

# Aligning Generative Music AI with Human Preferences: Methods and Challenges

**Dorien Herremans, Abhinaba Roy**

AMAAI Lab, Singapore University of Technology and Design  
8 Somapah Road  
487372 Singapore  
dorien\_herremans@sutd.edu.sg, abhinaba\_roy@sutd.edu.sg

## Abstract

Recent advances in generative AI for music have achieved remarkable fidelity and stylistic diversity, yet these systems often fail to align with nuanced human preferences due to the specific loss functions they use. This paper advocates for the systematic application of preference alignment techniques to music generation, addressing the fundamental gap between computational optimization and human musical appreciation. Drawing on recent breakthroughs including MusicRL’s large-scale preference learning, multi-preference alignment frameworks like diffusion-based preference optimization in DiffRhythm+, and inference-time optimization techniques like Text2midi-InferAlign, we discuss how these techniques can address music’s unique challenges: temporal coherence, harmonic consistency, and subjective quality assessment. We identify key research challenges including scalability to long-form compositions, reliability amongst others in preference modelling. Looking forward, we envision preference-aligned music generation enabling transformative applications in interactive composition tools and personalized music services. This work calls for sustained interdisciplinary research combining advances in machine learning, music-theory to create music AI systems that truly serve human creative and experiential needs.

## Introduction

The emergence of large-scale generative models for music has transformed the landscape of artificial creativity, achieving unprecedented fidelity in audio synthesis and remarkable diversity in stylistic expression. Contemporary systems such as MusicLM (Agostinelli et al. 2023), MusicGen (Copet et al. 2023), Mustango (Melechovsky et al. 2023) and Jukebox (Dhariwal et al. 2020) demonstrate capabilities that would have been unbelievable just years ago, generating high-quality musical compositions from textual descriptions with impressive technical proficiency. Yet despite these remarkable achievements, these systems face a fundamental and persistent challenge: aligning their outputs with the nuanced, subjective nature of human musical preferences.

Traditional supervised training approaches in music generation rely predominantly on likelihood-based objectives that optimize for training distributions. While these objectives successfully capture surface-level musical patterns and

enable coherent generation, they fundamentally fail to capture the deeper qualities that make music aesthetically pleasing, emotionally resonant to human listeners (Juslin and Västfjäll 2008). This alignment gap becomes particularly pronounced in music generation, where subjective appreciation depends on complex, multi-layered factors including cultural context, emotional resonance, personal taste, and the intricate interplay between musical elements across multiple temporal scales (Agres, Forth, and Wiggins 2016).

The challenges inherent to musical preference alignment extend far beyond those encountered in other generative domains. Unlike text generation where semantic correctness and factual accuracy provide objective evaluation criteria (Blei, Ng, and Jordan 2003), or image generation where visual fidelity offers measurable benchmarks (Salimans et al. 2016; Heusel et al. 2017), musical quality encompasses a complex constellation of factors that resist straightforward quantification. These include temporal coherence across extended sequences that may span minutes or hours, harmonic consistency that must maintain theoretical soundness while supporting creative expression, rhythmic stability that provides structural foundation while allowing for expressive variation, timbral quality that affects emotional perception (McAdams 2013), and subjective aesthetic judgment that varies dramatically across individuals and cultural contexts (North, Hargreaves, and Krause 2016).

The concept of quality of music itself defies universal definition, being inherently user-dependent and context-sensitive. A caption such as “upbeat workout music” might legitimately map to diverse interpretations ranging from retro guitar solos to techno pop beats, from orchestral arrangements to electronic dance music, each carrying different emotional connotations and serving different functional purposes (Cideron et al. 2024). This inherent ambiguity in musical interpretation challenges traditional supervised learning paradigms and calls for specialized approaches that can integrate continuous human feedback into post-deployment model refinement.

Furthermore, music exists within complex social and cultural ecosystems that shape preference formation in ways that purely technical metrics cannot capture. Musical preferences are influenced by factors including cultural background, age demographics, social identity, historical context, and personal experiences that create deep emotional as-

