

Data Sonification as a Scientific Tool: How Astronomers Turn Light Curves into Music

Wonbin Yoon
santawonbin@gmail.com

ABSTRACT

Data sonification, the translation of numerical data into sound, has emerged over the past decade as a serious methodology in astronomy, driven by the fields' dependence on time-series data, the accessibility needs of blind and low-vision researchers, and the sheer scale of astrophysical data. This review synthesizes empirical and technical literature on astronomical sonification to evaluate whether it functions as a genuine scientific tool rather than merely an artistic or outreach exercise. It surveys the three dominant techniques and examines controlled signal-detection studies, accessibility and public-engagement surveys, and emerging machine-learning applications. The evidence indicates that sonification performs comparably to visual analysis at high signal-to-noise ratios and narrows the performance gap between expert and non-expert listeners but underperforms at low signal-to-noise ratios and exhibits a consistent liberal response bias toward false positives. Sonification has also demonstrated concrete value in accessibility, education, and public engagement, and in at least one documented case revealed a scientific feature that visual inspection had missed. However, the field remains constrained by the absence of standardized mapping conventions and reporting practices, small and inconsistent sample sizes, and a persistent gap between accessibility claims and genuine co-design with blind and low-vision communities. As data volumes from next-generation surveys such as the Vera Rubin Observatory and the Square Kilometer Array exceed human listening capacity, machine learning and standardization are identified as the critical next steps for sonification to move from a proof-of-concept tool to a routine component of the astronomer's toolkit.

INTRODUCTION

In 570 BCE, Pythagoras proposed that the movements of celestial bodies produce a harmony he called the music of the spheres (Godwin, 1993). Today, astronomers are making it literal, not philosophically but computationally.

Data sonification is the process of translating numerical or visual data into sound, allowing listeners to perceive patterns through pitch, volume, rhythm, and timbre (Kramer and Walker, 1999). It is not a new idea, but it has become a serious research methodology in the last decade, particularly in astronomy. Astronomy turns out to be an unusually good fit for sound. Many of its core datasets are time-based: a star's brightness measured night after night, a pulsar's radio pulses arriving at millisecond intervals, an exoplanet's transit dipping and recovering over hours. Sound is naturally suited to time-series data in a

Oct 2026
Vol 12. No 5.

way that static visual graphs are not. The human auditory system evolved to be sensitive to exactly the kinds of patterns astronomers care about such as rhythms, interruptions, gradual shifts, sudden anomalies. Our ears catch a skipped heartbeat before we consciously register it. The question is whether they can do the same for a skipped stellar pulse.

There is also an access problem that sonification directly addresses. Astronomy is one of the most visually dominated sciences, and that creates a barrier for blind and low-vision (BLV) researchers and students (Noel-Storr and Willebrands, 2022). Sonification does not just offer an alternative representation for people who cannot see visual data. In at least one documented case, it revealed a feature that sighted researchers had missed relying on vision alone, a case examined in detail in Section 4 (Quinton et al., 2020).

The practical pressure driving all of this is scale. The National Aeronautics and Space Administration's (NASA's) Transiting Exoplanet Survey Satellite (TESS) and the Kepler Space Telescope generate light curves, i.e. measurements of a star's brightness over time, for hundreds of thousands of stars (Borucki et al., 2010; Ricker et al., 2014). Each one needs to be screened for transiting planets, stellar oscillations, and other signals. Visual analysis at that volume is not scalable nor sustainable. Researchers are now asking whether the human auditory system, which processes temporal patterns with exceptional precision, can help where the eye is overwhelmed. Recent empirical studies in astronomy journals and audio engineering conferences suggest the answer is a qualified yes. This review examines how sonification works, what the evidence says about its scientific effectiveness, how it relates to music composition, and where the field is heading.

ASTRONOMICAL DATA AND THE CASE FOR SOUND

A light curve is a graph of how bright an astronomical object appears over time (**Figure 1A**). Although it sounds simple, the patterns hidden in light curves reveal many scientific details of stars, planets, and stellar events. When a planet passes in front of its host star, the star's apparent brightness drops by a small fraction, typically less than a percent, and then recovers as the planet moves on (Charbonneau et al., 2000). This dip repeats at the planet's orbital period. Finding it reliably against a noisy signal is the core challenge of exoplanet detection. The scale of that challenge is considerable. TESS observes more than 200,000 stars, each producing a light curve that needs to be screened for transits. García Riber and Serradilla demonstrated what a sonification pipeline can do with confirmed signals by sonifying 132 exoplanet transits (**Figure 1B**) from TESS Data Validation Transit files (Riber and Serradilla, 2022). The harder question of whether such a pipeline can work effectively on unconfirmed candidates still remains to be answered.

