

Ethical Challenges in AI-Driven Soundscape Monitoring: Balancing Ecological Insights and Privacy

Sridharan Sankaran^[0009-0005-1341-6384]

Fynro Advisory Services, Chennai, India
sridharan.sankaran@fynro.in

Abstract. As AI-powered acoustic sensing technologies proliferate across ecological and urban environments, machine listening systems are increasingly tasked with interpreting, classifying, and acting upon soundscapes once mediated through human perception. This paper explores the ethical and socio-technical dimensions of AI-driven soundscape monitoring and analytics, positioning it within the broader field of human auditory ecology. We argue that while such systems promise unprecedented ecological insight—detecting biodiversity shifts, illegal activity, and environmental change—they also pose significant risks of privacy intrusion, algorithmic bias, and acoustic surveillance. Bridging computational bioacoustics, auditory neuroscience, and critical data ethics, the paper examines how machine listening reconfigures the sensory and political dynamics of listening. Drawing on interdisciplinary scholarship, we develop a normative framework for ethical soundscape AI, foregrounding principles of acoustic privacy, data minimization, participatory governance, and environmental justice. We also propose a technical architecture that integrates edge computing, differential privacy, federated learning, and homomorphic encryption to operationalize these commitments. By interrogating who gets to listen, what is heard, and whose voices are silenced or amplified, this paper calls for a reimagining of soundscape AI—not as a tool of extractive surveillance, but as a relational, accountable, and pluralistic infrastructure. In doing so, we contribute to advancing human auditory ecology in the age of algorithmic listening.

Keywords: Auditory Ecology, Soundscape AI, Ethical AI.

1 Introduction

1.1 Listening Machines in Ecological Contexts

We are entering an era where machines are not only watching but also listening—pervasively, persistently, and at scale. AI-powered acoustic sensing systems are being deployed across ecological and urban environments to interpret soundscapes once mediated solely by human perception, with the application of machine learning techniques in ecoacoustics showing exponential growth in recent years (Nieto-Mora et al., 2023).



All rights remain with the authors under the Creative Commons Attribution 4.0 International License (CC BY 4.0).

Proc. of the 17th Int. Symposium on Computer Music Multidisciplinary Research,
London, United Kingdom, 2025

Enabled by advances in edge computing and machine learning, these “listening machines” form part of a broader sensor infrastructure that now monitors biodiversity, illegal activity, and environmental change in real time.

Ecological soundscape monitoring has emerged as a powerful tool within this paradigm. Traditionally reliant on field biologists and manual annotation, sound-scape analysis now benefits from AI models capable of processing massive volumes of acoustic data. For example, deep learning systems can detect chainsaws, bird calls, or species-specific vocalizations across diverse ecosystems (Stowell, 2022; Sethi et al., 2020). These developments align with the rise of soundscape ecology, which considers biophony (biological sounds), geophony (non-biological natural sounds), and anthrophony (human-generated sounds) as integral components of ecological systems (Pijanowski et al., 2011).

However, as Jennifer Gabrys (2016) warns, the sensor society also transforms citizens into unwitting participants in ambient data collection. In urban contexts, AI-equipped microphones in smartphones, surveillance systems, and smart cities capture everything from dawn choruses to protest chants. These machines do not simply hear—they mediate what becomes knowable and actionable within ecological and political frameworks.

1.2 Human Auditory Ecology and the Ethics of Listening

The rise of machine listening invites renewed attention to *human auditory ecology*—the study of how people perceive, interpret, and are shaped by their sonic environments. R. Murray Schafer’s notion of *schizophonia*—the separation of sound from its source—takes on new meaning as AI systems extract and analyze disembodied acoustic traces. Barry Truax’s emphasis on sound as a communicative, context-dependent phenomenon (Truax, 2001) contrasts with the reductionism of AI, which processes sound as signal, often devoid of cultural or emotional resonance.

Auditory neuroscience supports this critique: listening is embodied, affective, and multisensory (Zatorre et al., 2007). Jonathan Sterne’s concept of perceptual technics—where hearing is abstracted for algorithmic optimization—highlights how machine listening can prioritize clarity, efficiency, and extractive insight over relational depth (Sterne, 2012). Brandon LaBelle adds that listening is territorial and political, shaped by place and power. Together, these perspectives underscore a key concern: machine listening displaces context-rich human engagement with instrumental, and potentially intrusive, algorithmic hearing.

