



2023 EAAI Mentored Undergraduate Research Challenge: Human-Aware AI in Sound and Music

Richard G. Freedman (Smart Information Flow Technologies (SIFT); rfreedman@sift.net)

DOI: [10.1145/3511322.3511328](https://doi.org/10.1145/3511322.3511328)

Abstract

The topic for EAAI 2023's Mentored Undergraduate Research Challenge is Human-Aware AI in Sound and Music. What does that mean? Where are the applications? How can you get started? We break down the topic, discuss applications, and explore project ideas in this column.

Introduction

The EAAI Mentored Undergraduate Research Challenge invites undergraduate students to team up with research mentors and participate in an artificial intelligence (AI) research project from start to finish. For 2023, the specific challenge is for students to complete a feasible and creative research project within the topic of *Human-Aware AI in Sound and Music*. Have you ever imagined people performing music alongside their computers? How about having a program that could identify your emotions by what you play? Perhaps you have wondered whether AI could spice up your exercise routine with a custom soundtrack that matches your movements? These are just a few project ideas that relate to the challenge topic, but they are far from all the possibilities.

Not sure where to start? Do not worry, this column has you covered! Not only is there some code available to help you get started, but we also break down the buzzwords, clarify the concepts, and dig into the details so that you can identify your perfect research project for the challenge. First, we discuss what it means for AI to be human-aware and go over some examples (your smartphone probably is not one of them, believe it or not). Second, we review the history of automating sound and music to figure out how computers, and eventually AI, got involved in audio generation—the software code provided for the challenge is a synthesizer, which plays an important role in the story. Next, we explore possible project

ideas and related works now that you are a pro in human-aware AI and how AI works with sound and music. This column ends with more details about forming a team and registering for the 2023 EAAI Mentored Undergraduate Research Challenge and accessing the synthesizer software code. We look forward to seeing all the cool and creative projects you come up with for the challenge—happy researching!

When is AI Human-Aware?

Over the years, people and artificially-intelligent systems have worked together in a variety of relationships. In the earlier days of AI research, the system was usually a stand-alone program that solved problems by itself; there was no connection to people. Since various applications began to adopt AI, AI has often become a tool for people to use. The roles that people play while working with these tools determine how much the AI needs to understand them.

For example, human-in-the-loop AI relies on the fact that people are sources of information and gives people a role to play in the algorithm. Inverse reinforcement learning (Ng & Russell, 2000) identifies its own reward incentives based on observations of people performing a task. Active learning (Settles, 2009) can ask a person to classify something if the class label is unknown. Semi-autonomous systems (Zilberstein, 2015) let a person take over when the AI is not able to perform well autonomously. Although the AI algorithm has some line of pseudocode that says, “let the human do it,” it is not necessary for the AI to think about the person beyond this step—the system will operate almost the same with or without people during the rest of the algorithm. Similarly, mixed initiative systems (Allen, Guinn, & Horvitz, 1999; Burstein & McDermott, 1996) alternate between people and the AI to iteratively refine a solution. The AI only thinks about people to the extent of receiving their

refinement and making sure to not undo the person's changes without a good reason.

Many "smart" devices today apply AI to complete tasks and provide information at our requests. However, how much does the person's existence matter to the AI outside making the request? As of the writing of this column, it does not matter too much because the AI only processes the request and gathers context relevant to that request. "Find restaurants near me" requires sensing (GPS) to identify the user's location, but this is limited to resolving the "near me" part of the request. The system would not sense the user's location for irrelevant requests, and the AI would not use that information even if it did.

