - | - | - | 0.132 | 0.758 | 0.690 | -9.25e-13 | 1.293 |
| LGrS | - | - | - | - | - | - | 0.173 | 0.539 | 0.638 | -0.119 | <u>1.290</u> |
| MorphFader | - | - | - | - | - | - | 0.085 | 0.418 | 0.421 | -9.72e-13 | 1.430 |
| SoundMorpher | - | - | - | - | - | - | 0.242 | 0.591 | 0.641 | -9.64e-13 | 1.380 |
| Mix2Morph | 0.5 | 1 | ✓ | ✓ | ✓ | × | <u>0.150</u> | <u>0.725</u> | <u>0.648</u> | 0.436 | 1.220 |

Table 1: Ablation results (above) and baseline comparisons (below). Best scores bolded, second-best underlined (per section).

tune for 50k steps using the generated surrogate morph dataset (§2). Training uses 8s audio segments mixed via augmentation techniques; we generate 3s sounds for evaluation.

4.3. Effect of Training at Higher Timesteps

The base model, when given the two-concept infusion prompts, often collapses to single-concept outputs. The metrics reflect this imbalance: high directionality (0.525) and compressibility (LCS = 0.136), but weaker correspondence (0.678), reflecting limited joint representation of both sources. To test our hypothesis that noisy surrogate morphs are best utilized at higher timesteps, we restrict the mix dataset to specific diffusion timestep ranges during training. Outside the allocated window, training proceeds as in the base model, i.e., the single-source reconstruction objective. We fine-tune models with allocations of [0, 1], [0.25, 1], [0.5, 1], [0.75, 1], using RMS-only augmentation for simplicity, and compare against the unmodified base model (Table 1). Shifting augmented mixes toward higher timesteps improves performance relative to base, **peaking with $t \in [0.5, 1]$, which yields the best balance and all-around performance:** highest correspondence (0.721) and intermediateness (0.672) while maintaining strong directionality (0.296). Further constricting to [0.75, 1] increases directionality (0.364) but reduces intermediateness, indicating a trade-off between preserving semantic balance and strengthening prompt asymmetry. Qualitatively, this aligns with perceptual impressions: [0.5, 1] best incorporates both sources while retaining the primary sound’s quality, indicating blending, whereas [0.75, 1] loses shared concepts and reverts back toward base-model behavior. We therefore fix the allocation range to [0.5, 1] in subsequent experiments.

4.4. Effect of Augmentation Modes

Fixing the timestep range to $t \in [0.5, 1]$, we compare the effect of applying different augmentation modes to the mix at varying probabilities. In addition to the RMS-only model (probability= 1.0), we evaluate (i) a balanced configuration with RMS-only (0.5) and Both (0.5), and (ii) a diversified three-way configuration with RMS-only (0.33), Spectral-only (0.33), and Both (0.33). Results in Table 1 show that RMS-only improves upon the baseline, but both richer

augmentation configurations outperform RMS-only, with the **three-way configuration delivering the strongest gains across nearly all metrics** (with slightly lower but still competitive intermediateness, reflecting a trade-off with correspondence and directionality). It also yields the lowest FAD score, indicating high-quality generations. Including “None” as a fourth branch reduces performance across metrics. Qualitatively, the three-way configuration produces the most natural and well-blended infusions; the two-way and four-way branches both yielding more mix-like outputs. We thus adopt this three-way augmentation model, trained at $t \in [0.5, 1]$, as our final training setup for downstream evaluation, i.e. **Mix2Morph**.

