

HeartSong: Beat-Synchronized Haptic Feedback for Music Listening and Social Presence

Prithiv Premkumar, Shuyuan Liu, Ryan Clark, and Bruce N. Walker

Georgia Institute of Technology

Atlanta, Georgia, USA

ppremkumar3@gatech.edu, sliu828@gatech.edu, rclark309@gatech.edu, bruce.walker@psych.gatech.edu

Abstract—HeartSong is a system that automatically generates heartbeat-like haptic feedback from music and streaming media by mapping detected beat timing, tempo, and intensity to a synchronized tactile and visual pulse. The system runs in web browsers on mobile devices and is also implemented in WebXR, supporting individual listening and shared experiences over the Internet in which multiple geographically and temporally distant users in the same virtual environment feel the same heartbeat pattern aligned with the music. HeartSong leverages heartbeat entrainment principles so that the haptics promote natural synchronization between the user’s perceived pulse and bodily timing with the audio beat, thereby strengthening temporal awareness and supporting more intuitive regulation of effort and emotion during listening. The amplitude of the heartbeat is dependent on both the volume and the semantic context, increasing during louder or more emotionally salient sections and decreasing during quieter passages to produce a lifelike and expressive tactile profile. This work describes the system design, the sync-generation pipeline, and the runtime playback engine, and discusses application domains including exercise pacing with target heart rate guidance, brief rhythm training for cadence and timing practice, and collective embodied listening over the Internet where synchronized tactile feedback reduces social disconnectedness and enhances social presence among geographically distant listeners sharing streaming media.

Index Terms—music haptics, virtual reality, entrainment, accessibility, tactile feedback, rhythm, embodied interaction

I. INTRODUCTION

Music perception is inherently embodied: listeners often internalize rhythm through movement, anticipation, and physiological responses [1]. However, most playback systems remain primarily auditory, offering limited support for tactile or visual cues that reinforce temporal structure and shared attention [2]. This limitation is particularly evident in web-based and virtual reality (VR) environments, where synchronized bodily engagement is largely absent.

HeartSong addresses this gap by generating heartbeat-like haptic feedback from music, mapping beat timing, tempo, and intensity to synchronized tactile and visual pulses. The system operates in both browser-based and WebXR environments, supporting solo and shared listening experiences in which users can have identical, time-aligned haptic patterns. By representing musical structure as a heartbeat-like signal, HeartSong provides an intuitive mapping that strengthens temporal awareness and listener engagement.

This work-in-progress paper makes three contributions. First, it presents a deterministic pipeline for generating

heartbeat-like haptics from musical features and creating variations in intensity based on semantic and volume features. Second, it introduces a cross-platform runtime supporting both mobile web and WebXR playback. Third, it explores how synchronized haptic feedback can enhance embodied listening, accessibility, and shared musical experiences in networked environments.

II. RELATED WORK

HeartSong builds on work in musical haptics, affective haptics, and multisensory music listening in virtual environments. Early studies showed that vibrotactile stimulation can augment low-frequency auditory perception and influence pitch discrimination thresholds, establishing that touch can contribute meaningfully to fine-grained auditory tasks [3]. Subsequent work demonstrated that rhythm perception is shared between audio and haptics, with listeners able to detect and interpret rhythmic changes similarly across the two modalities [4]. Reviews of audio–tactile rendering and related systems describe technologies that translate music into vibrotactile signals for both enhancement and assistive listening [2]. Studies of audio–tactile congruence further show that matching tactile intensity and timing to auditory intensity and timing improves perceived enhancement, which motivates HeartSong’s use of beat-synchronous heartbeat pulses rather than arbitrary vibration patterns [5].