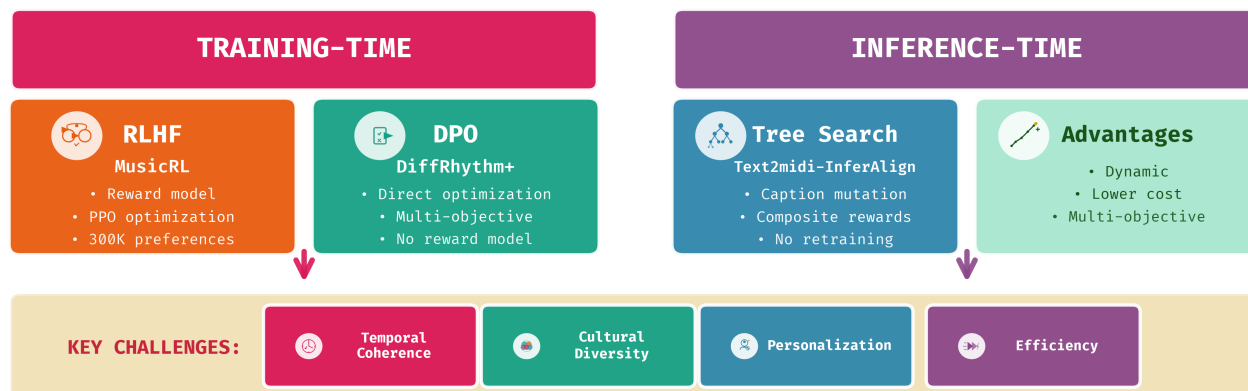


Figure 1: Overview of selected preference alignment approaches for music generation. Training-time methods (left) optimize models during development, while inference-time methods (right) adapt generation without retraining. Both address the fundamental gap between likelihood-based training and human musical preferences.

sociations with particular musical styles or elements (North, Hargreaves, and Krause 2016). These preferences evolve dynamically over time, both at individual and collective levels, requiring adaptive systems that can respond to changing aesthetic landscapes (Schedl et al. 2014).

Recent breakthroughs have demonstrated promising directions for addressing these multifaceted challenges through sophisticated preference alignment techniques. MusicRL (Cideron et al. 2024) is the first music generation system fine-tuned from human feedback at scale, collecting an unprecedented dataset of approximately 300,000 pairwise preferences to train models that significantly outperform baselines in comprehensive human evaluations. However, this preference dataset remains proprietary and is not publicly available, limiting reproducibility and broader adoption. This work demonstrates that large-scale preference learning can bridge the gap between computational optimization and human musical appreciation, while revealing the complex, multi-dimensional nature of musical preference that extends beyond simple quality metrics.

DiffRhythm+ (Chen et al. 2025) extends preference optimization to diffusion-based architectures, enabling multi-preference alignment for controllable full-length song generation that can simultaneously satisfy multiple preference criteria. This work demonstrates how modern generative architectures can be adapted to incorporate preference signals at multiple levels of musical structure, from local harmonic progressions to global compositional form.

Text2midi-InferAlign (Roy, Puri, and Herremans 2025) introduces inference-time optimization techniques that achieve substantial improvements in text-audio consistency without requiring expensive model retraining, opening new possibilities for dynamic preference adaptation. Through sophisticated tree search algorithm, this approach demonstrates that preference alignment can be achieved through clever inference-time computation, enabling real-time adaptation to user preferences and contextual requirements.

The field builds upon earlier foundational work in reinforcement learning for music generation. Jaques12 et al. (2016) used RL-based music generation by fine-tuning RNN-based models using rewards incorporating music theory rules, while Le Groux and Verschure (2010) developed adaptive interactive systems using RL agents for real-time emotional feedback optimization. Recent systems have continued this trajectory: JAM (Liu et al. 2025) employs DPO, NotaGen (Wang et al. 2025) introduces CLaMP-DPO for classical sheet music generation, and inference-time methods like DITTO (Novack et al. 2024) and SMITIN (Koo et al. 2025) demonstrate sophisticated optimization strategies during generation.

These advances collectively highlight three core technical approaches to preference alignment in music generation as shown in Figure 1: (1) large-scale preference learning through reinforcement learning from human feedback (RLHF) that can capture complex preference patterns from extensive human evaluation data, (2) inference-time optimization that balances multiple objectives during generation without requiring model modification, and (3) direct preference optimization integrated into modern generative architectures that can learn preference-aligned representations during training. Each approach offers distinct advantages and challenges, creating a rich landscape of technical possibilities for future development.

We provide a comprehensive examination of how these methods address music’s fundamental challenges while identifying key research directions for creating music AI systems that truly serve human creative needs. We argue that the future of music generation lies not in technical sophistication or computational scale, but in meaningful alignment with the rich complexity of human musical preference, cultural expression, and creative intention.

Preference alignment in generative models has emerged as a critical paradigm for creating systems that produce outputs aligned with human judgment rather than merely maxi-

mizing likelihood with respect to training distributions. This shift represents a fundamental reconceptualization of the generative modeling objective, moving from statistical fidelity to human-centered quality optimization (Christiano et al. 2017; Bai et al. 2022).

