Tuning Music Education: AI-Powered Personalization in Learning Music

Mayank Sanganeria
Independent Researcher
New York City, USA
mayank@ccrma.stanford.edu

Rohan Gala
Independent Researcher
Seattle, USA
rhngla@gmail.com

Abstract

Recent AI-driven step-function advances in several longstanding problems in music technology are opening up new avenues to create the next generation of music education tools. Creating personalized, engaging, and effective learning experiences is a continuously evolving challenge in music education. Here we present two case studies using such advances in music technology to address these challenges. In our first case study we showcase an application that uses Automatic Chord Recognition to generate personalized exercises from audio tracks, connecting traditional ear training with real-world musical contexts. In the second case study we prototype adaptive piano method books that use Automatic Music Transcription to generate exercises at different skill levels while retaining a close connection to musical interests. These applications demonstrate how recent AI developments can democratize access to high-quality music education and promote rich interaction with music in the age of generative AI. We hope this work inspires other efforts in the community, aimed at removing barriers to access to high-quality music education and fostering human participation in musical expression.

1 Introduction

Music holds a unique power to evoke emotions, foster creativity, and build cross-cultural connections. However, traditional music education often adopts a one-size-fits-all approach, emphasizing classical repertoires or mainstream genres that may not resonate with the diverse musical preferences of modern students. Rigid pedagogical methodology risks alienating learners whose tastes lie beyond the confines of the selected curriculum, hindering engagement and enthusiasm. Advances in AI research can be used to offer transformative solutions in this regard – personalized music education tailored to each student's unique musical identity. By curating lessons around an individual's favorite artists, genres, and songs, we can create inclusive environments where students can fulfill their musical potential.

AI music generation has garnered significant attention and made immense progress in recent years [1, 2, 3, 4]. We share the opinion that the practice of music is valuable far beyond just the final output [5, 6]. Using generative AI beyond music synthesis, towards analysis of musical concepts from audio data to personalize and support music education remains an under-explored and exciting area for the community.

In the following sections, we first explore the importance of leveraging students' individual musical preferences to enhance motivation and improve learning outcomes. We then present two case studies that showcase the potential of AI-driven music technology in education. One is an ear training application (app) that generates customized exercises based on students' favorite songs. The other is an AI-powered piano method book prototype that adapts to students' skill levels and musical interests.

38th Conference on Neural Information Processing Systems (NeurIPS 2024) Creative AI Track.

2 Background

Need for personalization

Conventional music curricula tend to prioritize mastery of Western classical music, folk traditions, and works of renowned composers. While this canonical approach provides a solid foundation, it can fail to captivate students whose musical interests lie in contemporary, non-Western, or niche genres [7]. The lack of representation of personal tastes in educational content can disengage learners, limiting their motivation and ability to connect with the material [8]. As an example, a student drawn into picking up the practice of music through their interest in electronic dance music may find it challenging to engage with a curriculum focused on Baroque-era compositions [9].

Extensive research highlights the benefits of personalized learning experiences [10]. Self determination theory suggests that relatedness and autonomy are key components to sustain intrinsic motivation in the pursuit of achievement and performance [11]. Incorporating the student's preferred music and designing ways to interact with it can enhance agency and connection, supporting the educational endeavor [12, 13, 14].

Moreover, the emotional resonance of music is thought to play a pivotal role in its educational efficacy [15]. Personalization in music education directly impacts skill acquisition, as engaged students are more likely to practice consistently and persevere through challenges [16, 17].

For an individual student, the limitations of traditional music education can be circumvented by working with an exceptional music teacher. Such teachers could incorporate the student's preferred genres with a deep understanding those genres, and with the dedication to create a highly personalized learning experience [18, 19, 20]. From the teacher's perspective, creating such experiences requires substantial time and effort towards transcribing songs, arranging them to suit the student's current skill level, and designing targeted exercises to develop specific abilities, such as ear training and piano technique. Tools can help ease this burden and enable better teaching of personalized content; even something as simple as automated harmonic analysis of symbolic music, like ChordNamer [1]. Given the demands on music teachers, this is uncommon and inaccessible to the majority of music students [21].