So, what is difference between these examples and human-aware AI? Everything mentioned above is aware of the human's existence, but none of them are *aware of the human's state of being*. People have motives, intents, and emotions that guide their behavior, and we often consider these factors when interacting with each other. Someone is probably hungry if they request, "find restaurants near me," and mentioning restaurants that are currently closed would not help the person because they cannot sate their hunger there. Demonstrations of the intelligent personal assistant Viv ([TechCrunch, 2016](#)) in 2016 expressed human-awareness as it used context from recent requests to infer why the user was making those requests, and then Viv offered additional details and requests based on the inferred motivations. AI systems using responsive planning ([Freedman & Zilberstein, 2017](#)) define their goals based on predictions about an observed person's goal, and the system then chooses actions that account for what the person seems to be doing. Although many AI systems try explaining their decisions using internal algorithms and data, model reconciliation ([Chakraborti, Sreedharan, Zhang, & Kambhampati, 2017](#)) approaches explainable AI as a search for reasons that a person would misunderstand the system's decisions and tries to establish a common understanding.

Rather than seeing people as a part of the equation and environment, human-aware AI reasons about people as autonomous agents and acts accordingly. However, this does not

mean that other methods, such as human-in-the-loop AI and semi-autonomous systems, cannot be human-aware. For example, the CoBots navigating Carnegie Mellon University have thought about which people to ask for help based on their past willingness and current availability ([Rosenthal, Veloso, & Dey, 2011](#)), and there is research on semi-autonomous systems deciding how to transfer control to the human so that they are prepared to take over ([Wray, Pineda, & Zilberstein, 2016](#)).

Why AI in Sound and Music?

Research in automated music began well before computers existed. The Museum Speelklok (<https://www.museumspeelklok.nl/>) in The Netherlands curates various machines that were engineered to play instruments and generate sounds with respect to some modular device that activates certain notes at specific times. Modern instances of these machines include player pianos and music boxes; Martin Molin of music group Wintergatan recently built such a machine that manipulates marbles to play a variety of percussion instruments ([Wintergatan, 2016](#)).

Besides these analogue machines, electronic instruments such as synthesizers and EWIs (electronic wind instruments) can generate unique sounds through connecting circuits that alter the sound waves coming out of electronic speakers. The number of possible sounds is almost endless because the circuit connections produce various mathematical changes that modify and/or combine sound waves, including addition, subtraction, multiplication, and even function composition (<https://digitalsoundandmusic.com/>) ([Burg, Romney, & Schwartz, 2017](#)). These electronic devices can plug into computers as an interface for playing music via MIDI (<https://www.midi.org/>), and sequencers are electronic devices or computer programs that store information about when to play specific sounds. Before computer memory was cheap enough to record, store, and playback a song's-worth of raw audio data, sound cards and chips were effectively synthesizers of various complexity that played music via sequencer programs—many retro videogame consoles used these, and 8-bit

music's iconic beeps were products of their sound chip hardware (Collins, 2008).

These automated music players are just like traditional computer programs, playing exactly what the human composer created—can computers also create music on their own? Even without AI, the answer is yes! For starters, mathematicians have studied properties of music theory and composition for quite some time (American Mathematical Society, n.d.), and this presents the opportunity to create mathematical models that produce original music. Computing has also explored ways to automatically create novel music via *algorithmic composition*. Algorithms can be as simple as programming a sequencer, but many also consider adding additional control flow (loops, conditions, etc.), relational representations (grammars, databases of rules, etc.), and random processes (Markov chains, bag of musical phrases, etc.) for dynamic changes to the notes, lengths, and sounds/instruments (Taube, 2004). Research in procedural content generation for audio also explores these techniques with the goal of constantly producing unique music, rather than composing a song that others may perform and replicate.