## Reinforcement Learning from Human Feedback

Reinforcement Learning from Human Feedback (RLHF) (Christiano et al. 2017) pioneered the preference alignment paradigm by decomposing the alignment problem into two distinct phases. In the first phase, a **reward model**  $r_\phi(x, y)$  is trained to predict human preferences using a dataset of preference comparisons  $\mathcal{D} = \{(x_i, y_i^w, y_i^l)\}_{i=1}^N$ , where  $y_i^w$  represents the preferred output and  $y_i^l$  the less preferred output for input  $x_i$ . The reward model is typically trained using the Bradley-Terry preference model (Bradley and Terry 1952):

$$p(y^w \succ y^l | x) = \sigma(r_\phi(x, y^w) - r_\phi(x, y^l))$$

where  $\sigma$  denotes the sigmoid function and  $\succ$  indicates preference ordering.

In the second phase, the generative policy  $\pi_\theta$  is optimized to maximize expected rewards while maintaining similarity to a reference policy  $\pi_{\text{ref}}$  through a regularization term:

$$\max_{\theta} \mathbb{E}_{x \sim \mathcal{D}, y \sim \pi_\theta(\cdot | x)} [r_\phi(x, y)] - \beta \mathbb{E}_{x \sim \mathcal{D}} [\text{KL}(\pi_\theta(\cdot | x) \parallel \pi_{\text{ref}}(\cdot | x))]$$

where  $\beta$  controls the strength of the Kullback–Leibler(KL) regularization. This objective is typically optimized using policy gradient methods such as Proximal Policy Optimization (PPO) (Schulman et al. 2017).

While RLHF has demonstrated remarkable success across multiple domains, it suffers from several limitations including training instability, high computational requirements due to the need for multiple model training phases, and potential reward hacking where the policy learns to exploit weaknesses in the reward model rather than genuinely improving quality (Bai et al. 2022).

## Direct Preference Optimization

Direct Preference Optimization (DPO) (Rafailov et al. 2023) addresses the limitations of RLHF by **eliminating the explicit reward model** training phase. DPO leverages a key theoretical insight: under the optimal policy for the RLHF objective, a closed-form solution exists between the optimal policy, the reference policy, and the implicit reward function.

Specifically, if  $\pi^*$  represents the optimal policy for the RLHF objective, then:

$$\pi^*(y|x) = \frac{1}{Z(x)} \pi_{\text{ref}}(y|x) \exp\left(\frac{1}{\beta} r(x, y)\right)$$

where  $Z(x)$  is a partition function. Rearranging this relationship yields:

$$r(x, y) = \beta \log \frac{\pi^*(y|x)}{\pi_{\text{ref}}(y|x)} + \beta \log Z(x)$$

Since the partition function is independent of  $y$ , preference comparisons eliminate this term, allowing DPO to directly optimize the policy using the reparameterized objective:

$$\mathcal{L}_{\text{DPO}}(\pi_\theta) = - \mathbb{E}_{(x, y^w, y^l) \sim \mathcal{D}} \left[ \log \sigma \left( \beta \log \frac{\pi_\theta(y^w|x)}{\pi_{\text{ref}}(y^w|x)} - \beta \log \frac{\pi_\theta(y^l|x)}{\pi_{\text{ref}}(y^l|x)} \right) \right]$$

This formulation enables stable optimization while maintaining the expressiveness of reward-based methods, offering significant computational advantages over traditional RLHF approaches.

## Inference-Time Alignment Paradigms

Inference-time alignment techniques offer complementary advantages by enabling **preference optimization without model retraining**. These methods modify the generation process itself to satisfy preference constraints through techniques such as:

**Contrastive Decoding:** Amplifying probability mass from preferred directions while suppressing undesired outputs through contrastive search procedures that balance exploration and exploitation during generation (Li et al. 2023).

**Preference-Conditioned Sampling:** Modifying sampling distributions to favor outputs that satisfy specified preference criteria, often through energy-based reweighting or guided generation techniques (Bai et al. 2022).

**Control Vector Steering:** Manipulating intermediate representations within generative models to guide outputs toward desired characteristics without explicit retraining (Subramani, Suresh, and Peters 2022).

In music generation, this flexibility proves essential for balancing competing preferences such as text adherence, audio quality, stylistic consistency, and temporal coherence. The ability to dynamically adjust these trade-offs during inference enables responsive adaptation to user preferences and contextual requirements.

## Musical Preference Complexity

Traditional evaluation metrics such as Fréchet Audio Distance (FAD), Inception Score (IS), or CLAP-based text-audio similarity fail to capture musical appreciation’s subjective nature. Musical preferences emerge from:

**Temporal Structure:** Music unfolds over time with hierarchical patterns spanning multiple scales from local rhythmic patterns (milliseconds to seconds) to phrase-level structure (seconds to minutes) to global compositional form (minutes to hours) (Lerdahl and Jackendoff 1996). Preference alignment must account for coherence across all these temporal scales simultaneously.

