

RISE: ADAPTIVE MUSIC PLAYBACK FOR REALTIME INTENSITY SYNCHRONIZATION WITH EXERCISE

Alexander Wang
University of Michigan
alexhci@umich.edu

Chris Donahue
Carnegie Mellon University
chrisdonahue@andrew.cmu.edu

Dhruv Jain
University of Michigan
profdj@umich.edu

ABSTRACT

We propose a system to adapt a user’s music to their exercise by aligning high-energy music segments with intense intervals of the workout. Listening to music during exercise can boost motivation and performance. However, the structure of the music may be different from the user’s natural phases of rest and work, causing users to rest longer than needed while waiting for a motivational section, or lose motivation mid-work if the section ends too soon. To address this, our system, called RISE, automatically estimates the intense segments in music and uses cutpoint-based music rearrangement techniques to dynamically extend and shorten different segments of the user’s song to fit the ongoing exercise routine. Our system takes as input the rest and work durations to guide adaptation. Currently, this is determined either via a pre-defined plan or manual input during the workout. We evaluated RISE with 12 participants who compared our system to a non-adaptive music baseline while exercising in our lab. Participants found our rearrangements seamless, intensity estimation accurate, and many recalled moments when intensity alignment helped them push through their workout.

1. INTRODUCTION

The right music can transform an ordinary workout into a more engaging and motivating experience. Research has shown that listening to music while working out can improve physical performance and reduce perceived exertion [1]. Consequently, many commercial applications offer music recommendations to fit the user’s exercise routines [2, 3]. However, current applications do not consider musical structure, treating songs as uniform in intensity. This can lead to suboptimal scenarios, such as a brief calm section of an otherwise high-energy song playing just as the listener begins weightlifting. Conventional music playback remains static and unresponsive to the listener’s dynamic behavior and context, highlighting the need for adaptive music systems that customize the listening experience in real time.

We propose RISE, a real-time intensity synchronization system that adapts the playback of user-selected music to better align with exercise plans (Figure 1). RISE integrates several Music Information Retrieval (MIR) techniques to segment songs, estimate segment intensity, and identify cutpoints, enabling seamless looping and skipping to align segment durations with user exercise. RISE operates in two scenarios: in the *guided scenario*, it adapts a music recording to fit a predefined workout plan, while in the *un-guided scenario*, it adapts in real-time to spontaneous exercise changes, offering flexibility at the potential cost of adaptation quality. In this work, we focus on the music analysis and adaptive playback, assuming that the ground truth exercise state is given as input to the system.

We validate our solution through a two-pronged evaluation. First, we assessed the seamlessness of our system’s adaptations by generating excerpts from 270 songs and conducting an in-lab study in which participants rated transition naturalness. Results showed that our system’s adaptations were perceived as seamless, scoring similarly to unmodified excerpts. Second, in a controlled user study, 12 participants compared exercising with adaptive music to nonadaptive music. Participants found our adaptations seamless and intensity segmentation accurate. They appreciated how the alignment of music with workout enhanced motivation and overall experience, particularly when extending high-energy segments helped them push through.

In summary, our work makes two primary contributions. First, we propose a novel music adaptation system for exercises by applying several existing MIR techniques. Second, we offer insights from two user studies demonstrating the utility of music adaptation for exercise. Moreover, as the first to explore cutpoint-based rearrangement techniques for adapting music to exercise, our work lays the foundation for future research in applying adaptive music systems to scenarios beyond exercise.¹

2. RELATED WORK

Context-aware music systems enrich user experience by dynamically adapting music based on usage contexts, a focus of MIR research emphasizing user-centered applications [4–6]. For example, prior work has explored adaptive music for enhancing driving experiences [7, 8], notification experiences [9, 10], and narrative experiences [11, 12]. Given the well-documented benefits of music in



© A. Wang, C. Donahue, and D. Jain. Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). **Attribution:** A. Wang, C. Donahue, and D. Jain, “RISE: Adaptive Music Playback for Realtime Intensity Synchronization with Exercise”, in *Proc. of the 26th Int. Society for Music Information Retrieval Conf.*, Daejeon, South Korea, 2025.

¹ Supplemental video: <https://youtu.be/XZLBfpt6Lgg>.