Of course, we're most interested in the role AI can play in algorithmic composition and procedural content generation for audio. Given the mathematical and computational nature of the above methods that describe how computers play music, there has been some research in generating algorithms autonomously. Examples include learning grammars from sheet music datasets (Kohonen, 1989) and optimizing tables of if-then rules through global search (Tsubasa & Furukawa, 2012). Global search algorithms such as genetic algorithms can even generate sequences of notes directly (Matić, 2010). Since the recent revival of interest in neural networks via deep learning, directly generating music and sound with some deep model has taken off in popularity—it is almost all you will find when searching for "AI generating music" in a web browser! Whether the network is designed to process high-level representations like annotated sheet music (Yan, Lustig, Vanderstel, & Duan, 2018) and MIDI (Huang, Cooijmans, Roberts, Courville, & Eck, 2017), low-level representations like the raw audio signal (Engel et al., 2017), or a blend of the two (Hung, Chiang,

Chen, & Yang, 2019), there are many documented models from academics and hobbyists alike. You can even play with some online such as Coconet (Huang, Cooijmans, Dinulescu, Roberts, & Hawthorne, 2019), which runs in a Google Doodle to turn your simple music into a Bach-style harmony, and Blob Opera (Li, 2020), which lets you compose songs for an opera-singing quartet.

The Blob Opera application is a fantastic example where AI is involved for both sound and music. The music portion identifies note choices for the unassigned blobs to harmonize with the blobs currently assigned a note to sing. The sound portion generates the operatic voices you hear because human voices are very complex sound waves that are difficult to mimic through connecting circuits or tweaking parameters to send to a sound card (in present-day, this is typically done on the software side as part of a program that computes the sound wave). Furthermore, an instrument has a single sound wave shape that stretches and shrinks depending on the note, but each phoneme (spoken syllable) has a unique sound wave shape. Without AI to deal with the complex and constantly changing sound waves, audio engineers typically craft human-like vocal sounds for electronic instruments with a technique called sampling. Sampling stores a short recording of a sound wave, and the circuitry or program modifies the recording to match a specific note via pitch-shifting. Many pop songs today use sampling and pitch-shifting to create unique sounds, and a number of famous cartoon characters' voices come from pitch-shifting to manipulate the voice actors' recordings (Buckner & Sol, 2019). Yamaha even took this concept to the extreme with their Vocaloid software (<https://www.vocaloid.com/en/>), which contains samples for every phoneme in a supported language so that musicians can compose music with automated singers—think of composing for Vocaloid as programming a sequencer where every possible spoken syllable is its own instrument!

Generating sound goes beyond human voices to making anything audible using sound waves. Audio engineers and foley artists in film work hard to create the perfect sound effects using various materials and activities (Rooster Teeth, 2019; TV Tropes, n.d.), and a

Table 1: Interactive Experiences Involving Sound and Music

	Human-Aware AI with Audio	Human-Aware AI without Audio
Human with Audio	Accompaniment, Tutoring	Sentiment Analysis, Responsive Animation/Dance
Human without Audio	Personalized Performance, Responsive Audio	N/A (No Sound or Music)

number of sound effects also come from mathematically manipulating various sound waves. Though likely not surprising, research with deep neural networks has looked into generating sound effects for video sequences as well (Ghose & Prevost, 2021).

Let's Put Them Together!

When people and intelligent systems interact, the number of things they can do together exceeds the limitations of just music composition and sound generation. In particular, only one person or computer needs to do something with audio in order for human-aware AI to be involved in sound and music. The tasks they can each do depend on who interacts through audio as a medium—is the audio's role to trigger engagement, serve as a response, be the product of the interaction, or do something completely different? Table 1 presents a few possibilities for each allocation, but it is far from complete.

People and computers can generate audio at the same time when they are both directly involved with the sound and music. To avoid a cacophony from misaligned note choices that do not mix well, there needs to be some degree of coordination in the interaction. Specifically, one performer should take the lead while the others focus on *accompaniment* and choose what to play with respect to the lead. Human-aware accompaniment could focus on various aspects like harmonizing with a matching chord, keeping the rhythm through beats that match the tempo, or timing phrases to respond to the lead's melody. Several of these factors made it possible for the Shimon robot to improvise playing a marimba alongside a human musician (Hoffman & Weinberg, 2011). A cognitive model developed in Soar illustrated how an intelligent system can apply reinforcement learning from human drummers to eventually copy them (Derbinsky & Essl,

2012), and the future challenges listed at the end of the work present great opportunities for including human-awareness to become more adaptive. Besides playing together as peer performers, the human-aware AI could serve as an intelligent tutoring system that teaches the human how to perform music or generate sound (like an audio engineer). Existing work in intelligent tutoring systems for music focus almost exclusively on music theory (Angelides & Tong, 1995; Phon-Amnuaisuk & Siong, 2007; Tael, Barreto, & Hammond, 2015).