**Harmonic Content:** Western tonal music relies on sophisticated harmonic relationships that create expectations and resolutions. Preferences depend on whether these harmonic progressions feel satisfying, surprising, or emotionally appropriate within their musical context (Meyer 1954).

**Cultural Context:** Musical preferences are deeply embedded within cultural frameworks that shape aesthetic expectations. What constitutes “good” music varies dramatically across musical traditions, historical periods, and social communities (Liu, Hu, and Schedl 2018).

**Functional Considerations:** Music serves diverse functional roles—from background ambiance to focused listening, from dance accompaniment to game music and emotional expression—each requiring different optimization criteria (Herremans, Chuan, and Chew 2017).

These factors necessitate specialized alignment approaches that can handle the multi-dimensional nature of musical preference while respecting the inherent subjectivity and cultural specificity of musical aesthetics.

## Preference Optimization in Music

Application of preference optimization to music generation presents unique opportunities and challenges that distinguish it from text-based preference alignment. We discuss three prominent approaches here.

### Large-Scale Preference Learning with MusicRL

MusicRL (Cideron et al. 2024) pioneered **large-scale preference learning** in music generation by fine-tuning a pre-trained autoregressive MusicLM model using both expert-designed reward functions and massive human preference datasets. The system represents a groundbreaking implementation that demonstrates how preference alignment can bridge the gap between computational optimization and human musical appreciation at unprecedented scale.

The MusicRL framework implements two complementary optimization strategies that address different aspects of musical preference alignment. MusicRL-R employs explicit reward functions designed in collaboration with expert raters, focusing on sequence-level rewards related to text adherence and audio quality. These reward functions capture measurable aspects of musical quality including semantic alignment between textual descriptions and generated audio, perceptual audio quality measures that assess technical fidelity and artifacts, and musical structure coherence, evaluating temporal consistency and harmonic progression.

MusicRL-U incorporates human feedback at scale through Reinforcement Learning from Human Feedback (RLHF). The system collected a private dataset containing pairwise preferences from real users, creating the largest known dataset of human preferences for music generation. This massive preference collection enables training a sophisticated preference model that captures subtle aspects of subjectivity.

The preference data collection process involved presenting users with pairs of generated musical pieces and asking them to indicate preferences based on overall quality, musical coherence, and text adherence. This process revealed important insights about musical preference: human evaluators consistently demonstrate preferences that extend beyond simple audio quality or text matching, indicating aesthetic judgments that traditional metrics fail to capture.

Experimental results demonstrate that both MusicRL-R and MusicRL-U significantly outperform baseline MusicLM in human evaluations across diverse musical styles. The combined MusicRL-RU approach achieves the strongest performance. Crucially, ablation studies reveal that text adherence and audio quality account for only a portion of human preferences, underscoring the prevalence of subjectivity in musical appreciation and highlighting the inadequacy of current evaluation metrics.

### Multi-Preference Diffusion Optimization with DiffRhythm+

DiffRhythm+ (Chen et al. 2025) extends preference optimization to **diffusion-based architectures** for controllable full-length song generation, demonstrating how Direct Preference Optimization (DPO) can be effectively integrated into modern generative frameworks. The system employs multi-modal style conditioning through MuLan embeddings, enabling precise control over musical attributes while maintaining preference alignment.

The framework integrates Direct Preference Optimization directly into the diffusion training process using a modified DPO objective adapted for continuous latent spaces. Unlike discrete sequence models, diffusion models operate in continuous spaces, requiring careful adaptation of preference optimization techniques.

DiffRhythm+ incorporates metrics from multiple evaluation frameworks including SongEval (Yao et al. 2025) and Audiobox-aesthetic (Tjandra et al. 2025), enabling optimization for diverse preference criteria simultaneously. SongEval provides a more comprehensive assessment that includes structural coherence and memorability, evaluating how well generated music maintains logical harmonic progressions and consistent rhythmic patterns across extended compositions. Audiobox-aesthetic metrics focus on perceptual quality and aesthetic appeal, incorporating learned models that predict human aesthetic judgments.

The integration of DPO into diffusion architectures presents unique advantages for music generation. Unlike autoregressive models that generate sequentially, diffusion architectures can optimize global structure and long-range dependencies simultaneously (Wang et al. 2024), proving particularly valuable for music where harmonic progressions and rhythmic patterns must maintain coherence across extended temporal spans. Experimental validation demonstrates substantial improvements over baseline diffusion models, with particularly strong performance in long-form musical coherence and aesthetic appeal (Chen et al. 2025).