Figure 1. We present RISE, a real-time music intensity synchronization system for exercise. Our system takes in user music and workout phases as input and aligns the high-intensity segments of the music with the user’s exertion phases (and vice versa), to enhance the workout experience.

exercise contexts [1], prior work has also explored music adaptations for workouts, including song recommendation [2, 13, 14] and tempo adjustments based on running pace and heart rate [15–19].

While these approaches personalize music at the song level, they overlook the motivational affordances of specific segments within a song. Relaxing music is better suited for recovery, while upbeat, rhythmically prominent music stimulates physical activity [20]. These variations in intensity also occur within a single song, where high-energy sections like the ‘drop’ in electronic music feel particularly exciting [21] and exercisers naturally synchronize bursts of effort with these peaks [22]. We leverage these insights to develop music adaptation algorithms that dynamically adjust individual segments within a song, aligning high-intensity music with peak exertion.

RISE combines several MIR techniques for structure-aware adaptive playback. We use the All-in-one package [23] for *music structure analysis*, automatically segmenting songs into sections (e.g., intro, verse). Then, we use Spleeter [24] for *source separation*, isolating individual audio components (e.g., drum tracks) to estimate each segment’s drum prominence. Finally, we use *cutpoint identification* [25] to pinpoint pairs of transition points where cutting from one to the other wouldn’t sound abrupt. This enables RISE to dynamically skip or loop musical segments.

3. RISE

We present RISE, a system that synchronizes music intensity with workout intensity to enhance the exercise experience (Figure 2). RISE takes in a user’s music (audio) and workout (planned intervals in the guided setting or current status in the unguided setting) as input. The music is pre-processed to label high-intensity segments (Section 3.1) and points of transition (Section 3.2). This information is then used for real-time music adaptation by our Unity-based system for workout alignment (Section 3.3).

3.1 Preprocessing - Estimating Intense Segments in Music

Informed by prior work, we select chorus and instrumental sections as the high-intensity segments, as they are often perceived as more energetic and motivating [21, 22]. To ensure rhythmic drive [26, 27], we filter out sections lacking in drum presence to avoid arrangements where percussion is intentionally dropped or simplified for contrast.

We estimate these segments by leveraging two existing music analysis systems. First, we use a music structure analysis system [23] to segment each song into a set of functional sections $S = \{s_1, s_2, \dots, s_N\}$ and extract beat timestamps $B = \{b_1, b_2, \dots, b_M\}$. Here, S represents a partition of the song into N sections, where s_n marks the start timestamp of a section n and s_{n+1} marks both its end and the start of the next section. B is the time-ordered set of detected beat times. Functional section labels (e.g., chorus, instrumental, verse, bridge) are assigned to each section by the system.

To measure drum prominence, we first extract the drum track using a source separation system [24]. Given an audio input x , we define the isolated drum signal as $d := \text{SourceSepDrums}(x)$. We then compute LUFS loudness [28] at beat-level intervals. For each beat index m , we define its loudness as:

$$L_d(m) := \text{LUFS}(\text{Segment}(d, b_m, b_{m+1})),$$

where $\text{Segment}(d, b_m, b_{m+1})$ extracts the segment of the drum audio between consecutive beats b_m and b_{m+1} .

The average drum loudness for a section s_n is calculated by averaging the loudness of all beats in the segment:

$$L_d(s_n) := \frac{1}{|B_n|} \sum_{m \in B_n} L_d(m),$$

where $B_n \subseteq \{1, \dots, M\}$ is the set of beat indices that satisfy $s_n \leq b_m < s_{n+1}$.