AI powered solutions

We can build digital systems to serve as adaptive learning environments that cater to each student's unique musical interests and learning needs, at scale [22]. Through the analysis of audio tracks in each student's listening history, AI could enable creation of dynamic lesson plans and exercises tailored to individual interests as has been explored in language learning settings [23]. Moreover, these adaptive learning systems could monitor student progress, adjusting the content and difficulty of lessons to maintain challenge and interest through the learning process.

The particular technologies we use for our case studies are Automatic Chord Recognition (ACR), beat detection, and Automatic Music Transcription (AMT).

Automatic Chord Recognition (ACR) has evolved considerably from early knowledge-based systems towards data-driven machine learning approaches [24]. Current tools like ChordAI [2], Chordify [3], and libraries like crema [4], as well as data-sets like Chord-Annotations [5] are some of the resources that are available at present. Beat detection is often performed as a step on the way to chord recognition, although dedicated open-source tools such as BeatNet [25] also exist.

Automatic Music Transcription (AMT) is the process of converting audio recordings into symbolic musical representations such as sheet music [26]. Researchers have explored options ranging from signal processing techniques to machine learning algorithms [26] for AMT. In recent years, deep learning models have demonstrated remarkable progress in accurately transcribing polyphonic music [27, 28, 29], and may soon be application-ready. Some of the main tools available at this time are

[1] [https://github.com/e7mac/chord_namer]
[2] [https://chordai.net]
[3] [https://chordify.net]
[4] [https://github.com/bmcfee/crema]
[5] [https://github.com/tmc323/Chord-Annotations]

Piano Cover Generation (PiCoGen) [30, 31], Pop2Piano [32], piano-transcription [6] [29], Piano2Notes [7], and Audio to sheet music converter [8].

While it is clear that the underlying music technologies required for music education systems have been researched for some time now, their practical deployment has been limited by accuracy concerns. Students must be able to trust the system, which requires the underlying technologies to have very low error rates. With AI approaches these technologies are now approaching expert level performance and crossing over to a regime where they can effectively be deployed in educational contexts.

3 Case Studies

We now explore how to build such systems through two case studies. In the first study we describe our ear training app, which generates custom exercises based on students' favorite songs. In the second case study we present an approach to create adaptive and personalized piano method books.

Ear Training App

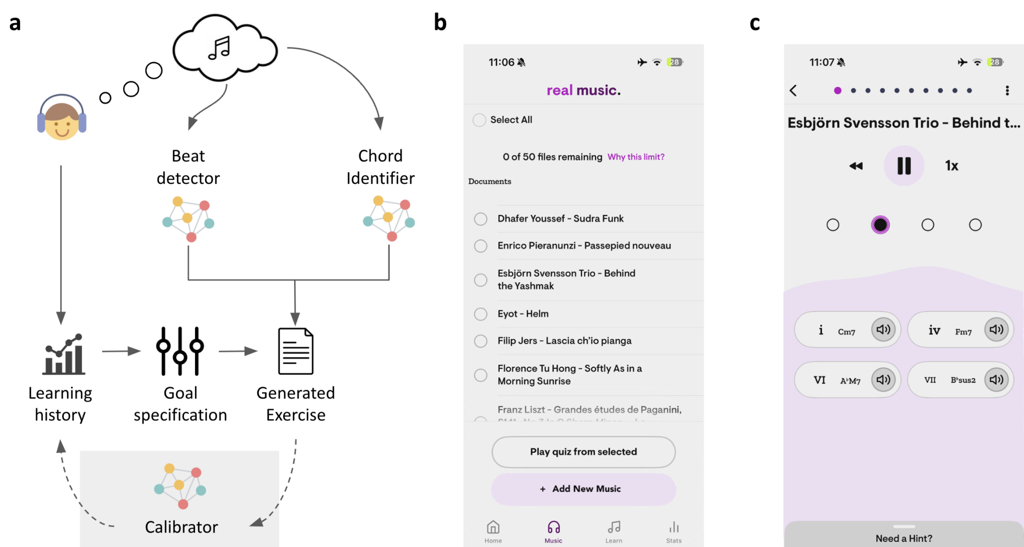

Figure 1: Overview of RealEarTrainer. **(a)** The student can select custom audio tracks. AI modules detect beats and identify chords within these tracks. The app generates personalized exercises using snippets sourced from the selected audio tracks. Calibrating exercises based on learning history, goal specification, and current performance could be achieved with another AI module. **(b)** RealEarTrainer interface to select preferred music from a list of available tracks. **(c)** During the exercise, the app plays a snippet from one of the selected tracks and prompts the student to identify the chords being played. The sound icon on each chord in the options plays a synthesized piano version that can be used by the student to make the harmonic content salient.