Both MusicRL and DiffRhythm+ demonstrate the critical importance of multi-objective optimization in musical preference alignment. Music generation must simultaneously satisfy competing constraints across multiple dimensions that often conflict with each other, requiring sophisticated balancing mechanisms.

The success of these methods shows that DPO can effectively balance generation objectives while maintaining diversity, essential for creative applications. However, the complexity of multi-objective optimization in music generation suggests opportunities for more sophisticated prefer-

ence modeling techniques that can explicitly represent and balance competing criteria during optimization.

## Inference-Time Alignment Techniques

Text2midi-InferAlign (Roy, Puri, and Herremans 2025) demonstrates inference-time preference optimization through a tree search methodology that balances multiple reward objectives at inference. The framework employs a composite reward function combining (weighted) text-audio consistency scores ( $S_{\text{text}}(y_t, x)$ ) with harmonic consistency scores ( $S_{\text{harmony}}$ ):

$$\text{Score}(y_t, x) = \alpha \cdot S_{\text{text}}(y_t, x) + \beta \cdot S_{\text{harmony}}(y_t)$$

The system incorporates caption mutation for exploration, generating semantically related variations of input captions to explore different musical interpretations while maintaining core meaning. The reward mechanism efficiently navigates the search space without exhaustive enumeration.

Experimental results demonstrate effectiveness: this approach achieves 29.4% improvement in CLAP scores compared to baseline text2midi generation (Bhandari et al. 2025) without any model modification, preserving diversity while enhancing quality. However, the tree search process requires some computational overhead, creating trade-offs between quality improvements and inference latency that future work must address for real-time applications.

## Benchmarking and Evaluation

Evaluating preference-aligned music generation requires metrics that capture both objective technical quality and subjective human judgment. Traditional metrics like Fréchet Audio Distance (FAD) (Kilgour et al. 2018) and Inception Score (IS) (Barratt and Sharma 2018) provide baselines but fail to capture musical-specific qualities or aesthetic appeal.

Recent work introduces dedicated frameworks: MusicRL’s pairwise preferences data demonstrates that automated metrics correlate only partially with human judgment. Text2midi-InferAlign employs CLAP scores for text-audio consistency alongside harmonic consistency metrics. DiffRhythm+ incorporates SongEval and Audiobox-aesthetic frameworks for comprehensive quality assessment.

Key evaluation challenges include developing metrics for diverse musical traditions and tastes, creating standardized benchmarks for long-form musical coherence, and establishing protocols for personalized preference evaluation. Cross-cultural validity remains important as many current frameworks often reflect recent western musical content (rock, pop, electronic primarily), limiting global applicability (Melechovsky et al. 2023).

## Technical Implementation Considerations

Practical deployment of preference-aligned music generation requires addressing computational scalability, data management, and user interface design. Large-scale models like MusicLM require billions of parameters, while preference learning adds overhead through RLHF (or DPO) training

and inference-time optimization creates additional computational costs at inference.

The architectural considerations for preference alignment systems extend beyond model parameters to encompass sophisticated data pipelines and training infrastructure. MusicRL’s success with preference data required specialized data collection platforms capable of presenting musical pairs to human evaluators while maintaining consistent quality assessment protocols. The system architecture must also have quality control mechanisms to identify inconsistent evaluators, and sophisticated aggregation algorithms that account for individual biases and cultural backgrounds.

Storage and retrieval systems for large-scale preference datasets present unique challenges given the multimodal nature of musical data. Unlike text preferences, musical preferences often require storing both symbolic representations such as MIDI (Melechovsky, Roy, and Herremans 2024) or audio (Roy et al. 2025), creating substantial storage requirements. Efficient indexing schemes must enable rapid retrieval of similar musical content for preference model training while supporting diverse query types including harmonic similarity, rhythmic patterns, and stylistic attributes.

Training infrastructure requirements vary significantly across different preference alignment approaches. DiffRhythm+’s integration of DPO into diffusion architectures demands careful memory management during reverse diffusion processes, as preference optimization requires maintaining gradients through the entire denoising chain. This creates memory scaling challenges that exceed those of standard diffusion training, necessitating techniques such as gradient checkpointing (Chen et al. 2016) and mixed-precision computation (Micikevicius et al. 2017) to enable training on practical hardware configurations.

Deployment considerations must account for the distinct computational demands of different preference alignment strategies. Training-time methods like MusicRL add computational costs during development but enable efficient deployment with minimal inference overhead. Conversely, inference-time optimization approaches like Text2midi-InferAlign add computational requirements after deployment, requiring sophisticated resource allocation to maintain consistent response times.

Quality assurance and monitoring present ongoing challenges for deployed preference-aligned systems. Unlike traditional generation systems where quality can be assessed through automated metrics, preference-aligned systems require continuous evaluation of alignment with human preferences that may evolve over time. This necessitates development of automated monitoring systems capable of detecting preference drift and triggering retraining procedures when alignment quality degrades below acceptable thresholds.