Sound offers a physiologically motivated alternative to visual inspection. The auditory system processes time-varying patterns with a resolution that complements, rather than duplicates, the spatial strengths of the visual system. A heartbeat that skips is noticed before it is consciously measured. If a light curve that

dips periodically can be rendered as a sound that drops in pitch at the moment of transit, the ear may catch what the eye overlooks. It was first scientifically demonstrated that variable stars pulsate according to acoustic wave equations that are mathematically identical to those governing musical instruments (Kolláth and Keuler, 2006). Their composition *Stellar Music No. 1*, in which every note is produced by a stellar model driven by real physical data, established that scientifically motivated sound mapping has legitimate scholarly standing. NASA's Chandra X-ray Observatory launched its System Sounds program in 2020, and coordinated institutional investment has grown steadily since (Fisher, 2023).




Figure 1. Light curves turned into sonified data. (A) Folded lightcurve and transit model for TIC220479565 as a sub-Neptune close to the transition between super-Earths and sub-Neptunes transiting the M2 dwarf. Reprinted from Ribera and Serradilla, licensed under CCBY-NC 4.0. (B) The pipeline from raw data to the final soundscape in visualized and sonified form. Reprinted from Ribera and Serradilla, licensed under CCBY-NC 4.0.

HOW ASTRONOMERS TURN DATA INTO SOUND

There are three main sonification techniques, each with different scientific and perceptual properties. Parameter mapping is the most common. A data variable such as stellar brightness is assigned to an audio parameter, most commonly pitch. As brightness falls during a transit, the tone drops. As brightness recovers, the pitch rises. The *astronify* tool, developed by the Space Telescope Science Institute, uses this approach and served as the test instrument in two of the field's most important studies, as shown in Figure 2A (Guiotto Nai Fovino, 2024; Tucker-Brown et al., 2022). The central design question that parameter mapping raises is which data-to-sound assignments are intuitive for listeners, and which introduce systematic perceptual distortions. Fovino et al. compared pitch-only and duration-only mappings and found that they produced measurably different detection rates (Guiotto Nai Fovino, 2024).

Additive synthesis, a more physically grounded approach, constructs complex tones by stacking overtones

derived directly from a star's oscillation data, rather than assigning brightness to pitch by convention. García Riber's *Sonifigrapher* tool uses this method; its acoustic output reflects a star's actual physical oscillations rather than an arbitrary mapping convention (García Riber, 2019). García Riber demonstrated this on Tabby's Star, which has unusual, aperiodic dips in its light curve, as shown in Figure 2B. The resulting sonification had value both as scientific analysis and as experimental music, not because the two goals were blended but because the physical mapping was transparent enough to serve both purposes separately.

A third method, binaural or 3D sonification, transforms raw data into immersive soundscapes rather than a single audio stream. García Riber and Serradilla encoded CoRoT (Convection, Rotation, and planetary Transits) light curves in three-dimensional audio with the help of deep learning, letting listeners locate signals in space rather than only in time (Riber and Serradilla, 2024). This is particularly useful when multiple stellar sources need to be distinguished simultaneously. Figure 2C shows the mechanism of how binaural sonification works to help listeners perceive audio signals in 3D.