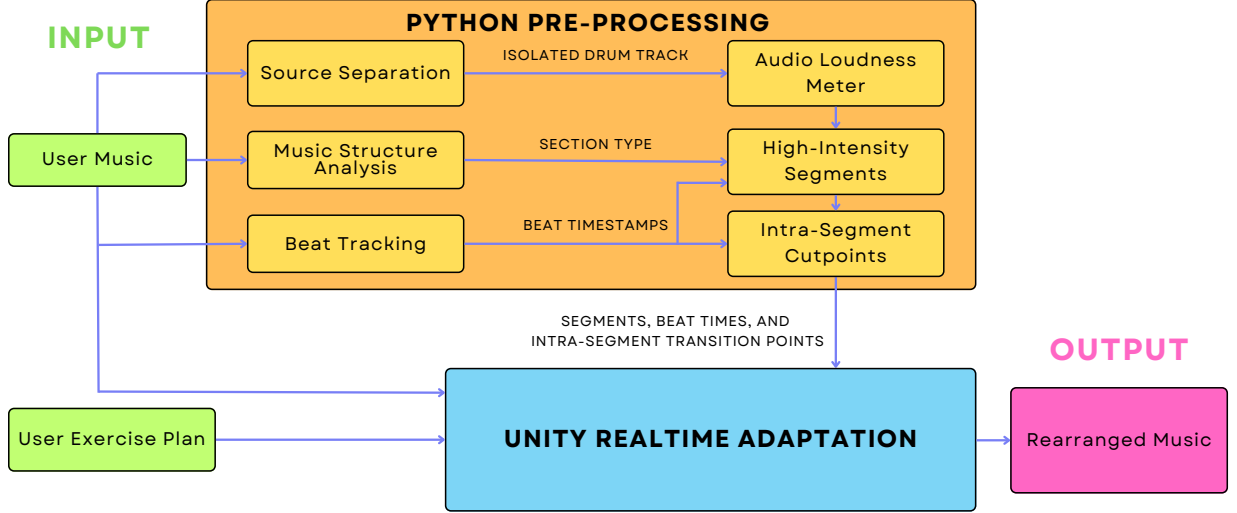


Figure 2. System overview. RISE takes user music and exercise plans as input, preprocesses the music to identify intense segments and intra-segment cutpoints, and sends this information to Unity for real-time adaptation.

We precompute a per-song loudness threshold τ , where any segment with a loudness above τ is considered high intensity, and otherwise we consider it low intensity. Threshold τ is computed relative to the loudest section in the song: $\tau = \max_{s_n \in S} L_d(s_n) - \delta$, where δ is a constant that determines the relative threshold. In our implementation, we set $\delta = 5$ decibels. If four or more consecutive segments are labeled high-intensity, we reassign the segment with the lowest loudness as low-intensity and merge consecutive segments with the same intensity labels. Ultimately, our analysis induces a partition of the full track into $\leq N$ segments and associated binary intensity labels $S' = \{(s'_1, i_1), (s'_2, i_2), \dots\}$, where s'_n are delineating timestamps and $i_i \in \{\text{Low, High}\}$.

3.2 Preprocessing - Estimating Intra-Segment Cutpoints

To better align intense segments of music with high-intensity workout phases, we estimate a set of *cutpoints* to facilitate seamless adaptation. A cutpoint is a pair of timestamps in the music recording, consisting of a starting timestamp and a destination timestamp. The objective is to estimate cutpoints that allow smooth musical transitions, such that if playback jumps from the start of a cutpoint to its destination, users experience minimal disruption.

We derive an initial set of cutpoints using the approach proposed by Plachouras and Miron [25], which analyzes recurrence matrices encoding the self-similarity of musical beats. Cutpoints are identified by detecting diagonals in these matrices that correspond to repeated patterns, pinpointing transitions between musically coherent sections. This results in an initial set of candidate cutpoints:

$$C = \{(c_i^{\text{orig.}}, c_i^{\text{dest.}}) \in B \times B\}$$

where C is the set of all estimated cutpoint pairs, and each cutpoint consists of a start time $c_i^{\text{orig.}}$ and an end time $c_i^{\text{dest.}}$, both aligned to detected beat timestamps B .

In prior work, cutpoints have been applied to rearrange music to fit external constraints, such as video duration [29]. These approaches allow cutpoints to cross section boundaries, maximizing flexibility at the potential cost of playback naturalness.

To prioritize naturalness in our system, we enforce an *intra-segment* constraint, ensuring that cutpoints only jump within a segment as opposed to across segments. We define the filtered set of intra-segment cutpoints as:

$$C' := \{(c_i^{\text{orig.}}, c_i^{\text{dest.}}) \in C \mid \exists s'_n \in S', s'_n \leq c_i^{\text{orig.}}, c_i^{\text{dest.}} < s'_{n+1}\}$$

This guarantees that every cutpoint's start and destination timestamps fall within the same functional section s'_n , preserving the structural integrity of the music.