Wearable and ambient musical haptics systems have explored how touch can augment live performances and listening in more social settings. Musical haptic wearables for audiences introduce networked armbands and gilets that deliver vibrotactile stimuli derived from performers’ actions, enriching live concerts and encouraging participatory experiences [6]. FeelMusic likewise shows that carefully designed audio–tactile mappings can increase emotional engagement with music in seated listening contexts, providing a blueprint for affect-aware tactile rendering of musical features [7]. For accessibility, systems such as the Haptic Chair and visual music displays have been shown to significantly enhance the musical experience of Deaf users by translating sound into whole-body vibration and visual cues, with many participants preferring vibrotactile-based listening modes over audio alone [8]. In VR, multisensory music listening in affective environments shows that combining audio, visual, and tactile cues can shape affective responses, groove, and presence, while head-based




Fig. 1. Example of HeartSong with visual entrainment with glow and syncing heart animation

spatial haptics can substitute or complement audio cues for sound localization [9], [10]. HeartSong adopts similar design values by treating touch as a primary channel for musical information rather than a secondary effect.

Affective haptics research shows that vibrotactile stimulation shapes emotional interpretation of audiovisual content by evoking calmness, excitement, or tension, and EEG studies confirm that vibrotactile cues modulate attention and emotional processing [11]–[13]. Research on heart rate sharing demonstrates that perceiving another’s heartbeat increases empathy, social closeness, and emotional presence [14], [15], and haptic heartbeat representations can shape emotional experience [16]. Entrainment refers to biological oscillators synchronizing with external rhythmic signals. Research documents heart rate convergence during close communication correlating with empathy, and false physiological feedback indicating elevated heart rate increases actual heart rate and perceived tension [16]–[18]. Although we do not test entrainment here, HeartSong’s heartbeat-like haptics and shared rhythmic pulses are motivated by this possibility of using tactile heartbeat cues to nudge users toward different arousal states and to strengthen social connection in both mobile and VR contexts.

III. SYSTEM OVERVIEW

HeartSong uses a two-stage architecture with an offline sync generation pipeline and a runtime playback engine. The offline stage processes each music track into a deterministic `.music.srt` file that stores beat timestamps, BPM metadata, and intensity values. The runtime stage consumes that file and drives synchronized audio, haptic, and visual feedback across both mobile web and WebXR environments. This design follows the Heartbeat SRT approach used in Resonant, where a textual timing representation coordinates heartbeat-driven feedback with media playback [19].

A. Offline Sync Generation and Song Selection

The sync generator analyzes each input track and emits an ordered sequence of beat events into a `.music.srt` file as

shown in Fig. 2. Each entry includes a timestamp, optional end time, intensity score, BPM estimate, and beat index, which together define the timing and strength of the heartbeat pulse at playback time. Because the output is deterministic, the same song always produces the same sync file when run with the same settings, making the file a stable master schedule that can be reused across devices and sessions [19].

Beat timing is estimated from the song’s rhythmic structure and aligned to a stable tempo grid. The pipeline first detects candidate beat peaks from audio energy and spectral change, then refines those peaks against a tempo estimate so the final sync aligns closely with the actual song. Haptic intensity is computed from audio energy and optionally modified by timestamped semantic salience from a context file. This allows stronger musical or lyrical moments to produce stronger pulses while quieter passages remain perceptible but less intense.

The intensity model combines these signals into a single value that drives runtime rendering:

$$I = \alpha S + (1 - \alpha)V \quad (1)$$

where S is semantic or contextual salience, V is audio energy, and α controls the weighting between them. When $\alpha = 0$, intensity is volume-driven; when $\alpha = 1$, it is semantics-driven. In practice, this lets the same sync format support both straightforward beat-driven playback and more expressive, emotionally informed mappings.

The prototype was tested with a diverse set of tracks to test different rhythmic structures and use cases. *Bohemian Rhapsody* (Queen, 72 BPM) was included specifically as a stress test for variable tempo handling: the song changes tempo across distinct sections, starting slowly in the intro and ballad passages, then moving into a guitar-driven section that effectively doubles the rhythmic intensity before returning to the original pace. Our sync generator must track these tempo shifts and preserve musical phrasing rather than forcing the song into a rigid constant BPM grid. In contrast, rhythmically stable tracks such as *Another One Bites the Dust* (Queen, 110 BPM), *Billie Jean* (Michael Jackson, 117 BPM), *Around the World* (Daft Punk, 121 BPM), and *Happy* (Pharrell Williams, 160 BPM) provide baseline evaluation under steady beat conditions. Additional tracks with use cases in exercise and pacing include *Come Together* (The Beatles, 84 BPM), *Back in Black* (AC/DC, 94 BPM), *Stayin’ Alive* (Bee Gees, 104 BPM), and *Eye of the Tiger* (Survivor, 109 BPM).