The three techniques trade off differently. Parameter mapping is the simplest to implement and the most thoroughly tested for signal detection, but its data-to-sound assignments are a design choice rather than a physical necessity, which is why different mappings of the same data can produce measurably different detection rates (Guiotto Nai Fovino, 2024). Additive synthesis gives up some of that simplicity for a mapping that is physically transparent, at the cost of having been demonstrated mainly in single case studies rather than controlled detection experiments. Binaural sonification is the least tested of the three as a detection tool, but it is the only one built to resolve multiple simultaneous sources in space rather than a single signal in time. Roughly 100 sonification projects and related techniques have been catalogued, but no consensus has been found on best practice (Zanella et al., 2022). Misdariis et al., drawing on research from IRCAM, the Paris institute where much of the modern science of music perception was developed, argued that how navigable a sonification feels directly affects whether it gets used, particularly for BLV listeners (Misdariis et al., 2022). Knowledge from music composition therefore contributes directly to the perceptual science and methodology design of sonification: composers working in Western tonal music have accumulated centuries of practical understanding about which intervals untrained listeners can distinguish reliably, which rhythmic patterns are memorable, and which timbral contrasts encode categorical difference consistently across different listening conditions. A sonification that is perceptually tiring or tonally ambiguous is impractical to use for extended analysis, no matter how accurately it encodes the data, which is why timbral contrast and rhythmic hierarchy function as core instruments of scientific communication rather than merely aesthetic choices.




Figure 2. Three main sonification techniques. (A) Astronify tool used to create synthetic light curves for participant testing with varying signal-to-noise ratios (SNRs). Reprinted from Tucker Brown et al. with permission. (B) Sonifigrapher tool used to create a light curve showing a planet transit. Reprinted from Riber, licensed under CCBY-NC 4.0. (C) Schematic drawing of binaural sonification. Listener explores the 3D soundscape and intuitively locates potential astrophysical signals in space and time, facilitating detection of weak or rare events. The binaural rendering presents sounds at virtual positions around the listener using head-related transfer functions and helps the listener perceive signals in 3D.

SCIENTIFIC EFFICACY EVIDENCE OF SONIFICATION

Several studies rigorously tested sonification as a signal detection tool. One study tested 192 participants, ranging from professional astronomers to members of the public, on synthetic light curves containing zero, one, or two transit-like signals across a range of signal-to-noise ratios (SNR); the results are now foundational to the field (Tucker-Brown et al., 2022). At high SNR values of 30 and 100, sonification achieved 85-100% successful signal detection, comparable to visual analysis. More interesting was what happened to the gap between experts and non-experts: in visual conditions, trained astronomers substantially outperformed non-astronomers, but in the sonification condition that gap narrowed considerably. If the specialized perceptual training that underpins visual light curve reading matters less when data is rendered as sound, sonification could meaningfully lower the barrier to participation in astronomical data analysis. At medium and low SNR, however, visual analysis remained superior, and sonification's detection rate fell as signal strength decreased, which sets a real ceiling on where the method is currently useful. The authors also documented a liberal response bias: participants in the

sonification condition were more likely to report a signal whether or not one was actually present, a pattern that raises concern about false positive rates if sonification were integrated into an automated pipeline.

A study by Fovino replicated and extended these findings (Guiotto Nai Fovino, 2024). Visual stimuli again led to better overall detection performance, but auditory performance stayed above chance, and the liberal response bias appeared again, confirming that it is a property of sonification-based analysis in general rather than an artifact of one study's design. The study also showed that different mapping strategies produced meaningfully different detection rates, indicating that sonification design choices carry measurable scientific consequences. Together, these two controlled studies converge on the same pattern: sonification approaches visual performance at high SNR, degrades faster than vision as SNR falls, and consistently biases listeners toward reporting signals that are not present.

A different kind of evidence comes from two case studies rather than controlled experiments, and it is worth being precise about what it does and does not show. In one, an astronomer analyzing an exoplanet's transmission spectrum had missed a water absorption feature during visual inspection; when the same data was rendered as sound, the feature was immediately identifiable despite the noise that had masked it visually (Quinton et al., 2020). A follow-up study showed that planetary orbital resonances in the asteroid belt could be reliably distinguished by ear (Quinton et al., 2021). These are demonstrations that sonification can support discovery in specific instances. They are not, on their own, evidence that it reliably produces new science: unlike the detection-rate studies above, they are single, self-reported cases without a controlled comparison. This distinction is critical. The controlled studies show that sonification detects known signal types as well as vision under good conditions; the case studies show that, on at least two occasions, a person heard something in the data they had not seen. The first is evidence of improved or comparable pattern detection under test conditions; the second is a suggestive but unreplicated report of discovery, and the two should not be cited interchangeably.