Traditional ear training tools have relied on the use of piano sounds or the student's primary instrument to facilitate the development of aural skills. In real-world musical contexts, students must contend with a myriad of textures, timbres, and parts distributed across multiple instruments. To address this challenge, some students resort to training without the aid of dedicated tools, relying solely on their ability to listen to music and reproduce it on their instrument. This approach, while valuable, can be time-consuming and may not always provide targeted feedback for improvement. We've found app-based tools, such as our personal favorite Chet [9], to be effective and helpful not only for building

[6][https://github.com/bytedance/piano_transcription]
[7][https://piano2notes.com]
[8][https://latouchemusicale.com/en/tools/audio-to-sheet-music-converter/]
[9][https://chetapp.io]

foundational ear training skills but also for bridging the gap between learning environments and real-world musical contexts.

We developed an ear training application that uses beat detection and ACR to analyze student-provided audio files, and generate a personalized ear training curriculum that features snippets and examples drawn directly from the student's preferred music. This app is called RealEarTrainer [`https://realeartrainer.com`] and is available for download for iOS. Figure 1a-c shows an overview of the app and interface, that we also describe below:

The student begins by selecting their preferred audio tracks within the app. AI modules identify chords and align them with the beats of each track. This analysis is used to generate a series of tailored exercises. In each exercise, the student listens to a short snippet sourced from one of their chosen tracks and attempts to identify the chords being played. The app optionally provides synthesized piano chords aligned with chord changes in the audio snippet. This feature is designed to highlight the harmonic content, assisting students until they can confidently identify chords from the original audio alone.

While the core experience of connecting ear training quizzes to the student's preferred music is already functional in the app, we recognize that there is a lot of room in refining how this content is presented to the student. In particular, tuning exercise difficulty based on past performance and specific goals (e.g. distinguishing between specific chord families) could be achieved by a calibrator module in future versions.

Personalized piano method book

Aspiring musicians often find their initial spark of inspiration in a beloved song or piece of music. This passion serves as a driving force behind their decision to learn an instrument, towards the goal of being able to play that piece. Often, a first step is to pick up a method book for the instrument.

For piano, method books have a rich pedagogical history and intended to teach essential theoretical knowledge and technical skills. However, they are limited from a personalization point of view. Consequently, students may find themselves with pieces and exercises that, while valuable from an educational perspective, lack the personal resonance and emotional connection that initially drew them to the instrument. This opens up the prospect of AI-powered piano method books that adapt to each student's musical interests and preferences. We think personalization can be achieved through two primary approaches.

First, AI can be employed to analyze songs chosen by the student and generate custom arrangements that align with their current skill level and learning objectives. This can be achieved through adapting and simplifying the original compositions. It could, for instance, remove complex ornamentations, simplify arpeggio patterns and left-hand accompaniment to block chords, among many other simplifications.

Second, AI can be used to create exercises that address technical demands specific to the student's chosen music. As opposed to providing a fixed suite of drills, a next generation method book would analyze chord progressions, rhythmic patterns, and melodic elements in the selected songs to construct exercises that directly build the skills necessary to play those specific pieces. For example, if a particular song features a complex syncopated rhythm in the left-hand accompaniment, the generated exercises target development of coordination and timing required to execute that specific rhythm accurately. These targeted exercises would be presented alongside the simplified arrangement of the song, allowing the student to develop the necessary skills in a focused and purposeful manner.

We demonstrate this concept with a short excerpt from Yann Tiersen's "Comptine D'un Autre Été L'après". We focus on measures 30 - 34 of the piece and prototype an AI system that can simplify the arrangement and generate targeted exercises for a beginner.

Our prototype simplifies the left hand part to block chords using the notes present in each measure, and removes the 16^{th} note ornaments in the right hand, Figure 2b. This maintains the overall structure and melody of the original piece while being more accessible for a beginner to play. This simplified arrangement serves as an achievable milestone on the path to mastering the original composition.