B. Mobile Web Player

As shown in Fig. 1, the mobile version is a browser-based player that runs on standard phones and desktops. It loads the audio track and the generated `.music.srt` file, then schedules heartbeat pulses at the exact beat timestamps defined in the sync file. On supported phones, vibration is triggered through the browser vibration API, which provides patterned device vibration directly from the web runtime [20]; if vibration is unavailable, the interface still renders the visual layer so the beat remains visible and the experience remains functional.

```

TITLE:Queen - Another One Bites the Dust
SOURCE:offline_audio_beat_analysis
AUDIO_DURATION_SEC:222.772
ESTIMATED_BPM:110.00
BEAT_COUNT:396
FIRST_BEAT:00:00:07,513
LAST_BEAT:00:03:42,392

```

```

1
00:00:07,513 --> 00:00:08,062
BEAT_START:00:00:07,513
BEAT_END:00:00:08,062
BEAT_INDEX:1
INTENSITY:0.87
PEAK_SCORE:0.0828
SOURCE:predetermined_audio_analysis

```

```

2
00:00:08,062 --> 00:00:08,592
BEAT_START:00:00:08,062
BEAT_END:00:00:08,592
BEAT_INDEX:2
INTENSITY:0.76
PEAK_SCORE:0.0645
SOURCE:predetermined_audio_analysis

```

Fig. 2. Excerpt from a generated `.music.srt` file for *Another One Bites the Dust* by Queen.

The mobile player also includes a pulsing orb for visual entrainment. The orb expands, contracts, and changes brightness in synchrony with the beat stream, giving the user a simple visual anchor for rhythm. The effect is intentionally restrained and is meant to feel like a disco-ball-style pulse in the center of the interface rather than a decorative animation. The same intensity value that drives the haptic pulse also controls the orb size and glow so that tactile and visual cues stay aligned.

The browser player supports both solo and shared listening. In solo mode, one user receives the complete haptic pattern locally. In shared mode, multiple users can join the same session and experience the same timing pattern at the same musical moments, allowing the mobile experience to function both as a personal music companion and as a lightweight shared listening space. The supplementary videos contain a video of the mobile application in action.

C. Shared Sessions and WebXR

HeartSong supports shared listening through a lightweight peer-to-peer connection layer built on PeerJS. For browser-based sessions, users can connect by simply entering a room code (a short alphanumeric string) that PeerJS uses to coordinate a shared signaling channel. Once one user creates a room and others join with the same code, the PeerJS layer establishes WebRTC connections between participants and aligns playback state, sync metadata, and session timing across devices. The peer layer is used only for lightweight coordination such as shared start time, track selection, and session alignment; the music sync file still drives local haptic and visual rendering on each device, keeping timing deterministic on the client side.




Fig. 3. Example of WebXR version with movable “heart.” The music and haptics originate from the sphere and as the user moves closer or further away the intensity of the haptics changes. The user can also pick up the sphere and move it around and the sphere affects the background flashing as well.

As shown in Fig. 3, the WebXR version using ThreeJS mirrors the mobile implementation but replaces phone vibration with controller haptics. The same `.music.srt` file is loaded at runtime, and each beat event is mapped to a short controller pulse through the WebXR input and haptic interfaces, keeping the timing logic identical across platforms while moving the tactile feedback into the headset environment [21], [22]. The same pulsing orb is used in VR as a floating visual object that pulses with the beat, and shared listening works identically: multiple headset users join using the same room code, receive the same beat schedule, and feel the same controller pulses at the same musical moments.