ACCESSIBILITY, EDUCATION, AND PUBLIC EDUCATION OF SONIFICATION

The literature on sonification's social impact supports two different kinds of claims, and they are worth separating. The first is that sonification is accessible: it gives BLV audiences a usable, non-visual way to engage with astronomical data, and audiences respond to it positively. One study surveyed 3,184 participants, both sighted and BLV, following exposure to NASA sonifications of three astronomical objects (Arcand, 2024). Significant self-reported learning gains were observed across both groups, alongside increased trust in the scientific data and strong positive experiential responses, including enjoyment and motivation to learn more. Because this sample was far larger than the detection studies in Section 4 (Guiotto Nai Fovino, 2024; Tucker-Brown et al., 2022), it is well suited to measuring broad social impact, but what it measured was attitude and self-reported learning, not task performance. Fovino's Edukoi project (2024), which engaged approximately 150 middle school students with

interactive sonification tools, shows engagement, not performance. Students without specialist training responded positively to data exploration through sound.

The second, stronger claim is that sonification improves scientific participation or performance specifically for BLV researchers, and the evidence for this is thinner. The controlled studies that do measure performance (Guiotto Nai Fovino, 2024; Tucker-Brown et al., 2022) were not run with BLV-specific samples, so they show what sonification can do for a general population under laboratory conditions, not what it does for BLV researchers doing real research.

Harrison et al. (2022) argued that closing this gap requires genuine interdisciplinary co-design, in which BLV individuals, sound designers, and scientists develop the methodology jointly from the outset; the Audible Universe workshop that resulted from this argument demonstrated that such collaboration is productive and publishable, but it is a demonstration of process, not a performance comparison. This distinction matters because the accessibility argument for sonification is currently ahead of the performance argument. Existing tools were largely designed about BLV users rather than with them (Zanella et al., 2022), and a separate line of critique has directly asked whether inclusion is being achieved in practice or only claimed in principle (Johanna Casado and García, 2024). Until studies measure BLV researchers' task performance directly, rather than general audience engagement, the claim that sonification improves scientific participation for BLV researchers remains plausible but unproven.

THE FUTURE OF SONIFICATION WITH BIG DATA, MACHINE LEARNING (ML), AND ARTIFICIAL INTELLIGENCE (AI)

Up to this point, the scale problems in this review have been manageable. However, survey data volumes keep growing, and at some point, the numbers will become too large for a human to process without the help of AI.

The Vera Rubin Observatory's Legacy Survey of Space and Time started full operations in 2025. It produces about fifteen terabytes of data every night and is expected to find around ten million transient events a year (Ivezić et al., 2019). The Square Kilometre Array, currently being built in South Africa and Australia, is projected to generate data at rates that rival or exceed global internet traffic (Scaife, 2020). Even if sonification only matches a fraction of the detection rates shown in Tucker-Brown et al. (2022) and Fovino (2024), there is no way a human listener could keep up with signals at this volume. Turning sonification into a tool that can effectively process such large volumes of surveys begs this question: who, or what, is actually doing the listening?

The most ambitious answer so far is the Auditory Virtual Observatory (García Riber, 2024). It is modeled on the Visual Virtual Observatory, a system that lets astronomers pull standardized data from dozens of observatories through one (Hanisch et al., 2015). The Auditory Virtual Observatory would do the same thing for sound. Sonified data would become something you can search for, download, and trust to follow

Oct 2026

Vol 12. No 5.

a common format (García Riber, 2024). The prototype could already retrieve data from the Kepler Objects of Interest catalogue and pull in stellar spectral libraries from the Spanish Virtual Observatory (García Riber, 2024). Machine learning was built into the prototype from the start because generating sonifications by hand for an archive this size is not realistic. The real shift is that it is the computer that does the listening at first.

sonoUno takes a completely different approach to the same scaling problem (J Casado and García, 2024). Instead of building one big centralized system, it drops sonification into the tools astronomers already use. It works with Sloan Digital Sky Survey (SDSS) spectroscopic data and All-Sky Automated Survey for Supernovae (ASAS-SN) variable star photometry right inside a normal Python environment, next to the graphs they're already looking at. Sonification becomes a routine step in data reduction, as one mode of inspection among several without a specialist software. This matters more than it sounds like it should. A tool that makes you leave your usual workflow gets used occasionally. A tool that's already present in the workflow gets used habitually.