To sum up, key implementation considerations include: collecting high-quality preference data at scale, ensuring privacy protection and consent for sensitive musical preference data, addressing bias in rater demographics and cultural representation, and designing intuitive user interfaces that enable effective preference specification while maintaining system explainability and quality assurance.

Challenge	Key Issues
<b>Scalability</b>	Long-form composition modeling, attention complexity, hierarchical structures
<b>Multi-modal</b>	Video-music sync, cross-cultural integration, real-time adaptation
<b>Personalization</b>	Few-shot learning, individual aesthetics, cultural awareness
<b>Robustness</b>	Adversarial attacks, bias amplification, quality degradation
<b>Efficiency</b>	Inference overhead, energy costs, interactivity latency
<b>Evaluation</b>	Preference learning, cross-domain transfer, evaluation paradox

Table 1: Key Research Challenges in Music Preference Alignment

## Challenges and Open Research Questions

Despite recent advances, fundamental challenges remain.

**Long structure** is critical as current approaches struggle with long-form compositions due to attention complexity and hierarchical structure modelling across temporal scales (Zixun, Makris, and Herremans 2021).

**Multi-modal alignment** poses challenges for video-music synchronization and cross-modal media integration while maintaining musical coherence (Kang, Poria, and Herremans 2024).

**Personalization** remains largely unexplored, requiring few-shot preference learning methods and efficient representation learning for individual aesthetic judgments (Huron 2008). Cultural awareness is essential here, as current systems predominantly reflect Western musical traditions, limiting global applicability and risking cultural appropriation.

**Computational efficiency** challenges stem from inference-time optimization overhead, energy consumption, and latency requirements for interactive applications.

**Preference representation learning** MusicRL’s findings reveal that traditional metrics capture only partial aspects of human musical judgment, indicating the need for more sophisticated models that capture implicit musical understanding, emotional responses, and contextual appropriateness that human evaluators demonstrate.

**Data scarcity:** MusicRL’s 300,000-preference dataset remains proprietary. The lack of open, large-scale preference datasets limits research reproducibility and cross-cultural model development.

**Cross-domain generalization** presents significant challenges as evidenced by the domain-specific nature of current evaluation frameworks. DiffRhythm+’s reliance on SongEval and Audiobox-aesthetic metrics highlights the limitation of genre-specific evaluation approaches, while Text2midi-InferAlign’s CLAP-based optimization may not generalize effectively across diverse musical styles and cultural contexts. The development of universal preference alignment techniques that maintain effectiveness across musical genres, cultural traditions, and application domains remains an open research question.

**Dynamic preference adaptation** represents a critical gap in current methodologies. While few approaches demonstrate inference-time optimization capabilities, the field lacks strategies for adapting to evolving user preferences over time. This challenge encompasses both technical aspects—developing efficient incremental learning algorithms for preference models—and theoretical foundations for un-

derstanding how musical preferences change through exposure, expertise development, and cultural evolution.

The **evaluation paradox** in preference-aligned systems creates circular dependencies where assessment of preference alignment quality itself requires human judgment, potentially introducing the same biases and limitations.

## Vision and Roadmap

**Applications:** Preference-aligned systems can enable interactive composition tools, adaptive film scoring, game audio, therapeutic music generation (Agres et al. 2021), and personalized music services.

**Technical Priorities:** (1) Creating open, large-scale preference datasets spanning diverse cultures and personalization dimensions; (2) developing unified inference-time frameworks for multi-objective optimization with reduced overhead; (3) building cross-cultural evaluation frameworks and culturally-competent models through collaboration with ethnomusicologists; (4) enabling real-time adaptive systems for interactive human-AI co-creation.

Success requires interdisciplinary collaboration combining machine learning, music theory, cognitive science, and ethics to serve diverse creative needs.

## Conclusion

This paper advocates systematic preference alignment in music generation, addressing the gap between computational optimization and human appreciation. Through examining large-scale preference learning, inference-time optimization, and multi-preference diffusion, we identified promising directions and challenges.

Preference alignment improves quality and relevance while preserving diversity. Realizing full potential requires addressing scalability, culturally-aware frameworks, and multi-modal capabilities.

We call for interdisciplinary research combining machine learning, music theory, and human-computer interaction to create music AI serving human creative needs in meaningful alignment with human musical preference and cultural expression.

## Acknowledgements

This work has received support from SUTD’s Kickstart Initiative under grant number SKI 2021 04 06 and MOE under grant number MOE-T2EP20124-0014.