If humans are the only ones who interact through music and sound, then intelligent systems can express their human-awareness as active listeners to what people play. Sentiment analysis lies at the intersection of AI and affective computing, identifying emotions through data—what could a human-aware AI do with knowledge about a person's emotions based on their performance? For example, sentiment analysis of movie scripts can provide parameters that guide algorithmic composition to suggest non-diegetic (background) music for movies (Kirke & Miranda, 2017). At the moment, sentiment analysis with music is limited to interpreting song lyrics (Shukla, Khanna, & Agrawal, 2017) or properties of the entire song's sound waves (Gómez & Cáceres, 2017). This presents an untapped area of research when it comes to real-time sentiment analysis of human musical performance. If the human-aware AI also has some form of visual output or actuation, then the system can respond to what the human plays through its motions or animations. Animusic (<https://www.animusic.com/>) generates music videos of fantastical instruments that perform the songs, but their techniques rely more on the song data than understanding the performer.

We can think about the opposite direction

when only the intelligent systems interact through sound and music. In this case, audio becomes a mode of response while the computer observes people through other sensors and input modalities. Students have shown quite a bit of creativity when given access to PoseBlocks (Jordan, Devasia, Hong, Williams, & Breazeal, 2021) for basic body tracking, facial expression recognition, and emotion recognition in a blocks-based programming environment. Could human-aware AI identify how people are feeling and use that to influence the sound and music, either matching the mood or trying to cheer people up? If people are being active, such as moving around, dancing, or playing a game, could an intelligent system generate a soundtrack and/or coordinate sound effects to accompany their activity? Movie soundtracks and modern videogame music both aim to ebb-and-flow with the displayed content, building up intensity during action sequences and dramatic moments alongside sound effects that emphasize even the simplest motions (Collins, 2008). With the AI being human-aware rather than only focusing on the state of the media, this form of interaction could act as an audio-based experience manager (Riedl, Stern, Dini, & Alderman, 2008).

Participating in the Challenge

The EAAI 2023 Mentored Undergraduate Research Challenge invites teams of students and mentors to work together on a research project involving human-aware AI in sound and music—the goal is to complete a feasible project and submit a paper about the research to EAAI 2023. As these projects can become ambitious, it is important that students focus on one idea of interest and think of a simple task within that idea. Mentors are expected to be involved as guides for the students to evaluate feasibility, provide tips and ideas, and teach the research pipeline from observation and ideation, to the scientific method, to presenting results in a full paper. To provide a starting point for teams, code for a synthesizer software platform will be made available so that students can focus on the AI research even if they are not familiar with audio. Because the synthesizer communicates through both a user interface (for people) and code (for computers), it is

already set up to handle any of the interactive experience styles described in Table 1. The code and up-to-date information about the challenge are available at https://www.yetanotherfreedman.com/resources/challenge_haaisam.html.

If you are interested in AI techniques that think about the people with whom they interact, the roles computers can play in creating or analyzing sound and music, and exploring how to bring these together, then we encourage you to consider participating in this challenge. Please make sure to form a team that meets the following requirements:

- At least one undergraduate researcher who has not completed a post-secondary education degree. Students in community college are also eligible for this role. Students in this category are expected to play a significant role in the project.
- At least one mentor who has received a Ph.D. and is actively engaged in research. This can be a faculty member at a university or a researcher in industry who has experience training undergraduate students in research. Mentors in this category are expected to be involved with the students regularly to guide them along their journey.
- As long as the above two roles are satisfied, additional team members are allowed. Additional members may include graduate students as long as the undergraduate researchers are actively involved in the research process (ideas, experiment design, paper writing, etc.). Graduate students may also provide additional mentorship to the undergraduate students, but it cannot serve as a substitution for the mentor's participation.