4.5. Baselines for Comparison

We compare Mix2Morph against five baselines: (i) base model, (ii) simple waveform mixing, (iii) latent granular resynthesis (LGrS) [10], (iv) MorphFader [12], and (v) SoundMorpher [7]. For simple mixing, the primary and secondary sources are generated independently with the base model (without Mix2Morph finetuning) and then mixed. LGrS is included as it explicitly addresses sound infusion; we re-implement it using a 48 kHz VAE (the same as our base diffusion model) instead of the original 44.1 kHz DAC and provide it with the same input pairs as simple mixing. MorphFader requires paired text prompts (e.g., “alien voice” and “lion roaring”), which we adapt for infusion-style comparisons, while SoundMorpher functions as an audio-to-audio method with only auxiliary text conditioning, so we provide it with the same input pairs as LGrS and mixing. We note that Mix2Morph, LGrS, and simple mixing operate at 48 kHz, whereas MorphFader and SoundMorpher are restricted to 16 kHz as an artifact of AudioLDM2.

4.6. Quantitative Results

Shown in Table 1, Mix2Morph consistently outperforms all baselines. It achieves higher correspondence and intermediateness than LGrS, MorphFader, and SoundMorpher while maintaining positive directionality, whereas the latter two collapse to near-zero or negative scores (though expected since they were not designed for infusion). Simple mixing attains high correspondence (0.758) but near-zero directionality, confirming it produces overlays rather than coherent morphs. Mix2Morph also yields higher LCS than the base model and mixing baselines (0.141 vs 0.127–0.136), indicating stronger morph tendencies, though slightly lower than LGrS and SoundMorpher, which is expected given their more aggressive latent operations. However, LGrS and SoundMorpher both tend to show lower correspondence, i.e., effectively producing morphs, but ones that are less perceptually coherent, whereas Mix2Morph achieves a better balance between morph strength and perceptual fidelity. Relative to the unmodified base model, Mix2Morph improves correspondence (0.721 vs 0.678) and intermediateness (0.672 vs 0.611), reduces directionality (0.296 vs 0.525), and raises LCS (0.141 vs 0.136), effectively reorienting the base model from its single-source bias toward more balanced blends that still preserve the primary sound’s identity. **Overall, Mix2Morph delivers the best balance of semantic coherence and perceptual relevance without sacrificing fidelity, enabling more coherent and perceptually balanced infusions than both the base model and prior baselines.**

4.7. Subjective Evaluation

We conducted a listening test with $N = 25$ participants (inc. audio producers, musicians, researchers, content creators, and enthusiasts). From the 100 infusion prompts in §4.1, we randomly select

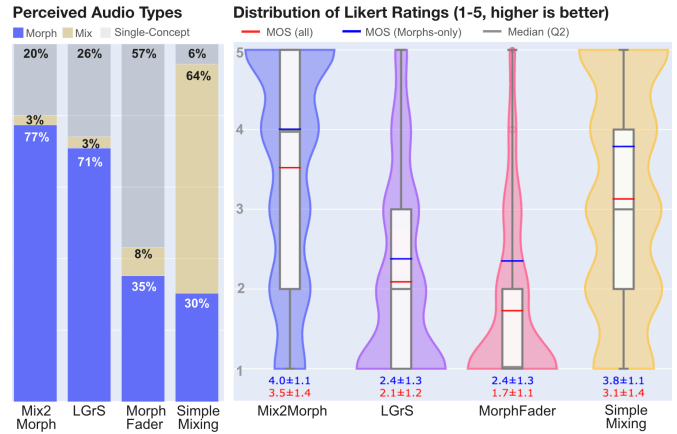


Fig. 2: Listener Study Results: Mix2Morph achieves the highest morph rate, with MOS skewed toward higher values.

20. To mitigate rater fatigue, we limit evaluation to three external baselines: simple mixing, LGrS [10], and MorphFader [12] (excluding SoundMorpher due to redundancy with AudioLDM2 and high computational cost).

Given a generated sound infusion and its target infusion prompt (e.g., “behavior of a dog barking with timbre like a car horn”), we asked participants to score each generated output on this question: *How successfully does the audio capture the morphing prompt above?* using a 1-5 Likert scale, where 5 is best. Participants were also asked to label each generated output as one of the following: *a Morph, a Mix, or Single-Concept*. Each participant evaluated 80 morphs (4 models \times 20 prompts), randomized per subject.