In addition to the simplified arrangement, our prototype generates targeted exercises to help the student develop specific skills required to play this excerpt. Figure 2c shows an example exercise

Figure 2: Re-imagined piano method books. **(a)** We obtain the score for bars 30 - 34 of Comptine D'un Autre Été L'après using Piano2Notes as our transcriber. **(b)** A mix of procedural and ACR AI modules is used to distill the piece by removing ornamentations and providing block chords that make the core idea easy to follow and play on the piano. **(c)** Modules to suggest scales over the block chord changes serve to build the related skills while maintaining a connection to the original piece.

focused on learning relevant scales over the chord changes. This exercise focuses on playing the scale of the piece, while transitioning between chord changes found in the original excerpt.

Further exercises could target specific skills such as playing the melodic rhythm with the right hand, mastering the rhythm interplay between both hands, and mastering the left-hand accompaniment pattern. Breaking down the original piece into manageable components supports the student's progress towards confidently learning the entire piece – and these can now be created using the original audio file!

4 Discussion

AI is poised to enable a paradigm shift in music education, away from rigid curricula prevalent in traditional approaches towards those where personalization take center stage. In our opinion, the impact of personalized approaches such as those showcased here for ear training and method books, extends far beyond simple convenience or novelty. By directly linking effort and practice to an improved ability to understand, appreciate, and play the music the student loves, such applications can support the learning process for skills that have much broader utility.

Access to skilled teachers is usually limited, and it tends to be expensive as well. Through applications like ours, students can receive tailored instruction at a fraction of the cost of private lessons. However, we'd like to note that AI-powered music education is not intended to replace human teachers. Rather, it can serve as a powerful tool to augment and enhance the work of music educators. AI systems can handle the time-consuming tasks of content creation and adaptation, freeing up teachers to focus on higher-level skills such as musical expression, creativity, and collaboration.

While our case studies demonstrate the potential of AI-powered personalization in music education, we acknowledge that comprehensive assessment of these approaches' effectiveness remains as future work. Rigorous evaluation through controlled studies comparing traditional and AI-enhanced learning methods, along with longitudinal studies tracking student progress and engagement over time, will be crucial to validate these approaches.

Here we have demonstrated an engaging ear training experience, and a conceptual approach to tailoring method books at various difficulty levels, both enabled by AI. These efforts are directed towards lowering barriers, increasing personalization, and we present it with a hope that it inspires other attempts in music education.

Acknowledgments

We are grateful to Stochastic Labs for their support and for fostering an innovative environment that has been instrumental in developing RealEarTrainer. Their commitment to exploring the intersection of AI and human creativity aligns with our vision for transforming music education.

We would like to express our sincere gratitude to Vivien Seguy, creator of ChordAI, for his remarkable on Automatic Chord Recognition (ACR), which has been crucial inspiration for these ideas.

References

[1] Koray Tahiroğlu, Shenran Wang, Eduard Tampu, and Jackie Lin. Deep learning with audio: An explorative syllabus for music composition and production. In *AIMC 2023 Proceedings: Conference on AI and Music Creativity 2023*, 2023. Conference on AI and Music Creativity.

[2] Jade Copet, Felix Kreuk, Itai Gat, Tal Remez, David Kant, Gabriel Synnaeve, Yossi Adi, and Alexandre Défossez. Simple and controllable music generation. *Advances in Neural Information Processing Systems*, 36, 2024.

[3] Prafulla Dhariwal, Heewoo Jun, Christine Payne, Jong Wook Kim, Alec Radford, and Ilya Sutskever. Jukebox: A generative model for music. *arXiv preprint arXiv:2005.00341*, 2020.

[4] Adam Roberts, Jesse Engel, Colin Raffel, Curtis Hawthorne, and Douglas Eck. A hierarchical latent vector model for learning long-term structure in music. In *International conference on machine learning*, pages 4364–4373. PMLR, 2018.

[5] Oliver Bown. Music ai's potential impact: Scoping the terms of the debate about value. In *AIMC 2023 Proceedings: Conference on AI and Music Creativity 2023*, 2023. Conference on AI and Music Creativity.

[6] Nina Kraus and Bharath Chandrasekaran. Music training for the development of auditory skills. *Nature reviews neuroscience*, 11(8):599–605, 2010.