Herakoi points in a third direction, which is real-time interactive image sonification (Ginolfi et al., 2024). Herakoi is a hand-tracking, webcam-based system where a user moves their hand over an astronomical image to sonify pixel colors in real time. For a survey like Legacy Survey of Space and Time (LSST), where catching an event within hours can matter scientifically, being able to listen to incoming data as it arrives instead of checking a visual alert feed every so often is appealing. But this possibility has not yet been tested, and the current tools don't support streaming alert data. Tucker-Brown et al. (2022) and Fovino (2024) only tested people listening to pre-recorded sounds in a controlled setting. No research has yet empirically validated that a person can reliably detect a real signal in a live, streaming audio feed.

Machine learning (ML)'s role in sonification can work in two directions. The first is automating the parameter-mapping step. Instead of a researcher manually deciding that brightness should map to pitch, a model learns from data which mappings produce reliable responses from listeners (García Riber, 2024). The second role would be using ML for automated listening: training a model to recognize transit-like patterns directly in sonified audio, flagging events without any human listener at all. No published work has yet attempted this specific approach. Multiple studies have found that human listeners tend to over-report signals that aren't there (Guiotto Nai Fovino, 2024; Tucker-Brown et al., 2022). If a machine listener is trained on data annotated by humans who have that bias, there's a high probability that the model inherits the same bias.

Generative AI introduces a complex problem, arguably the most important one for the long-term credibility of the field. Models trained on large volumes of music can now produce audio that sounds complex, intentional, and genuinely good by ear. Applied to astronomical data, such models can generate sonifications that are perceptually engaging and compositionally sophisticated without meaningful and direct mapping between the data and sound parameters. This is what researchers warned about as scientifically problematic (Zanella et al., 2022). The problem becomes significant if it can sound convincing while carrying zero real information. If a listener can't tell the difference between a sound

built from an honest physical mapping and one a model just generated to sound plausible, that's not a trivial mistake. It breaks the entire reason sonification was supposed to be trustworthy in the first place.

Dealing with this will probably require the kind of standardization discussed in this review: shared benchmark datasets, agreed mapping conventions, and minimum reporting standards that specify the complete computational path from raw data to audio output for any published sonification. As AI tools that can fake plausible-sounding astronomical audio get more powerful and widely accessible, this kind of paper trail becomes more necessary.

The most optimistic version of this future isn't one where AI and ML replace the human ear but extend the conditions under which human listening is productive. Imagine a model that filters ten million LSST alerts down to the ten thousand sonically anomalous events, and presents them to humans for a final check. This setup would preserve the perceptual strengths of human auditory processing, including temporal sensitivity, pattern recognition, and anomaly detection. At the same time, it would solve the volume problem that makes unassisted human listening impractical at LSST or SKA (Square Kilometre Array) scale. Whether the detection results from controlled lab studies hold up under noisier, higher-volume, real-world conditions is the question the field now needs to answer.

LIMITATIONS AND FUTURE OUTLOOK

The current protocols and design methodologies of sonification have five limitations that should temper how the evidence above is read.

First, there is no standardized protocol or musification boundary. Different studies use different tools, different mapping conventions, and different participant pools so direct comparison across studies requires caution, and a true meta-analysis is not yet possible. Without an agreed line between scientifically grounded sonification and artistic output, reproducibility and peer review consistency are both compromised.

Second, participant populations are small and inconsistent across studies. The benchmark detection studies involve dozens to low hundreds of participants (Guiotto Nai Fovino, 2024; Tucker-Brown et al., 2022), and their samples differ in composition, professional astronomers against members of the public against students, in ways that are rarely matched across studies. Only Arcand (2024) approaches a sample size ($N = 3,184$) suited to social-impact claims, and that study measures attitudes rather than detection performance. This combination of small samples and inconsistent populations limits how far effect sizes and bias estimates can be generalized.

Third, the liberal response bias documented in Section 4, the tendency of listeners to report a signal whether or not one is present, is a limitation as well as a finding. It has been replicated across at least two independent studies (Guiotto Nai Fovino, 2024; Tucker-Brown et al., 2022), but no published work has yet identified its source, quantified how it changes with training or experience, or established how much it would inflate false-positive rates in an automated sonification pipeline.