3.3 Adaptation With Cutpoints

In addition to music audio and exercise plan, our real-time adaptation system takes as input the following information estimated during pre-processing: (1) musical segments and corresponding intensities, (2) seamless cutpoints, and (3) beat timestamps. To adapt music in real time, we define a **state machine** that governs playback behavior. The system operates in one of three possible states (Figure 3):

- **Loop State:** If the system determines that a segment should be extended, it selects a cutpoint c_i where $c_i^{\text{orig.}} > t_{\text{current}}$ and $c_i^{\text{dest.}} < t_{\text{current}}$, looping previously played sections to increase the duration of the current intensity segment.
- **Skip State:** If a segment duration needs to be shortened, the system selects a cutpoint c_i where $c_i^{\text{orig.}} > t_{\text{current}}$ and $c_i^{\text{dest.}} > c_i^{\text{orig.}}$, skipping forward to reduce the segment duration.
- **Unmodified State:** If no transition is required, the music plays continuously without alteration.

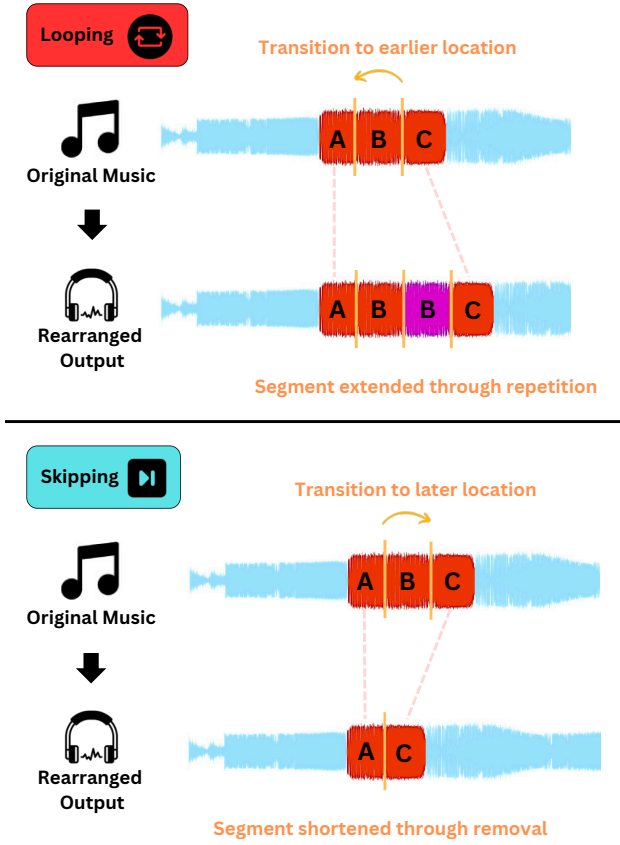


Figure 3. Visualization of adaptation modes. Top: Loop mode extends a segment by jumping back. Bottom: Skip mode shortens a segment by jumping forward.

Playback state transitions are determined dynamically based on the user’s exercise plan, enabling the system to adjust segment duration in real time.

3.3.1 Filter-based transitions

Sometimes, the system cannot immediately transition to the target intensity because there are no available cutpoints that matches the desired transition. In these cases, we use a *filter-based transition*, gradually removing high-frequency content before the transition and restoring it afterward, similar to DJ fade techniques. These transitions are noticeable and less seamless compared to cutpoint transitions. We currently apply filter-based transitions only in unguided mode (section 3.4.1), when the exercise calls for a timely transition to a high-intensity music state.

3.4 Usage Modes

Workout habits may vary greatly across different types of exercises—weight training can require minutes of rest to fully recover and the actual timing can vary greatly depending on how exhausted the user is. Guided interval workouts, on the other hand, emphasize short bursts followed by short rest periods that are strictly timed to maximize time efficiency. Motivated by this observation, we designed two usage modes for two scenarios. An *unguided*




Figure 4. Adaptation mode for different scenarios.

mode where the user is free to rest as long as they need, and a *guided mode* where the user follows a predefined exercise plan. We detail the system design of each mode below.

3.4.1 Unguided use

The unguided mode allows users to freely choose start times and work/rest durations, but sometimes sacrifice adaptation quality by using filter-based transitions when no cutpoints are immediately available. The system takes real-time workout state as input and adjusts the music accordingly (binary: work/rest). Currently, users manually indicate state changes by pressing a button. As depicted in Figure 4, RISE transitions between playback states depending on the current workout status:

1. **Starting exercise during a low-intensity segment:** The system enters **skip mode** to quickly transition to a high-intensity segment.
2. **During exercise in a high-intensity segment:** The system activates **loop mode** to sustain high-intensity music until the user begins resting.
3. **Starting rest in a high-intensity segment:** The system disables looping and switches to **skip mode** to exit high-intensity segments.
4. **During rest in a low-intensity segment:** The system enters **unmodified mode**, allowing the music to play naturally until the user resumes exercising.