Once you have formed a team, please contact the author of this column with the names, e-mail addresses, and roles (mentor, undergraduate student, etc.) of all team members to register your team in the challenge.

There are no limits to team sizes or number of teams per institution. However, due to conference logistics, there will be a limit on the number of accepted papers for publication and presentation at EAAI 2023. All submitted manuscripts will undergo peer review for writing quality, evidence of quality research at the undergraduate level, and relevance to the

topic of human-aware AI in sound and music. We look forward to seeing your exciting and creative research on this topic!

Acknowledgments

The author would like to thank Dr. Todd W. Neller for establishing the Mentored Undergraduate Research Challenge and all his efforts organizing past challenges; his time and guidance while preparing for the 2023 challenge has been invaluable. The author also thanks Dr. Todd W. Neller, Dr. Scott Alfeld, Dr. Jason R. Wilson, Dr. Tom Williams, Dr. Vasanth Sarathy, Dr. Ross Mead, Dr. Samuel S. Cho, Dr. Eric Aaron, and Dr. Shlomo Zilberstein for their discussions and feedback on the topic and challenge details. Additionally, a special thank you to K. Beth Lane for taking the time to review the column and offer feedback on clarity.

References

- Allen, J. F., Guinn, C. I., & Horvitz, E. (1999). Mixed-initiative interaction. *IEEE Intelligent Systems and their Applications*, 14(5), 14–24. Retrieved from <https://www.microsoft.com/en-us/research/publication/mixed-initiative-interaction/> doi: 10.1109/5254.796083
- American Mathematical Society. (n.d.). *Mathematics & music*. <https://www.ams.org/publicoutreach/math-and-music>. (Retrieved September 3, 2021)
- Angelides, M. C., & Tong, A. K. (1995). Implementing multiple tutoring strategies in an intelligent tutoring system for music learning. *Journal of Information Technology*, 10(1), 52–62. Retrieved from <https://doi.org/10.1177/026839629501000107> doi: 10.1177/026839629501000107
- Buckner, L., & Sol, N. (2019). *Pitch shifting in music: From the Chipmunks to Kanye | Sound Field* [Video]. PBS Learning Media. <https://tpt.pbslearningmedia.org/resource/pitch-shifting-from-chipmunks-to-kanye-video/sound-field/>.
- Burg, J., Romney, J., & Schwartz, E. (2017). *Digital sound & music: Concepts, applications, and science*. Portland, Oregon, USA: Franklin, Beedle.
- Burstein, M. H., & McDermott, D. V. (1996). Issues in the development of human-computer mixed-initiative planning systems. In B. Gorayska & J. L. Mey (Eds.), *Cognitive technology: In search of a human interface* (pp. 285–303). North Holland: Elsevier Science.
- Chakraborti, T., Sreedharan, S., Zhang, Y., & Kambhampati, S. (2017, 8). Plan explanations as model reconciliation: Moving beyond *Explanation as Soliloquy*. In *Proceedings of the twenty-sixth international joint conference on artificial intelligence* (pp. 156–163). Melbourne, Australia: International Joint Conferences on Artificial Intelligence Organization.
- Collins, K. (2008). *Game sound*. Cambridge, Massachusetts, USA: MIT Press.
- Derbinsky, N., & Essl, G. (2012, May). Exploring reinforcement learning for mobile percussive collaboration. In *Proceedings of the twelfth international conference on new interfaces for musical expression* (pp. 70–75). Ann Arbor, Michigan, USA.
- Engel, J., Resnick, C., Roberts, A., Dieleman, S., Norouzi, M., Eck, D., & Simonyan, K. (2017). Neural audio synthesis of musical notes with wavenet autoencoders. In *Proceedings of the thirty-fourth international conference on machine learning* (pp. 1068–1077). Sydney, NSW, Australia: JMLR.org.
- Freedman, R. G., & Zilberstein, S. (2017). Integration of planning with recognition for responsive interaction using classical planners. In *Proceedings of the thirty-first aai conference on artificial intelligence* (p. 4581–4588). San Francisco, California, USA: AAAI Press.
- Ghose, S., & Prevost, J. J. (2021). AutoFoley: Artificial synthesis of synchronized sound tracks for silent videos with deep learning. *IEEE Transactions on Multimedia*, 23, 1895–1907. doi: 10.1109/TMM.2020.3005033
- Gómez, L. M., & Cáceres, M. N. (2017). Applying data mining for sentiment analysis in music. In *Trends in cyber-physical multi-agent systems. the paams collec-*