1.3 From Ecological Insight to Ethical Dilemma

The same infrastructure that enables ecological insight may also enable surveillance. When AI systems capture speech, emotion, or identity from ambient sound, they risk violating privacy—especially in public or contested spaces. Listening, in this context, becomes a political act. As Shoshana Zuboff (2019) and Ruha Benjamin (2019) argue, systems designed for benign purposes can reinforce surveillance capitalism and racialized listening hierarchies.

This paper responds to these tensions by interrogating how AI reconfigures the ethics of listening in ecological contexts. We argue that AI-driven soundscape monitoring must not only be technically effective but also ethically grounded—respecting privacy, context, and the pluralistic meanings of sound. Our goal is to advance human auditory ecology by designing systems that listen with, not just to, the world.

2 Context and Prior Work

2.1 Soundscape Ecology and the Politics of Listening

Soundscape ecology has emerged as a distinct field that integrates ecological, acoustic, and sociocultural dimensions of environmental sound. Pijanowski et al. (2011) proposed a tripartite framework—biophony, geophony, and anthrophony—to understand how all sounds in a landscape contribute to ecological meaning. This expands traditional bioacoustics, which focused on species-specific signals, to encompass entire acoustic environments as indicators of ecosystem health and human activity.

Building on Schafer’s (1994) foundational work on tuning the world, Truax’s (2001) concept of communicative listening, and studies exploring human perceptions of soundscapes (Smith & Pijanowski, 2014), contemporary scholars emphasize that sound is never neutral—it is shaped by perception, power, and place. LaBelle (2010) frames listening as a territorial practice, where soundscapes become arenas of inclusion and exclusion. Sterne’s (2012) critique of perceptual technics reveals how machine listening often abstracts sound from its social and bodily context in the pursuit of optimization. These perspectives are crucial for understanding how AI systems reshape auditory ecologies—not only by analyzing sound, but by mediating what is heard and by whom.

2.2 AI in Acoustic Monitoring

Advances in machine learning have enabled scalable acoustic analysis, with common data mining pathways now including the use of acoustic indices and deep learning models for automated sound source detection (Pijanowski et al., 2024). Deep learning models, particularly convolutional neural networks (CNNs) and attention mechanisms, are now widely used to detect species presence, identify illegal activity, and monitor ecosystem changes (Stowell, 2022; Sethi et al., 2020; Nieto-Mora et al., 2023). Autonomous recording units (ARUs), coupled with these models, allow for persistent monitoring even in remote locations.

However, these systems often rely on datasets shaped by dominant linguistic and ecological contexts, risking false generalization or cultural misclassification. Urban deployments, too, may record human voices, affective states, or culturally significant sound without consent—raising privacy and ethical concerns that are largely unaddressed in current practice. Moreover, systems optimized for accuracy often neglect questions of accountability, consent, and social impact.

2.3 Surveillance, Privacy, and Algorithmic Bias

AI-driven acoustic monitoring exists within broader infrastructures of data surveillance and behavioral extraction. Zuboff's (2019) concept of surveillance capitalism describes how technologies designed for benign purposes are repurposed to model and monetize behavior. In soundscape AI, this risk manifests in the incidental capture of human voices and emotional cues under the guise of environmental sensing.

Nautsch et al. (2019) show that even anonymized audio can contain rich biometric information—gender, emotion, location—making audio data especially vulnerable. Traditional notions of privacy and anonymity, rooted in visual and textual paradigms, are inadequate for the acoustic domain. Lyon (2018) and Benjamin (2019) further argue that surveillance systems often disproportionately impact marginalized communities, reinforcing racialized and spatial hierarchies in access and audibility.

2.4 Bridging Ethical and Technical Frameworks

While general AI ethics frameworks emphasize fairness, transparency, and accountability, few address the specificities of acoustic sensing. Listening is not merely data collection—it is a sensory, cultural, and relational act (Nissenbaum, 2009; Crawford, 2021). Thus, ethical soundscape monitoring must bridge computational models with social values such as privacy, data sovereignty, and community control (Gabrys, 2016).

This paper builds on emerging work in critical data ethics, privacy engineering, and participatory design (Costanza-Chock, 2020) to propose an integrated socio-technical framework. By foregrounding ethical commitments—acoustic privacy, participatory governance, environmental justice—and pairing them with feasible technical mechanisms such as edge computing, federated learning, and differential privacy, we aim to reimagine machine listening as a relational, pluralistic practice.