3.4.2 Guided use

The guided mode provides seamless adaptations but requires a precise exercise plan. The system takes two durations (in seconds) for work and rest. We iterate through available cutpoints and select the transition that results in an adaptation closest to the desired duration. We allow one transition per segment to maximize naturalness. The results are close to the specified duration but rarely perfect (e.g., adjusting a 50-second segment to 32 seconds for a 30-second target). We fade songs in at the beginning and out at the end, adjusting the intro to match the rest duration.

4. QUANTITATIVE EVALUATION

We conducted an in-lab quantitative listening evaluation to assess the seamlessness of transitions.

4.1 Procedure

We collected 270 different songs, 30 songs each from nine different official Spotify workout playlists with different genre preferences. 13 songs were removed from the study because they had similar intensity throughout the entire song or had no available intra-segment cutpoints. For the remaining 257 songs, we generated two 10-second audio clips per song, one with a transition and one unmodified. Each clip is selected from a random segment of the song. We randomized the transition timing to occur between two and eight seconds within the 10-second clip.

This evaluation was performed with six human listeners. Each clip was randomly assigned to two listeners who rated transition naturalness on a scale of 1-5 (1 = very jarring, 5 very seamless/unnoticeable). We compare the ratings of modified and unmodified clips using a paired t-test.

4.2 Results

We found no statistically significant differences between the clips with transitions ($M = 4.3/5$, $SD = 1.1$) and the baseline clips ($M = 4.5/5$, $SD = 1.0$), $t(256) = -1.63$, $p = .10$, suggesting that our transitions are highly seamless and comparable to the unmodified clips. We observed a small decrease in the average rating of transition clips (4.5 vs. 4.3) for two reasons. First, the beat detection algorithm is not perfect. We found instances where the transitions were not perfectly aligned, causing a slight jump in rhythm. Second, familiarity with a song influenced the detection of transitions. Raters noted that, even when transitions were completely natural, they perceived differences in expected progression (e.g., altered lyrics) in songs they knew well. We also found that unmodified clips did not receive perfect scores. This was due to structural elements such as syncopated rhythms and abrupt breaks, intended to surprise the listener, being perceived as transitions by raters, despite these elements being part of the original compositions.

5. USER STUDY

We conducted a user study to explore how users experience RISE in both guided and unguided exercise settings.

5.1 Study Design

Participants exercised to both unmodified music and our adaptive system across two blocks: guided interval training and unguided weight training. Within each block, they experienced both adaptive and non-adaptive conditions, with order fully counterbalanced. Each condition included a brief tutorial, 8 minutes of exercise, and optional rest. Participants selected two songs from a curated pool of 25 tracks, played identically across all conditions. The adaptive system modified the music in response to user activity, while the non-adaptive version left the music unchanged. The full session lasted approximately 90 minutes.

Guided interval training. Participants performed interval exercises of their choice (e.g., jumping jacks) according to a 40s work / 30s rest schedule. For the non-adaptive system, these intervals are strict. For the adaptive system, the actual timer may vary by seconds depending on the available transitions in each section. Instructions were shown on a screen with countdown visuals and sounds, modeled after popular workout timer videos [30].

Unguided weight training. Participants used dumbbells to perform any freeform weight exercises (e.g., bicep curls). In adaptive conditions, they verbally indicated when they were about to begin or end a work segment, allowing the researcher to input music adaptation state changes. No interface was shown.

Interview procedure. After introducing the study and obtaining participant demographic information and consent, we gave them a verbal description of our system and recorded their first impressions through a short interview. After experiencing all conditions, we showed participants how the system operated by replaying the music for them along with visualizations of both segment intensity labels and cutpoint transitions. This was done at the end of the study to avoid priming participants to focus on specific manipulations and to evaluate whether the adaptations were perceptible without guidance. The study concluded with a semi-structured exit interview for verbal feedback.