## References

- Agostinelli, A.; Denk, T. I.; Borsos, Z.; Engel, J.; Verzetti, M.; Caillon, A.; Huang, Q.; Jansen, A.; Roberts, A.; Tagliasacchi, M.; et al. 2023. Musiclm: Generating music from text. *arXiv:2301.11325*.
- Agres, K.; Forth, J.; and Wiggins, G. A. 2016. Evaluation of musical creativity and musical metacreation systems. *Computers in Entertainment (CIE)*, 14(3): 1–33.
- Agres, K. R.; Schaefer, R. S.; Volk, A.; Van Hooren, S.; Holzapfel, A.; Dalla Bella, S.; Müller, M.; De Witte, M.; Herremans, D.; Ramirez Melendez, R.; et al. 2021. Music, computing, and health: a roadmap for the current and future roles of music technology for health care and well-being. *Music & Science*, 4: 2059204321997709.
- Bai, Y.; Kadavath, S.; Kundu, S.; Askell, A.; Kernion, J.; Jones, A.; Chen, A.; Goldie, A.; Mirhoseini, A.; McKinnon, C.; et al. 2022. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*.
- Barratt, S.; and Sharma, R. 2018. A note on the inception score. *arXiv preprint arXiv:1801.01973*.
- Bhandari, K.; Roy, A.; Wang, K.; Puri, G.; Colton, S.; and Herremans, D. 2025. Text2midi: Generating symbolic music from captions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 23478–23486.
- Blei, D. M.; Ng, A. Y.; and Jordan, M. I. 2003. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan): 993–1022.
- Bradley, R. A.; and Terry, M. E. 1952. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4): 324–345.
- Chen, H.; Jiang, Y.; Ma, G.; Hao, C.; Wang, S.; Yao, J.; Ning, Z.; Meng, M.; Luan, J.; and Xie, L. 2025. DiffRhythm+: Controllable and Flexible Full-Length Song Generation with Preference Optimization. *arXiv:2507.12890*.
- Chen, T.; Xu, B.; Zhang, C.; and Guestrin, C. 2016. Training deep nets with sublinear memory cost. *arXiv preprint arXiv:1604.06174*.
- Christiano, P. F.; Leike, J.; Brown, T.; Martic, M.; Legg, S.; and Amodei, D. 2017. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30.
- Cideron, G.; Girgin, S.; Verzetti, M.; Vincent, D.; Kastelic, M.; Borsos, Z.; McWilliams, B.; Ungureanu, V.; Bachem, O.; Pietquin, O.; et al. 2024. Musicrl: Aligning music generation to human preferences. *arXiv:2402.04229*.
- Copet, J.; Kreuk, F.; Gat, I.; Remez, T.; Kant, D.; Synnaeve, G.; Adi, Y.; and Défossez, A. 2023. Simple and controllable music generation. *Advances in Neural Information Processing Systems*, 36: 47704–47720.
- Dhariwal, P.; Jun, H.; Payne, C.; Kim, J. W.; Radford, A.; and Sutskever, I. 2020. Jukebox: A generative model for music. *arXiv:2005.00341*.
- Herremans, D.; Chuan, C.-H.; and Chew, E. 2017. A functional taxonomy of music generation systems. *ACM Computing Surveys (CSUR)*, 50(5): 1–30.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30.
- Huron, D. 2008. *Sweet anticipation: Music and the psychology of expectation*. MIT press.
- Jaques12, N.; Gu13, S.; Turner, R. E.; and Eck, D. 2016. Generating music by fine-tuning recurrent neural networks with reinforcement learning.
- Juslin, P. N.; and Västfjäll, D. 2008. Emotional responses to music: The need to consider underlying mechanisms. *Behavioral and brain sciences*, 31(5): 559–575.
- Kang, J.; Poria, S.; and Herremans, D. 2024. Video2music: Suitable music generation from videos using an affective multimodal transformer model. *Expert Systems with Applications*, 249: 123640.
- Kilgour, K.; Zuluaga, M.; Roblek, D.; and Sharifi, M. 2018. Fr\`echet audio distance: A metric for evaluating music enhancement algorithms. *arXiv preprint arXiv:1812.08466*.
- Koo, J.; Wichern, G.; Germain, F. G.; Khurana, S.; and Le Roux, J. 2025. Smitin: Self-monitored inference-time intervention for generative music transformers. *IEEE Open Journal of Signal Processing*.
- Le Groux, S.; and Verschure, P. 2010. Towards adaptive music generation by reinforcement learning of musical tension. In *proceedings of the 6th sound and music conference, Barcelona, Spain*, volume 134.
- Lerdahl, F.; and Jackendoff, R. S. 1996. *A Generative Theory of Tonal Music, reissue, with a new preface*. MIT press.
- Li, K.; Patel, O.; Viégas, F.; Pfister, H.; and Wattenberg, M. 2023. Inference-time intervention: Eliciting truthful answers from a language model. *Advances in Neural Information Processing Systems*, 36: 41451–41530.
- Liu, M.; Hu, X.; and Schedl, M. 2018. The relation of culture, socio-economics, and friendship to music preferences: A large-scale, cross-country study. *PloS one*, 13(12): e0208186.
- Liu, R.; Hung, C.-Y.; Majumder, N.; Gautreaux, T.; Bagherzadeh, A. A.; Li, C.; Herremans, D.; and Poria, S. 2025. JAM: A Tiny Flow-based Song Generator with Fine-grained Controllability and Aesthetic Alignment. *arXiv:2507.20880*.
- McAdams, S. 2013. Musical timbre perception. *The psychology of music*, 3.
- Melechovsky, J.; Guo, Z.; Ghosal, D.; Majumder, N.; Herremans, D.; and Poria, S. 2023. Mustango: Toward controllable text-to-music generation. *arXiv:2311.08355*.
- Melechovsky, J.; Roy, A.; and Herremans, D. 2024. MidiCaps: A large-scale MIDI dataset with text captions. *arXiv preprint arXiv:2406.02255*.
- Meyer, L. B. 1954. *Emotion and meaning in music*. Ph.D. thesis, The University of Chicago.
- Micikevicius, P.; Narang, S.; Alben, J.; Diamos, G.; Elsen, E.; Garcia, D.; Ginsburg, B.; Houston, M.; Kuchaiev, O.; Venkatesh, G.; et al. 2017. Mixed precision training. *arXiv preprint arXiv:1710.03740*.