[7] David Hebert and Mikolaj Rykowski. *Music glocalization: Heritage and innovation in a digital age*. Cambridge Scholars Publishing, 2018.

[8] Robert H Woody. Music education students' intrinsic and extrinsic motivation: A quantsitative analysis of personal narratives. *Psychology of Music*, 49(5):1321–1343, 2021.

[9] Gareth Dylan Smith, Bryan Powell, David Lee Fish, Irwin Kornfeld, and Kat Reinhert. Popular music education: A white paper by the association for popular music education. *Journal of Popular Music Education*, 2(3):289–298, 2018.

[10] Atikah Shemshack and Jonathan Michael Spector. A systematic literature review of personalized learning terms. *Smart Learning Environments*, 7(1):33, 2020.

[11] Edward L Deci and Richard M Ryan. The" what" and" why" of goal pursuits: Human needs and the self-determination of behavior. *Psychological inquiry*, 11(4):227–268, 2000.

[12] Peter C Brown, Henry L Roediger III, and Mark A McDaniel. *Make it stick: The science of successful learning*. Harvard University Press, 2014.

[13] Candace Walkington and Matthew L Bernacki. Motivating students by "personalizing" learning around individual interests: A consideration of theory, design, and implementation issues. In *Motivational interventions*, volume 18, pages 139–176. Emerald Group Publishing Limited, 2014.

[14] James M. Renwick and Johnmarshall Reeve. Supporting motivation in music education. In *Oxford handbook of music education*, volume 1, pages 143–162. Oxford University Press, 2012.

[15] Angelika Güsewell. Music education and emotional engagement. *Well-being in education systems*, 2019.

[16] Susan H Cogdill. Applying research in motivation and learning to music education: What the experts say. *Update: Applications of Research in Music Education*, 33(2):49–57, 2015.

[17] Paul Evans. Self-determination theory: An approach to motivation in music education. *Musicae Scientiae*, 19(1):65–83, 2015.

[18] Carlos Xavier Rodriguez. *Coming of age: Teaching and learning popular music in academia*. Michigan Publishing, University of Michigan Library, 2017.

[19] David A Williams. The elephant in the room. *Music Educators Journal*, 98(1):51–57, 2011.

[20] Martina Vasil, Lindsay Weiss, and Bryan Powell. Popular music pedagogies: An approach to teaching 21st-century skills. *Journal of Music Teacher Education*, 28(3):85–95, 2019.

[21] John Kratus. Music education at the tipping point. *Music educators journal*, 94(2):42–48, 2007.

[22] David Brian Williams. Reaching the "other 80%:" using technology to engage "non-traditional music students" in creative activities. In *Proceedings of the Tanglewood II" Technology and Music Education" Symposium. Minnesota*, 2007.

[23] Locky Law. Application of generative artificial intelligence (genai) in language teaching and learning: A scoping literature review. *Computers and Education Open*, page 100174, 2024.

[24] Johan Pauwels, Ken O'Hanlon, Emilia Gómez, Mark Sandler, et al. 20 years of automatic chord recognition from audio. In *Proceedings of the 20th International Society for Music Information Retrieval Conference (ISMIR)*, 2019.

[25] Mojtaba Heydari, Frank Cwitkowitz, and Zhiyao Duan. Beatnet: Crnn and particle filtering for online joint beat downbeat and meter tracking. In *Proceedings of the 22nd International Society for Music Information Retrieval Conference (ISMIR)*, 2021.

[26] Emmanouil Benetos, Simon Dixon, Zhiyao Duan, and Sebastian Ewert. Automatic music transcription: An overview. *IEEE Signal Processing Magazine*, 36(1):20–30, 2018.

[27] Josh Gardner, Ian Simon, Ethan Manilow, Curtis Hawthorne, and Jesse Engel. Mt3: Multi-task multitrack music transcription. *arXiv preprint arXiv:2111.03017*, 2021.

[28] Hayato Sumino, Adrien Bitton, Lisa Kawai, Philippe Esling, and Tatsuya Harada. Automatic music transcription and instrument transposition with differentiable rendering. In *Proceedings of the 2020 Joint Conference on AI Music Creativity*, 2020.

[29] Qiuqiang Kong, Bochen Li, Xuchen Song, Yuan Wan, and Yuxuan Wang. High-resolution piano transcription with pedals by regressing onset and offset times. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:3707–3717, 2021.