3 Ethical Challenges in AI-Driven Soundscape Monitoring

3.1 The Right to Acoustic Anonymity

AI models trained on conservation-biased datasets often misinterpret culturally diverse or marginalized sonic environments, as current research is heavily concentrated on datasets from a few countries like Australia, the United States, and Brazil, with a strong focus on specific species such as birds and anurans (Nieto-Mora et al., 2023). As Ruha Benjamin (2019) warns, the New Jim Code embeds racial and social hierarchies into seemingly neutral technologies. In soundscape AI, this can mean flagging religious chants, children playing, or regional speech as “anomalous” or “noise,” reinforcing exclusion and stigma.

Bias is not simply technical misclassification—it has spatial and political consequences. Communities near conservation zones may be unjustly monitored or displaced based on flawed acoustic assessments. Similarly, models trained on standard English speech may fail to recognize Indigenous, regional, or non-normative sound practices.

Algorithmic listening thus reshapes auditory hierarchies, determining whose voices are audible, credible, and actionable.

3.2 Function Creep and Surveillance Repurposing

While developed for biodiversity monitoring, these systems are vulnerable to function creep, as the stated goal of soundscape analytics is "the extraction of implicit, previously unknown, and potentially useful information from data" (Pijanowski et al., 2024), which can easily be repurposed for policing or behavioral tracking. Zuboff (2019) characterizes this as the logic of surveillance capitalism, where behavioral data is commodified for control or prediction. AI systems designed to detect chainsaws in forests can just as easily identify human voices, protest chants, or conflict events in cities.

David Lyon (2018) emphasizes that surveillance is no longer exceptional—it is normalized, participatory, and often hidden under socially accepted agendas like safety or sustainability. Soundscape systems risk becoming dual-use infrastructures, where environmental insight masks expansive data extraction. Without strong governance, AI listening may shift from ecological stewardship to ambient social control.

3.3 Algorithmic Bias in the Soundscape

AI models trained on Western, urban, or conservation-biased datasets often misinterpret culturally diverse or marginalized sonic environments. As Ruha Benjamin (2019) warns, the New Jim Code embeds racial and social hierarchies into seemingly neutral technologies. In soundscape AI, this can mean flagging religious chants, children playing, or regional speech as "anomalous" or "noise," reinforcing exclusion and stigma.

Bias is not simply technical misclassification—it has spatial and political consequences. Communities near conservation zones may be unjustly monitored or displaced based on flawed acoustic assessments. Similarly, models trained on standard English speech may fail to recognize Indigenous, regional, or non-normative sound practices. Algorithmic listening thus reshapes auditory hierarchies, determining whose voices are audible, credible, and actionable.

3.4 Community Consent and Data Sovereignty

Ambient listening captures more than sound—it captures context, culture, and sometimes ceremony, as social norms and identities shape the acoustic properties of soundscapes (Smith & Pijanowski, 2014). For Indigenous and underrepresented communities, sound is a medium of history, resistance, and meaning. Yet most AI deployments rely on individual, one-time consent models—insufficient for relational and collective sonic spaces.

The CARE Principles (Collective benefit, Authority to control, Responsibility, Ethics) offer an alternative rooted in Indigenous data sovereignty. These principles demand that communities—not institutions—set the terms of data use, access, and interpretation. Technologies like soundscape AI must embed these commitments: co-designed protocols, differentiated permissions, and revocable access. OCAP® (Ownership, Control, Access, Possession), developed in Canada, also provides governance frameworks

that affirm collective rights over acoustic data, especially when recordings capture culturally significant soundscapes (Konczi & Bill, 2024).

3.5 Environmental Justice and the Politics of Listening

Soundscape monitoring is often portrayed as ecologically neutral—but technologies never enter a vacuum. Environmental justice research (Bullard, 1990; Schlosberg, 2007) shows that marginalized communities disproportionately bear the burden of ecological harm and technological oversight. Acoustic sensors, if deployed without justice-centered design, may replicate these patterns: monitoring low-income or racialized areas while privileging elite “quiet zones.”

Barry Truax reminds us that soundscapes reflect power, a perspective reinforced by studies showing how social norms and identities influence the perception and management of acoustic environments (Smith & Pijanowski, 2014; Truax, 2001). When machine listening defines certain sounds as “pollution” and others as “natural,” it reshapes what is valued and protected. Sasha Costanza-Chock (2020) argues that design justice demands a shift from extractive to liberatory technology—centering the voices and priorities of those historically excluded from decision-making. In soundscape AI, this means respecting cultural sound, supporting community-led monitoring, and avoiding “noise” narratives that justify displacement or erasure.