- tion - fifteenth international conference (pp. 198–205). Porto, Portugal: Springer, Cham. doi: 10.1007/978-3-319-61578-3_20
- Hoffman, G., & Weinberg, G. (2011, June 2). Interactive improvisation with a robotic marimba player. *Autonomous Robots*, 31(1), 133–153. Retrieved from <https://doi.org/10.1007/s10514-011-9237-0> doi: 10.1007/s10514-011-9237-0
- Huang, C.-Z. A., Cooijmans, T., Dinculescu, M., Roberts, A., & Hawthorne, C. (2019, March 20). *Coconet: The ML model behind today's Bach Doodle*. <https://magenta.tensorflow.org/coconet>.
- Huang, C.-Z. A., Cooijmans, T., Roberts, A., Courville, A., & Eck, D. (2017, October). Counterpoint by convolution. In *Proceedings of the eighteenth international society for music information retrieval conference* (pp. 211–218). Suzhou, China.
- Hung, Y.-N., Chiang, I.-T., Chen, Y.-A., & Yang, Y.-H. (2019). Musical composition style transfer via disentangled timbre representations. In *Proceedings of the twenty-eighth international joint conference on artificial intelligence* (pp. 4697–4703). Macao, China: AAAI Press.
- Jordan, B., Devasia, N., Hong, J., Williams, R., & Breazeal, C. (2021, May). PoseBlocks: A toolkit for creating (and dancing) with AI. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(17), 15551–15559. Retrieved from <https://ojs.aaai.org/index.php/AAAI/article/view/17831>
- Kirke, A., & Miranda, E. (2017, March 1). Aiding soundtrack composer creativity through automated film script-profiled algorithmic composition. *Journal of Creative Music Systems*, 1(2). doi: 10.5920/JCMS.2017.01
- Kohonen, T. (1989). A self-learning musical grammar, or “associative memory of the second kind”. In *International 1989 joint conference on neural networks* (Vol. 1, pp. 1–5). Washington DC, USA. doi: 10.1109/IJCNN.1989.118552
- Li, D. (2020, December). *Blob opera*. <https://experiments.withgoogle.com/blob-opera>.
- Matić, D. (2010). A genetic algorithm for composing music. *Yugoslav Journal of Operations Research*, 20(1), 157–177. doi: 10.2298/YJOR1001157M
- Ng, A. Y., & Russell, S. J. (2000). Algorithms for inverse reinforcement learning. In *Proceedings of the seventeenth international conference on machine learning* (pp. 663–670). San Francisco, California, USA: Morgan Kaufmann Publishers Inc.
- Phon-Amnuaisuk, S., & Siong, C. (2007). Web-based music intelligent tutoring systems. In K. Ng & P. Nesi (Eds.), *Interactive multimedia music technologies* (pp. 231–248). IGI Global. doi: 10.4018/978-1-59904-150-6.ch011
- Riedl, M. O., Stern, A., Dini, D. M., & Alderman, J. M. (2008). Dynamic experience management in virtual worlds for entertainment, education, and training. *International Transactions on System Science and Applications*, 3(1), 23–42. Retrieved from <http://www.cc.gatech.edu/~riedl/pubs/itssa.pdf>
- Rooster Teeth. (2019, January 31). *CRWBY ep 12: The other 50% | rooster teeth* [Video]. YouTube. <https://www.youtube.com/watch?v=ci5EJN3xP8w>.
- Rosenthal, S., Veloso, M., & Dey, A. K. (2011, July). Is someone in this office available to help me? Proactively seeking help from spatially-situated humans. *Journal of Intelligent & Robotic Systems*, 66(1), 205–221. Retrieved from <https://doi.org/10.1007/s10846-011-9610-4> doi: 10.1007/s10846-011-9610-4
- Settles, B. (2009). *Active learning literature survey* (Computer Sciences Technical Report No. 1648). Madison, Wisconsin, USA: University of Wisconsin–Madison.
- Shukla, S., Khanna, P., & Agrawal, K. K. (2017). Review on sentiment analysis on music. In *2017 international conference on infocom technologies and unmanned systems (trends and future directions)* (pp. 777–780). Dubai, United Arab Emirates. doi: 10.1109/ICTUS.2017.8286111
- Taele, P., Barreto, L., & Hammond, T. (2015). Maestoso: An intelligent educational sketching tool for learning music theory. In *Proceedings of the twenty-seventh*