Oct 2026

Vol 12. No 5.

Fourth, the SNR threshold below which sonification becomes unreliable is unknown. The consistent finding that sonification underperforms visual analysis at low SNR does not, by itself, specify where that threshold sits.

Fifth, blind and low-vision researchers remain underrepresented in study design rather than only in study samples. The BLV co-design deficit noted by Zanella et al. (2022), tools built about BLV users rather than with them, means that accessibility claims in the literature outrun the evidence base for them, as discussed in Section 5. Technical accessibility is not the same as usability, and it remains untested whether repeated engagement with sonification builds durable analytical skill in professional astronomers, BLV scientists, or students.

These limitations point to where the field is headed. Automation at scale is the most technically immediate priority: the Auditory Virtual Observatory concept and tools like herakoi (Ginolfi et al., 2024), which processes astronomical images in near real time using machine learning, trace a path from demonstration to pipeline. Connecting sonification to existing observatory data-reduction workflows would enable routine use alongside visual analysis. Standardization is the methodological priority: the field needs shared benchmark datasets, agreed parameter-mapping conventions for common data types, and a minimum reporting standard for sonification studies, without which the evidence base remains fragmented. Co-design with BLV communities is the equity priority: moving from single-session usability testing to longitudinal studies, in which BLV scientists use sonification tools as part of real research workflows over extended periods, is the missing step for evaluating both accessibility and performance.

CONCLUSION

Data sonification in astronomy has moved from Greek philosophy to peer-reviewed empirical science in a remarkably short period, with Zanella et al. (2022) marking the clearest institutional milestone and studies spanning signal detection, education, and public engagement building on it since.

Sonification is a complement to visual analysis, not a replacement. It performs reliably at high signal-to-noise ratios, narrows the expert–non-expert performance gap, delivers genuine value in education and public engagement, and has, in at least one documented case, revealed a scientific result that visual inspection missed. These strengths are balanced by real limitations, among them the absence of standardized protocols, a persistent liberal response bias, small and inconsistent samples, and a co-design deficit that leaves some accessibility claims ahead of the evidence for them.

The field's cross-disciplinary character, spanning astronomy, audio engineering, human–computer interaction, and music, is both its greatest source of creativity and the reason its evidence base remains fragmented. Standardization, sustained co-design with BLV communities, and integration into existing data-reduction infrastructure are what will determine whether sonification becomes a routine part of the astronomer's toolkit or remains a compelling demonstration.

REFERENCES

- Arcand, K.K., 2024. A Universe of Sound: Processing NASA Data into Sonifications to Explore Participant Response. *Frontiers in Communication* 9. <https://doi.org/10.3389/fcomm.2024.1288896>
- Borucki, W.J., et al., 2010. Kepler Planet-Detection Mission: Introduction and First Results. *Science* 327, 977–980. <https://doi.org/10.1126/science.1185402>
- Casado, J., García, B., 2024a. Astronomical images sonification: inclusion or outreach? <https://doi.org/10.48550/ARXIV.2402.00952>
- Casado, J., García, B., 2024b. A multimodal approach to data analysis in astronomy: sonouno applications in photometry and spectroscopy. *RAS Techniques and Instruments* 3, 625–635. <https://doi.org/10.1093/rasti/rzae042>
- Charbonneau, D., Brown, T.M., Latham, D.W., Mayor, M., 2000. Detection of Planetary Transits Across a Sun-like Star. *The Astrophysical Journal* 529, L45–L48. <https://doi.org/10.1086/312457>
- Fisher, D.G., 2023. NASA Launches the Chandra X-Ray Observatory.
- García Riber, A., 2024. Toward an Auditory Virtual Observatory. *Journal of the Audio Engineering Society* 72, 341–351. <https://doi.org/10.17743/jaes.2022.0146>
- García Riber, A., 2019. Sonifigrapher. Sonified Light Curve Synthesizer, in: *Proceedings of the 25th International Conference on Auditory Display (ICAD 2019)*.
- Ginolfi, M., Di Mascolo, L., Zanella, A., 2024. herakoi: a sonification experiment for astronomical data. <https://doi.org/10.48550/ARXIV.2412.09152>
- Godwin, J., 1993. *The harmony of the spheres: A sourcebook of the Pythagorean tradition in music*. Inner Traditions, Rochester VT.
- Guiotto Nai Fovino, L.G.N., 2024. Evaluating the effectiveness of sonification in science education using Edukoi. *Personal and Ubiquitous Computing*. <https://doi.org/10.1007/s00779-024-01809-5>
- Hanisch, R.J., Berriman, G.B., Lazio, T.J.W., Emery Bunn, S., Evans, J., McGlynn, T.A., Plante, R., 2015. The Virtual Astronomical Observatory: Re-engineering access to astronomical data. *Astronomy and Computing* 11, 190–209. <https://doi.org/10.1016/j.ascom.2015.03.007>
- Harrison, C., Zanella, A., Bonne, N., Meredith, K., Misdariis, N., 2022. Audible universe. *Nat Astron* 6, 22–23. <https://doi.org/10.1038/s41550-021-01582-y>
- Ivezić, Ž., et al., 2019. LSST: From Science Drivers to Reference Design and Anticipated Data Products. *ApJ* 873, 111. <https://doi.org/10.3847/1538-4357/ab042c>