Participant and apparatus. We recruited 12 participants (5 female, 7 male, age $M = 24.8$ years; $SD = 2.1$) from a local university. Participants regularly exercise (3: 1-2x/week, 7: 3-4x/week, 2: 5-6x/week) for considerable durations (1: 15-30 min, 4: 30-45 min, 1: 45-60 min, 5: 1-2 hours, 1: 2 hours+) in a variety of exercise types (8: steady cardio, 3: interval training, 10: strength training, 4: other). The study took place in a controlled lab environment with music played through speakers for participant safety. Participants received \$50 for their participation.

5.2 Findings

Interview transcripts were thematically analyzed using a coding reliability approach [31] with the Recal2 tool [32]. The analysis resulted in an average raw accuracy of 86.1% and an average Krippendorff's alpha of 0.72 between two raters, indicating acceptable agreement ($\alpha > 0.66$). We present key findings below.

Participants were excited by the premise of their music adapting to their workouts. After we explained

the study to participants but before they experienced our system, we asked participants about their initial reactions to the description of the study and our proposed system. Some participants had already made manual efforts to align their workout with music ($n = 8$), such as coordinating specific music sections with their workout ($n = 4$). All participants expressed that music intensity alignment with workouts could be helpful to them. Participants believed alignment can increase motivation, or help them feel more energized and “pumped” during the workout ($n = 11$).

Estimated high-intensity segments aligned with their intuitions. We showed participants a visualization of our system’s estimated high-intensity music segments after the study, and all participants found the results to be aligned with their expectations ($n = 12$).

Participants found our cutpoint adaptation technique seamless, subverting their expectations. Before the study, we shared a high-level description of our workout adaptation system with participants. Many expressed initial skepticism that the modifications might detract from the naturalness of the music playback ($n = 6$). Out of the participants who had concerns regarding the naturalness of music, all but one remarked that their concerns were overturned after experiencing our system ($n = 5$). While the filter-based modifications were regarded as more noticeable, most participants either did not notice or barely noticed any cutpoint-based modifications until they were shown a visualization at the end of the study ($n = 10$).

Participants appreciated the extra motivation our alignment approach provided. All but one participant found the intensity alignment to enhance their workout experience ($n = 11$), giving them more motivation to push harder in their workouts. P2 noted that when the music intensity peaked just as they were about to reach failure in their workout, it helped them push through and get an extra one or two reps. Similarly, P4 recalled an instance where they expected the music to shift to a lower-intensity segment in the middle of their work interval, but instead the upbeat segment repeated, providing them the energy to push through the remainder of their workout.

Participants varied in their sensitivity to both music modifications and alignment with exercise. For instance, P8 barely noticed music changes but felt their movements were more consistent in the adaptive condition, while the nonadaptive one felt chaotic. In contrast, P3, who regularly coordinates music with workouts, immediately noticed misalignments in the nonadaptive condition and even felt the urge to stop the music when high-energy segments played during their rest. Sensitivity could also depend on exercise. P6 mentioned that they treated music as background during simple exercises but valued music alignment for motivation during more challenging exercises, such as the bench press. This indicates the need to better understand how different modes of music engagement influence tolerance and demand for system-driven changes.

Participants preferred adaptive music but identified issues with the unguided experience. Out of 12 participants, 11 preferred our adaptive system over nonadaptive

music: 6 preferred it in both scenarios, while 5 favored it only in the guided setting. Those who preferred the system only in the guided setting found two main issues with the unguided experience. First, the input method—requiring them to notify the researcher before exertion—was distracting and impractical ($n = 7$). Second, while cutpoint transitions were deemed seamless and often unnoticeable, filter transitions disrupted the natural flow of the music. These filter transitions were especially common during exercises with longer work durations and shorter rest periods, where excessive looping also occurred, leading to unnatural adaptations of the music. As a result, participants recommended improvements in system customizability to support different workout habits (*e.g.*, long work, short rest) and music adaptation preferences ($n = 8$), manipulation seamlessness ($n = 4$), and expanding the available music pool or allowing user input ($n = 3$).

6. LIMITATION AND FUTURE WORK

Our study revealed several limitations, including users finding manual input impractical and filter transitions disrupting. We propose future work below:

Fully automating the system with sensing technologies. To eliminate the need for manual interaction during workouts, we propose automating the system using real-time sensing technologies. By integrating activity recognition, the system could automatically detect exercise phases without manual input. The system could also model user behavior and predict upcoming exercise states based on past patterns to prepare adaptation plans in advance.