Both runtimes use the sync file as the master schedule. Audio playback provides the transport clock, and haptic and visual events are dispatched relative to that clock so that the beat, pulse, and orb remain aligned. Because the timing is precomputed offline, the system avoids ad hoc beat estimation during playback and reduces drift across devices, making the mobile and WebXR versions easier to evaluate using the same musical inputs. Each beat event is rendered as a short pulse followed by a pause, producing the familiar lub-dub feel without requiring extra user configuration. Currently, the WebXR and mobile versions are separate, but future implementations will support cross-platform sessions using the same room-code-based infrastructure. The supplementary videos contain a video of the WebXR application in action.

D. Bidirectional Heartbeat Mapping

HeartSong works in both directions: forward, it maps music to heartbeat-like haptics; reverse, it uses a desired heartbeat value to select music whose rhythmic profile matches that target. In our prototype, the system chooses the song closest to the user’s current BPM, but only updates the target if the estimated heart rate changes consistently over at least one minute to avoid constant song switching; users can also manually request a refresh.

This bidirectional design is especially useful in XR, where cardiac signals can be estimated from headset motion by fil-

tering and normalizing accelerometer and gyroscope streams, then applying frequency analysis to extract a dominant rhythm in beats per minute [23]. This allows HeartSong to act as both a music-to-motion interface (rendering heartbeat haptics) and a motion-to-music interface (selecting songs based on estimated or target heartbeat), treating rhythm as a bidirectional control signal between motion, physiological state, and sound.

IV. POTENTIAL APPLICATIONS

HeartSong may support applications beyond entertainment where rhythm and bodily timing matter.

Exercise and pacing: Runners and other athletes often use music to stay on pace and regulate effort. A synchronized heartbeat cue could strengthen cadence awareness and make pacing more intuitive, especially in noisy or distracting environments. The tactile pulse may also help users feel the drive of the song even when audio volume must stay low, and the bidirectional heartbeat-to-music mapping allows users to choose songs whose tempo matches a desired effort level.

Emotion and entrainment: Music shapes arousal, tension, and release. HeartSong could make these changes more physically perceptible by matching haptic intensity to musical structure so that climaxes, drops, and quiet passages feel different in the body. In the shared WebXR version, these cues can become a group event where several people experience the same rhythmic rise and fall together, supporting collective entrainment and social connection around music listening.

CPR training: Rhythm is critical in CPR, where chest compressions must stay within a target rate. Prior work shows that using songs such as “Stayin’ Alive” as a mental metronome helps providers maintain compression rates near guideline recommendations [24]. HeartSong can extend this idea by providing a tactile pacing aid that tracks the beat of CPR-relevant songs, helping learners internalize compression timing through feel as well as sound and making musical pacing cues easier to follow in noisy or stressful settings.

Accessibility: For Deaf and hard-of-hearing users, or for anyone who wants a tactile route into music, HeartSong can make rhythm more accessible by adding a parallel sensory channel for beat and structure. Haptic chairs, sleeves, and head-based systems have already shown that vibrotactile feedback can enrich musical experiences and support sound localization for DHH listeners [8], [10], [25]. HeartSong’s heartbeat pulses and pulsing orb can be used in similar ways across mobile and VR, supporting inclusive social participation by letting multiple users experience the same tactile music cues together.

Performers and ensemble coordination: Groups such as musical ensembles or dance pairs (for example tango dancers) could use HeartSong to feel tempo, beat, and intensity in addition to hearing the music. Shared haptic cues could provide a common rhythmic reference that helps performers align timing and dynamics, especially in situations where auditory cues are limited or when silent practice is needed.

Large-scale shared experiences: In stadium or venue settings, HeartSong could turn a song into a shared embodied

event where every audience member feels the same heartbeat pattern at the same time. This coordinated tactile layer could deepen feelings of togetherness and collective presence, making musical moments more immersive and socially connected even when listeners are physically distant or in different virtual spaces.

V. CONCLUSION

HeartSong automatically converts music into heartbeat-like haptics using deterministic sync generation and runtime playback across mobile web and WebXR. The system supports both solo listening and shared sessions, and it adds a pulsing visual orb plus synchronized haptic feedback to make rhythm feel more embodied. By using a single sync file to drive music, touch, and visual entrainment, the system expands music playback into a more physical, social, and expressive experience.