[30] Chih-Pin Tan, Shuen-Huei Guan, and Yi-Hsuan Yang. Picogen: Generate piano covers with a two-stage approach. In *Proceedings of the 2024 International Conference on Multimedia Retrieval*, pages 1180–1184, 2024.

[31] Chih-Pin Tan, Hsin Ai, Yi-Hsin Chang, Shuen-Huei Guan, and Yi-Hsuan Yang. Picogen2: Piano cover generation with transfer learning approach and weakly aligned data. *arXiv preprint arXiv:2408.01551*, 2024.

[32] Jongho Choi and Kyogu Lee. Pop2piano: Pop audio-based piano cover generation. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2023.

A Supplementary material

`simplify.py` is a python script that accepts a MusicXML file as input, replaces the left hand piano part with one consisting only of block chords and the right hand part by modifying 16^{th} note pairs with an 8^{th} note of the first pitch. Also included are the output from the Piano2Notes service, and a MusicXML file corresponding to the 4 bars mentioned in the main text.

B Limitations and Future Work

While this paper explores promising applications of AI in music education, it's important to acknowledge several limitations of our approach and the current state of technology:

Accuracy of AI Models: Despite recent improvements, AI models make mistakes and lead to incorrect feedback or exercises.

Limited Scope of Case Studies: Our case studies focus on ear training and piano education and other areas e.g. composition, music theory, other instruments, require further investigation.

Role of Human Music Teachers: While AI-powered tools offer benefits, they cannot fully replicate the nuanced guidance, emotional support, and real-time adaptability of human music teachers.

Cultural Bias: The AI models used in these applications are likely trained on datasets that may not represent the full diversity of global musical traditions.

Limited Assessment: While we present promising applications, this work does not include comprehensive assessment of their effectiveness. Controlled studies comparing these AI-enhanced approaches with traditional methods are needed. Additionally, while personalization can increase engagement, the long-term effects of these AI-powered approaches on music skill development and retention have not been thoroughly studied. Future work should include:

- Comparative studies between traditional and AI-enhanced curricula.
- Quantitative metrics for learning outcomes and skill development.
- User experience studies with diverse student populations.
- Longitudinal research to validate effectiveness over time.

Technical Barriers: The proposed solutions require access to devices capable of running sophisticated AI models, which may not be available to all students, potentially exacerbating educational inequalities.

Privacy Concerns: Analyzing a student's listening history and practice sessions raises important privacy considerations that need to be carefully addressed.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The abstract and introduction motivate the need for personalized music instruction, and talk about the case studies presented in the paper.

   Guidelines:

   - The answer NA means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: There is a dedicated section outlining the limitations of the claims in the paper.

   Guidelines:

   - The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory Assumptions and Proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [NA]

Justification: No theoretical results in the paper.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental Result Reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: We have added outputs from the services used, and the scripts used to manipulate them. As well, links to all the software products, libraries and papers used.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: The scripts used for the piano exercises have been provided. However, the service used (Piano2Notes) is not open source, so we have provided the output we received from the service. Further, the AI model used in Real Ear Trainer is closed source, but the software is freely available on iOS.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental Setting/Details**

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [NA]

Justification: There are no experiments presented in the paper.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment Statistical Significance**

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [NA]

Justification: There are no experiments presented in the paper.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)

- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments Compute Resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [NA]

   Justification: There are no experiments presented in the paper.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code Of Ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: Reviewed and confirmed that the paper conforms to the Ethics Guidelines.

   Guidelines:

   - The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader Impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: The paper outlines the positive social impact of making music education more accessible, as well as highlight the potential negative impact on music teachers.

    Guidelines:

    - The answer NA means that there is no societal impact of the work performed.
    - If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

    Answer: [NA]

    Justification: No data or models are being released.

    Guidelines:

    - The answer NA means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: Credit has been given, and alternative options have also been provided wherever any assets are mentioned.

    Guidelines:

    - The answer NA means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [NA]

    Justification: No new assets released.

    Guidelines:

    - The answer NA means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [NA]

    Justification: No crowdsourcing or research with human subjects used in the paper.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [NA]

    Justification: No crowdsourcing or research with human subjects used in the paper.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.