Incorporating other types of modifications. While cutpoints enable seamless transitions, they may be unavailable when timely adaptation is required. Future work could complement our segment-sensitive approach with existing techniques, such as pace adjustment and song recommendation, to compensate for minor time discrepancies. Additionally, exploring audio inpainting with generative models could enable smooth transitions between arbitrary segments, providing greater flexibility and precision.

Determining high-intensity segments. RISE currently aligns drum-prominent chorus and instrumental sections with user exertion. While this approach worked well in our study, future work can further explore how segment-level musical variations influence exercise and how these effects may differ by genre or listener preference.

7. CONCLUSION

We present RISE, a novel system that adapts music to align with exercise phases. A user study involving 12 participants revealed that, despite initial skepticism, most users appreciated the alignment and preferred it over nonadaptive music for their workouts. Our work represents a step towards expanding the design space of adaptive music, making tailored music experiences, once limited to video games and precomposed soundtracks, applicable to real-world scenarios like workouts.

8. REFERENCES

- [1] P. C. Terry, C. I. Karageorghis, M. L. Curran, O. V. Martin, and R. L. Parsons-Smith, "Effects of music in exercise and sport: A meta-analytic review." *Psychological bulletin*, vol. 146, no. 2, p. 91, 2020.
- [2] S. Mitroff, "Hitting the pavement with spotify running (hands-on) - cnet," <https://www.cnet.com/tech/services-and-software/hitting-the-pavement-with-spotify-running-hands-on/>, 2015.
- [3] Apple, "Apple fitness+ - apple," <https://www.apple.com/apple-fitness-plus/>, 2024.
- [4] P. Knees, M. Schedl, and M. Goto, "Intelligent user interfaces for music discovery: The past 20 years and what's to come." in *Proceedings of the 20th International Society for Music Information Retrieval Conference, ISMIR 2019, Delft, The Netherlands, November 4-8, 2019*, 2019, pp. 44–53.
- [5] M. Goto, "Grand challenges in music information research," in *Dagstuhl Follow-Ups: Multimodal Music Processing*, M. Müller, M. Goto, and M. Schedl, Eds. Dagstuhl Publishing, 2012, pp. 217–225.
- [6] M. Schedl and A. Flexer, "Putting the user in the center of music information retrieval." in *Proceedings of the 13th International Society for Music Information Retrieval Conference, ISMIR 2012, Mosteiro S.Bento Da Vitória, Porto, Portugal, October 8-12, 2012*, 2012, pp. 385–390.
- [7] M. Kari, T. Grosse-Puppenthal, A. Jagaciak, D. Bethge, R. Schütte, and C. Holz, "Soundsride: Affordance-synchronized music mixing for in-car audio augmented reality," in *The 34th Annual ACM Symposium on User Interface Software and Technology*, 2021, pp. 118–133.
- [8] L. Baltrunas, M. Kaminskas, B. Ludwig, O. Moling, F. Ricci, A. Aydin, K.-H. Lüke, and R. Schwaiger, "In-carmusic: Context-aware music recommendations in a car," in *E-Commerce and Web Technologies: 12th International Conference, EC-Web 2011, Toulouse, France, August 30-September 1, 2011. Proceedings 12*. Springer, 2011, pp. 89–100.
- [9] A. Wang, Y. F. Cheng, and D. Lindlbauer, "Maringba: Music-adaptive ringtones for blended audio notification delivery," in *Proceedings of the CHI Conference on Human Factors in Computing Systems, CHI 2024, Honolulu, HI, USA, May 11-16, 2024*. New York, NY, USA: Association for Computing Machinery, 2024.
- [10] A. Wang, D. Lindlbauer, and C. Donahue, "Towards music-aware virtual assistants," in *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology, UIST 2024, Pittsburgh, PA, USA, October 13-16, 2024*. New York, NY, USA: Association for Computing Machinery, 2024.
- [11] J. Shriram, M. Tapaswi, and V. Alluri, "Sonus texere! automated dense soundtrack construction for books using movie adaptations," in *Proceedings of the 23rd International Society for Music Information Retrieval Conference, ISMIR 2022, Bengaluru, India, December 4-8, 2022*, 2022, pp. 535–542.
- [12] S. Rubin, F. Berthouzoz, G. Mysore, W. Li, and M. Agrawala, "Underscore: musical underlays for audio stories," in *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology*, ser. UIST '12. New York, NY, USA: Association for Computing Machinery, 2012, p. 359–366.
- [13] N. Masahiro, H. Takaesu, H. Demachi, M. Oono, and H. Saito, "Development of an automatic music selection system based on runner's step frequency," in *ISMIR 2008, 9th International Conference on Music Information Retrieval, Drexel University, Philadelphia, PA, USA, September 14-18, 2008*, 2008, pp. 193–198.
- [14] G. T. Elliott and B. Tomlinson, "Personalsoundtrack: context-aware playlists that adapt to user pace," in *CHI'06 extended abstracts on Human factors in computing systems*, 2006, pp. 736–741.
- [15] B. Moens, L. van Noorden, and M. Leman, "D-jogger: Syncing music with walking," in *7th Sound and music computing Conference*. Universidad Pompeu Fabra, 2010, pp. 451–456.
- [16] J. Hockman, M. M. Wanderley, and I. Fujinaga, "Real-time phase vocoder manipulation by runner's pace." in *NIME*, 2009, pp. 90–93.
- [17] N. Oliver and L. Kreger-Stickles, "Papa: Physiology and purpose-aware automatic playlist generation." in *ISMIR 2006, 7th International Conference on Music Information Retrieval*, 2006, pp. 250–253.
- [18] B. van der Vlist, C. Bartneck, and S. Mäueler, "mobeat: Using interactive music to guide and motivate users during aerobic exercising," *Applied psychophysiology and biofeedback*, vol. 36, pp. 135–145, 2011.
- [19] Y. Chen, C.-C. Chen, L.-C. Tang, and W.-H. Chieng, "Enhancing running exercise with iot, blockchain, and heart rate adaptive running music," *IEEE Access*, 2024.
- [20] C. I. Karageorghis and D. Holland, "Music in the exercise domain: A review and synthesis (part ii)," *International Review of Sport and Exercise Psychology*, vol. 5, no. 1, pp. 67–84, 2012.
- [21] A. Turrell, A. R. Halpern, and A.-H. Javadi, "When tension is exciting: an electroencephalogram exploration of excitement in music," *bioRxiv*, p. 637983, 2019.
- [22] D.-L. Priest and C. I. Karageorghis, "A qualitative investigation into the characteristics and effects of music accompanying exercise," *European physical education review*, vol. 14, no. 3, pp. 347–366, 2008.