The current prototype also suggests that HeartSong can work in reverse, where a heartbeat value can guide music selection and pacing. This makes the system interesting not only for music listening but also for motion to music interaction, exercise support, and adaptive pacing. In XR, the possibility of estimating heartbeat from headset motion opens an even broader design space for systems that infer state from movement and then respond with rhythm. This creates a novel loop between motion, body state, and music.

Future work will focus on collecting user data and validating the current prototype. We plan to test the system across a broader range of genres, tempos, and structural patterns so we can understand where the heartbeat mapping works best and where it needs refinement. We also want to compare solo listening, shared web listening, and WebXR listening more directly and test across a diverse audience to see if there are actually benefits in accessibility, emotion, and entrainment across users. Additional studies should evaluate whether the pulsing orb improves visual entrainment, whether controller haptics improve XR immersion, and whether users prefer fixed heartbeat pacing or flexible target-based pacing. More controlled testing will also help determine how well the system supports exercise and CPR-related timing tasks along with collaborative music experiences where many people listen and feel the same haptics at once.

REFERENCES

- [1] E. Van Dyck, B. Moens, J. Buhmann, M. Demey, E. Coorevits, S. Dalla Bella, and M. Leman, “Spontaneous entrainment of running cadence to music tempo,” *Sports medicine-open*, vol. 1, no. 1, p. 15, 2015.
- [2] B. Remache-Vinueza, A. Trujillo-León, M. Zapata, F. Sarmiento-Ortiz, and F. Vidal-Verdú, “Audio-tactile rendering: A review on technology and methods to convey musical information through the sense of touch,” *Sensors*, vol. 21, no. 19, p. 6575, 2021.
- [3] G. W. Young, D. Murphy, and J. Weeter, “Haptics in music: The effects of vibrotactile stimulus in low frequency auditory difference detection tasks,” *IEEE Transactions on Haptics*, vol. 10, no. 1, pp. 135–139, 2017.
- [4] C. Bernard, J. Monnoyer, M. Wiertelowski, and S. Ystad, “Rhythm perception is shared between audio and haptics,” *Scientific Reports*, vol. 12, no. 1, p. 4188, 2022.