We report audio type labels and Mean Opinion Score (MOS) in Fig. 2. A repeated-measures ANOVA showed a significant main effect of model on MOS ($F(3, 72) = 76.4, p < 0.001$). Tukey tests confirmed that Mix2Morph scored significantly higher than LGrS and MorphFader (both $p < 0.001$). The morph rate is the percentage of generated outputs labeled *Morph*. **Mix2Morph achieved the highest morph rate (77%) and MOS (3.52 overall; 4.00 morph-MOS)**, consistent with its strong objective metric scores (§4.6). LGrS produced a similar morph rate (71%) but much lower quality (MOS = 2.09), reflecting weaker perceptual coherence despite decent correspondence in objective scoring. Simple mixing was often judged as mixes (64%) but still reached moderate MOS (3.13 overall; 3.79 morph-MOS), suggesting it can suffice for certain infusion types (e.g., impact sounds with aligned onsets). Consistent with objective scores, MorphFader reached a morph rate of 35%, producing more single-concept and mix-like outputs, and the lowest MOS (1.73), likely due to its 16kHz backbone. Overall, Mix2Morph outperformed all baselines in MOS and morph frequency.

5. CONCLUSION

We introduce Mix2Morph, a TTA model that leverages surrogate mixes to adapt pretrained audio diffusion models for morphing without a dedicated morph dataset. By leveraging augmented mixes as noisy training data at higher diffusion timesteps, Mix2Morph achieves consistent improvements in morphing quality across both objective and subjective evaluations, outperforming existing baselines. This represents a substantial step toward more controllable and concept-driven tools for sound design. Future work may explore more intuitive control mechanisms (e.g., audio-to-audio).