- [23] T. Kim and J. Nam, “All-in-one metrical and functional structure analysis with neighborhood attentions on demixed audio,” in *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 2023.
- [24] R. Hennequin, A. Khlif, F. Voituret, and M. Moussallam, “Spleeter: a fast and efficient music source separation tool with pre-trained models,” *Journal of Open Source Software*, vol. 5, no. 50, p. 2154, 2020.
- [25] C. Plachouras and M. Miron, “Music rearrangement using hierarchical segmentation,” in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [26] C.-W. Li and C.-G. Tsai, “The presence of drum and bass modulates responses in the auditory dorsal pathway and mirror-related regions to pop songs,” *Neuroscience*, vol. 562, pp. 24–32, 2024.
- [27] G. Madison, “Experiencing groove induced by music: consistency and phenomenology,” *Music perception*, vol. 24, no. 2, pp. 201–208, 2006.
- [28] C. J. Steinmetz and J. D. Reiss, “pyloudnorm: A simple yet flexible loudness meter in python,” in *150th AES Convention*, 2021.
- [29] Adobe, “Remix in premiere pro,” <https://helpx.adobe.com/premiere-pro/using/remix-audio-in-premiere-pro.html>, 2024.
- [30] W. M. W. Timer, “Interval timer with music | 40 sec rounds 30 sec rest | mix 107,” YouTube video, 2021, accessed: September 1, 2024. [Online]. Available: https://www.youtube.com/watch?v=lnBOQnc_p-E
- [31] R. E. Boyatzis, *Transforming qualitative information: Thematic analysis and code development*. Sage, 1998.
- [32] D. Freelon, “Recal2: Reliability for 2 coders,” <http://dfreelon.org/utis/recalfront/recal2/>, 2010.