North, A. C.; Hargreaves, D. J.; and Krause, A. E. 2016. Music and consumer behavior.

Novack, Z.; McAuley, J.; Berg-Kirkpatrick, T.; and Bryan, N. J. 2024. Ditto: Diffusion inference-time t-optimization for music generation. *arXiv preprint arXiv:2401.12179*.

Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2023. Direct preference optimization: Your language model is secretly a reward model. *Advances in neural information processing systems*, 36: 53728–53741.

Roy, A.; Liu, R.; Lu, T.; and Herremans, D. 2025. Jamendo-maxcaps: A large scale music-caption dataset with imputed metadata. *arXiv preprint arXiv:2502.07461*.

Roy, A.; Puri, G.; and Herremans, D. 2025. Text2midi-InferAlign: Improving Symbolic Music Generation with Inference-Time Alignment. *arXiv:2505.12669*.

Salimans, T.; Goodfellow, I.; Zaremba, W.; Cheung, V.; Radford, A.; and Chen, X. 2016. Improved techniques for training gans. *Advances in neural information processing systems*, 29.

Schedl, M.; Gómez, E.; Urbano, J.; et al. 2014. Music information retrieval: Recent developments and applications. *Foundations and Trends® in Information Retrieval*, 8(2-3): 127–261.

Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.

Subramani, N.; Suresh, N.; and Peters, M. E. 2022. Extracting latent steering vectors from pretrained language models. *arXiv preprint arXiv:2205.05124*.

Tjandra, A.; Wu, Y.-C.; Guo, B.; Hoffman, J.; Ellis, B.; Vyas, A.; Shi, B.; Chen, S.; Le, M.; Zacharov, N.; et al. 2025. Meta audiobox aesthetics: Unified automatic quality assessment for speech, music, and sound. *arXiv preprint arXiv:2502.05139*.

Wang, Y.; Ren, S.; Gao, R.; Yao, L.; Guo, Q.; An, K.; Bai, J.; and Sun, X. 2024. LaDiC: Are diffusion models really inferior to autoregressive counterparts for image-to-text generation? *arXiv preprint arXiv:2404.10763*.

Wang, Y.; Wu, S.; Hu, J.; Du, X.; Peng, Y.; Huang, Y.; Fan, S.; Li, X.; Yu, F.; and Sun, M. 2025. Notagen: Advancing musicality in symbolic music generation with large language model training paradigms. *arXiv preprint arXiv:2502.18008*.

Yao, J.; Ma, G.; Xue, H.; Chen, H.; Hao, C.; Jiang, Y.; Liu, H.; Yuan, R.; Xu, J.; Xue, W.; et al. 2025. SongEval: A Benchmark Dataset for Song Aesthetics Evaluation. *arXiv preprint arXiv:2505.10793*.

Zixun, G.; Makris, D.; and Herremans, D. 2021. Hierarchical recurrent neural networks for conditional melody generation with long-term structure. In *2021 international joint conference on neural networks (IJCNN)*, 1–8. IEEE.