6. REFERENCES

- [1] Marcelo Caetano, *Morphing isolated quasi-harmonic acoustic musical instrument sounds guided by perceptually motivated features*, Ph.D. thesis, Paris 6, 2011.
- [2] Laurie M Heller and Lauren Wolf, “When hybrid sound effects are better than real recordings,” in *Proceedings of Meetings on Acoustics*. Acoustical Society of America, 2022, vol. 46, p. 050002.
- [3] Andrew J Oxenham, “How we hear: The perception and neural coding of sound,” *Annual review of psychology*, vol. 69, no. 1, pp. 27–50, 2018.
- [4] Zachary M Smith, Bertrand Delgutte, and Andrew J Oxenham, “Chimaeric sounds reveal dichotomies in auditory perception,” *Nature*, vol. 416, no. 6876, pp. 87–90, 2002.
- [5] P. Iverson and C. Krumhansl, “Isolating the dynamic attributes of musical timbre,” *The Journal of the Acoustical Society of America*, vol. 94 5, pp. 2595–603, 1993.
- [6] Satvik Dixit, Sungjoon Park, Chris Donahue, and Laurie M Heller, “Learning perceptually relevant temporal envelope morphing,” in *Proceedings of the IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 2025.
- [7] Xinlei Niu, Jing Zhang, and Charles Patrick Martin, “Sound-morpher: Perceptually-uniform sound morphing with diffusion model,” *arXiv preprint arXiv:2410.02144*, 2024.
- [8] William A Sethares and James A Bucklew, “Kernel techniques for generalized audio crossfades,” *Cogent Mathematics*, vol. 2, no. 1, pp. 1102116, 2015.
- [9] Marcelo Caetano and Naotoshi Osaka, “A formal evaluation framework for sound morphing,” in *ICMC*, 2012.
- [10] Nao Tokui and Tom Baker, “Latent granular resynthesis using neural audio codecs,” *arXiv preprint arXiv:2507.19202*, 2025.
- [11] Savvas Kazazis, Philippe Depalle, and Stephen McAdams, “Sound morphing by audio descriptors and parameter interpolation,” in *Proceedings of the 19th International Conference on Digital Audio Effects (DAFx-16)*. Brno, Czech Republic, 2016.
- [12] Purnima Kamath, Chitralekha Gupta, and Suranga Nanayakkara, “Morphfader: Enabling fine-grained controllable morphing with text-to-audio models,” in *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2025, pp. 1–5.
- [13] Marcelo Caetano, “Morphing musical instrument sounds with the sinusoidal model in the sound morphing toolbox,” in *International Symposium on Computer Music Multidisciplinary Research*. Springer, 2019, pp. 481–503.
- [14] Chris Donahue, Tom Erbe, and Miller S Puckette, “Extended convolution techniques for cross-synthesis,” in *ICMC*, 2016.
- [15] Aaron Valero Puche and Sukhan Lee, “Caesynt: Real-time timbre interpolation and pitch control with conditional autoencoders,” *2021 IEEE 31st International Workshop on Machine Learning for Signal Processing (MLSP)*, pp. 1–6, 2021.
- [16] Chitralekha Gupta, Purnima Kamath, Yize Wei, Zhuoyao Li, Suranga Nanayakkara, and Lonce Wyse, “Towards controllable audio texture morphing,” in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [17] Haohe Liu, Yi Yuan, Xubo Liu, Xinhao Mei, Qiuqiang Kong, Qiao Tian, Yuping Wang, Wenwu Wang, Yuxuan Wang, and Mark D Plumbley, “Audioldm 2: Learning holistic audio generation with self-supervised pretraining,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 32, pp. 2871–2883, 2024.
- [18] Adrien Bitton, Philippe Esling, and Tatsuya Harada, “Neural granular sound synthesis,” in *Proceedings of the International Computer Music Conference (ICMC)*, 2020.
- [19] Christopher J. Tralie and Ben Cantil, “The concatenator: A bayesian approach to real time concatenative musaicing,” in *Proceedings of the 25th Conference of the International Society for Music Information Retrieval (ISMIR)*, 2024.
- [20] Giannis Daras, Adrian Rodriguez-Munoz, Adam Klivans, Antonio Torralba, and Constantinos Daskalakis, “Ambient diffusion omni: Training good models with bad data,” *arXiv preprint arXiv:2506.10038*, 2025.
- [21] Rithesh Kumar, Prem Seetharaman, Alejandro Luebs, Ishaan Kumar, and Kundan Kumar, “High-fidelity audio compression with improved rvqgan,” *Advances in Neural Information Processing Systems*, vol. 36, pp. 27980–27993, 2023.
- [22] Yusong Wu, Christos Tsirigotis, Ke Chen, Cheng-Zhi Anna Huang, Aaron Courville, Oriol Nieto, Prem Seetharaman, and Justin Salamon, “FLAM: Frame-wise language-audio modeling,” in *Forty-second International Conference on Machine Learning*, 2025.
- [23] Kevin Kilgour, Mauricio Zuluaga, Dominik Roblek, and Matthew Sharifi, “Fréchet audio distance: A reference-free metric for evaluating music enhancement algorithms,” in *Interspeech 2019*, 2019, pp. 2350–2354.
- [24] Neil Hillman and Sandra Pauletto, “The craftsman: The use of sound design to elicit emotions,” *The Soundtrack*, vol. 7, no. 1, pp. 5–23, 2014.
- [25] Trevor Pinch and Karin Bijsterveld, *The Oxford handbook of sound studies*, OUP USA, 2012.
- [26] Michel Chion, “The three listening modes,” *The sound studies reader*, vol. 1, 2012.
- [27] Zach Evans, Julian Parker, CJ Carr, Zack Zukowski, Josiah Taylor, and Jordi Pons, “Long-form music generation with latent diffusion,” in *International Society for Music Information Retrieval Conference*, 2024.
- [28] Hugo Flores García, Oriol Nieto, Justin Salamon, Bryan Pardo, and Prem Seetharaman, “Sketch2sound: Controllable audio generation via time-varying signals and sonic imitations,” in *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2025, pp. 1–5.