- Kolláth, Z., Keuler, J., 2006. Stellar acoustics as input for music composition. *Musicae Scientiae* 10. <https://doi.org/10.1177/1029864906010001081>
- Kramer, G., Walker, Br., 1999. Sonification Report: Status of the Field and Research Agenda. Prepared for the National Science Foundation by members of the International Community for Auditory Display. The Int. Community for Auditory Display.
- Misdariis, N., Özcan, E., Grassi, M., Pauletto, S., Barrass, S., Bresin, R., Susini, P., 2022. Sound experts' perspectives on astronomy sonification projects. *Nature Astronomy* 6, 1249–1255. <https://doi.org/10.1038/s41550-022-01821-w>
- Noel-Storr, J., Willebrands, M., 2022. Accessibility in astronomy for the visually impaired. *Nat Astron* 6, 1216–1218. <https://doi.org/10.1038/s41550-022-01691-2>
- Quinton, M., McGregor, I., Benyon, D., 2021. Sonification of Planetary Orbits in Asteroid Belts, in: *Proceedings of the 16th International Conference on Audio Mostly (AM '21)*. pp. 72–80. <https://doi.org/10.1145/3478384.3478390>
- Quinton, M., McGregor, I., Benyon, D., 2020. Sonification of an Exoplanetary Atmosphere, in: *Proceedings of the 15th International Conference on Audio Mostly (AM '20)*. pp. 191–198.
- Riber, A.G., Serradilla, F., 2024. Corot Light Curve Ambsonics/Binaural Sonification, in: *Proceedings of the 29th International Conference on Auditory Display (ICAD2024)*. Presented at the ICAD 2024: The 29th International Conference on Auditory Display, International Community for Auditory Display, Rensselaer Polytechnic Institute, pp. 66–72. <https://doi.org/10.21785/icad2024.010>
- Riber, G.A., Serradilla, G.F., 2022. Sonification of TESS data validation timeseries files, in: *The 27th International Conference on Auditory Display (ICAD 2022)*. Presented at the International Conference on Auditory Display.
- Ricker, G.R., et al., 2014. Transiting Exoplanet Survey Satellite. *J. Astron. Telesc. Instrum. Syst* 1, 014003. <https://doi.org/10.1117/1.JATIS.1.1.014003>
- Scaife, A.M.M., 2020. Big telescope, big data: towards exascale with the Square Kilometre Array. *Phil. Trans. R. Soc. A* 378, 20190060. <https://doi.org/10.1098/rsta.2019.0060>
- Tucker-Brown, J., Harrison, C.M., Zanella, A., Trayford, J., 2022. Evaluating the efficacy of sonification for signal detection in univariate, evenly sampled light curves using ASTRONIFY. *Monthly Notices of the Royal Astronomical Society* 516, 5674–5683. <https://doi.org/10.1093/mnras/stac2526>
- Zanella, A., Harrison, C.M., Lenzi, S., Cooke, J., Damsma, P., Fleming, S.W., 2022. Sonification and sound design for astronomy research, education and public engagement. *Nat Astron* 6, 1241–1248. <https://doi.org/10.1038/s41550-022-01721-z>