- [5] S. C. Aker, H. Innes-Brown, K. F. Faulkner, M. Vatti, and J. Marozeau, "Effect of audio-tactile congruence on vibrotactile music enhancement," *Journal of the Acoustical Society of America*, vol. 152, no. 6, pp. 3396–3409, 2022.
- [6] L. Turchet, T. West, and M. M. Wanderley, "Touching the audience: musical haptic wearables for augmented and participatory live music performances," *Personal Ubiquitous Comput.*, vol. 25, no. 4, p. 749–769, Mar. 2020. [Online]. Available: <https://doi.org/10.1007/s00779-020-01395-2>
- [7] A. Haynes, J. Lawry, C. Kent, and J. Rossiter, "Feelmusic: Enriching our emotive experience of music through audio-tactile mappings," *Multimodal Technologies and Interaction*, vol. 5, no. 6, p. 29, 2021.
- [8] S. Nanayakkara, E. Taylor, L. Wyse, and S. H. Ong, "An enhanced musical experience for the deaf: design and evaluation of a music display and a haptic chair," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ser. CHI '09. New York, NY, USA: Association for Computing Machinery, 2009, p. 337–346. [Online]. Available: <https://doi.org/10.1145/1518701.1518756>
- [9] L. Turchet, C. Marotta, and A. Boem, "Multisensory music listening in affective virtual environments," *IEEE Transactions on Affective Computing*, vol. 16, no. 4, pp. 3304–3315, 2025.
- [10] P. K. Chelladurai, Z. Li, M. Weber, T. Oh, and R. L. Peiris, "Soundhapticvr: Head-based spatial haptic feedback for accessible sounds in virtual reality for deaf and hard of hearing users," in *Proceedings of the 26th International ACM SIGACCESS Conference on Computers and Accessibility*, ser. ASSETS '24. New York, NY, USA: Association for Computing Machinery, 2024. [Online]. Available: <https://doi.org/10.1145/3663548.3675639>
- [11] A. Mazzoni and N. Bryan-Kinns, "Mood glove: A haptic wearable prototype system to enhance mood music in film," *Entertainment Computing*, vol. 17, pp. 9–17, 2016.
- [12] V. Cerdán-Martínez, García-López, P. Revuelta-Sanz, T. Ortiz, and R. Vergaz, "Haptic stimulation during the viewing of a film: an EEG-based study," *Multimedia Tools and Applications*, vol. 83, no. 26, pp. 67 673–67 686, Jan. 2024. [Online]. Available: <https://link.springer.com/10.1007/s11042-024-18218-8>
- [13] M. Xu, T. Tachibana, N. Suzuki, E. Hoshino, Y. Terasawa, N. Miki, and Y. Minagawa, "The effect of haptic stimulation simulating heartbeats on the regulation of physiological responses and prosocial behavior under stress: The influence of interoceptive accuracy," *Biological Psychology*, vol. 164, p. 108172, 2021.
- [14] R. M. Winters, B. N. Walker, and G. Leslie, "Can you hear my heart-beat?: Hearing an expressive biosignal elicits empathy," in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM, 2021.
- [15] Y. Ju, D. Zheng, D. Hynds, G. Chernyshov, K. Kunze, and K. Minamizawa, "Haptic empathy: Conveying emotional meaning through vibrotactile feedback," in *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, ser. CHI EA '21. New York, NY, USA: Association for Computing Machinery, 2021. [Online]. Available: <https://doi.org/10.1145/3411763.3451640>
- [16] J. Costa, A. T. Adams, M. F. Jung, F. Guimbretière, and T. Choudhury, "Emotioncheck: leveraging bodily signals and false feedback to regulate our emotions," in *Proceedings of the 2016 ACM international joint conference on pervasive and ubiquitous computing*, 2016, pp. 758–769.
- [17] A. T. Adams, J. Costa, M. F. Jung, and T. Choudhury, "Mindless computing: designing technologies to subtly influence behavior," in *Proceedings of the 2015 ACM international joint conference on pervasive and ubiquitous computing*, 2015, pp. 719–730.
- [18] J. Costa, F. Guimbretière, M. F. Jung, and T. Choudhury, "Boostmeup: Improving cognitive performance in the moment by unobtrusively regulating emotions with a smartwatch," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 3, no. 2, pp. 1–23, 2019.
- [19] K. Shibata, P. Premkumar, I. C. Valderas, R. Clark, and B. N. Walker, "Resonant: A web-based extended reality platform for heartbeat sharing via headset ballistocardiography and mobile remote photoplethysmography," in *Proceedings of the 2026 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (IEEEVR 2026 VRW)*, Daegu, Korea, March 21–23 2026, pp. 186–189.
- [20] Mozilla, "Vibration api," https://developer.mozilla.org/en-US/docs/Web/API/Vibration_API, accessed 2026-04-09.
- [21] MDN Web Docs. (2026) Webxr device api. Mozilla Foundation. Accessed 2026-07-07. [Online]. Available: https://developer.mozilla.org/en-US/docs/Web/API/WebXR_Device_API
- [22] three.js contributors. (2026) Webxr haptics example in three.js. three.js. Accessed 2026-07-07. [Online]. Available: https://threejs.org/examples/webxr_xr_haptics.html
- [23] I. Roura, J. Raposo, C. Abreu, M. R. d. Santos, P. Campos, C. Mas-saroni, H. Silva, and D. Lopes, "Ghost in the vr shell: Capturing spectral cardio-respiratory rates from subtle vr device movements," in *Proceedings of the 31st ACM Symposium on Virtual Reality Software and Technology (VRST '25)*. ACM, 2025, pp. 1–11.
- [24] J. W. Hafner, J. L. Sturgell, D. L. Matlock, E. G. Bockewitz, and L. T. Barker, "'stayin' alive": A novel mental metronome to maintain compression rates in simulated cardiac arrests," *Journal of Emergency Medicine*, vol. 43, no. 5, pp. e373–e377, 2012.
- [25] L. Turchet, "Musical haptics for the listener, part i: Tactile music enhancement systems for hearing and dhh users," *IEEE Transactions on Haptics*, 2026.