conference on innovative applications of artificial intelligence (pp. 3999–4005). Austin, Texas, USA.

Taube, H. K. (2004). *Notes from the metalevel: An introduction to computer composition*. New York, New York, USA: Routledge. Retrieved from <https://doi.org/10.4324/9781315078182> doi: 10.4324/9781315078182

TechCrunch. (2016, May 9). *Beyond Siri: The world premiere of Viv with Dag Kittlaus* [Video]. YouTube. <https://www.youtube.com/watch?v=Rblb3sptgpQ>.

Tsubasa, T., & Furukawa, K. (2012, July). Automatic melodic grammar generation for polyphonic music using a classifier system. In *Proceedings of the ninth sound and music computing conference* (pp. 1–7). Copenhagen, Denmark.

TV Tropes. (n.d.). *The coconut effect*. <https://tvtropes.org/pmwiki/pmwiki.php/Main/TheCoconutEffect>. (Retrieved September 7, 2021)

Wintergatan. (2016, March 1). *Wintergatan - marble machine (music instrument using 2000 marbles)* [Video]. YouTube. <https://www.youtube.com/watch?v=Rblb3sptgpQ>.

Wray, K. H., Pineda, L., & Zilberstein, S. (2016, 7). Hierarchical approach to transfer of control in semi-autonomous systems. In *Proceedings of the twenty-fifth international joint conference on artificial intelligence* (pp. 517–523). New York, New York, USA: International Joint Conferences on Artificial Intelligence Organization.

Yan, Y., Lustig, E., Vanderstel, J., & Duan, Z. (2018, September). Part-invariant model for music generation and harmonization. In *Proceedings of the nineteenth international society for music information retrieval conference* (pp. 204–210). Paris, France.

Zilberstein, S. (2015). Building strong semi-autonomous systems. In *Proceedings of the twenty-ninth conference on artificial intelligence* (pp. 4088–4092). Austin, Texas. Retrieved from <http://rbr.cs.umass.edu/shlomo/papers/Zaaai15.html>



Richard G. Freedman is a Researcher at SIFT. In his quest to understand how machines can more naturally interact and play with people, Freedman studies the intersection of artificial intelligence and human-computer/robot

interaction. Freedman enjoys sharing opportunities that inspired him, introducing STEAM to K-12 and research to undergraduates.