4 Toward an Ethical Framework for Soundscape AI

Across these five domains—privacy, surveillance, bias, consent, and justice—AI-powered listening systems pose layered ethical risks. They amplify questions not just of what is heard, but who is allowed to listen, how sound is interpreted, and whose environments are reshaped by algorithmic ears. Without deliberate, participatory, and justice-oriented design, soundscape AI risks deepening existing inequities under the guise of environmental care.

This section proposes a normative framework for addressing these tensions—moving from critique to constructive ethical architecture.

4.1 Design Principles for Ethical Listening Systems

We propose five core principles for ethical soundscape AI: acoustic privacy, data minimization, transparency, participatory governance, and environmental justice. Each principle translates ethical concerns into actionable design and policy criteria, grounded in interdisciplinary perspectives from ecology, auditory science, and data ethics.

Acoustic Privacy. Listening systems must respect the contextual boundaries of sound. This includes avoiding the recording of private conversations in public spaces, excluding identifiable human speech where possible, and offering acoustic “opt-out” zones or filters. Inspired by Nissenbaum’s contextual integrity, privacy in soundscape AI should be understood not as data removal alone, but as the preservation of social and spatial norms around listening.

Transparency and Explainability. Acoustic AI models must offer intelligibility—not only to developers, but to communities and decision-makers affected by them. This includes transparency around what sounds are being collected, how they are interpreted, and what actions they trigger. Model explainability is especially crucial when outputs are used for regulatory or conservation decisions that affect human and non-human lives.

Participatory Governance. Communities should co-determine how soundscape AI is designed, deployed, and evaluated, building on policy approaches that consider public acceptability of soundscape management strategies (Smith & Pijanowski, 2014). This means engaging local stakeholders in defining relevant sounds, setting consent protocols, and interpreting results. Models of participatory sensing and community-led monitoring offer promising paths forward. Such approaches challenge the assumption that technical expertise alone suffices, and instead center cultural knowledge and lived experience as essential to ethical system design.

Environmental and Sonic Justice. Soundscape technologies must account for unequal power dynamics in how environments are heard, classified, and acted upon. This includes ensuring that “noisy” or “non-compliant” communities are not punished through acoustic policing, and that culturally significant sounds are not erased through narrow labeling schemes. Justice requires that benefits (e.g., conservation, urban planning improvements) and burdens (e.g., surveillance exposure, noise stigmatization) are equitably distributed.

4.2 From Principles to Practice

While these principles provide normative guidance, ethical soundscape AI requires implementation mechanisms that embed them from the ground up. Section 5 details a set of technical strategies—spanning edge computing, privacy-enhancing technologies, and federated learning—that enable ethical commitments to be realized without compromising core functionality.

Ultimately, this framework envisions soundscape AI not as a passive environmental monitor, but as a socio-technical infrastructure that listens with communities and ecologies—not simply to them.

5 Technical Overview: Operationalizing Ethical Listening

To translate ethical commitments into functional design, we outline a socio-technical architecture for privacy-preserving soundscape AI. This architecture balances ecological insight with community protection by incorporating technical mechanisms that enforce acoustic privacy, minimize unnecessary data capture, and distribute power in model development and deployment. We focus on four integrated components.

5.1 Edge Computing: Listening Without Storing

Edge computing, identified as a "new frontier in passive acoustic monitoring" where analytics are embedded in the sensor itself (Pijanowski et al., 2024), involves processing data locally and is a cornerstone of our proposed privacy-preserving architecture. In soundscape monitoring, edge devices can convert raw audio into labeled sound events (e.g., "bird call," "chainsaw," "car engine") without recording the actual waveform.

This offers several benefits:

- Privacy: Raw speech or identifying sounds never leave the device.
- Efficiency: Network bandwidth and energy usage are reduced.
- Context Sensitivity: Devices can adapt processing thresholds to environmental or community-defined parameters (e.g., ignore human voices or ceremonial music).

By aligning with data minimization and acoustic privacy, edge computing turns listening from an extractive act into a context-aware interaction.

5.2 Federated Learning: Decentralized Model Training

The standard machine learning pipeline in ecoacoustics centralizes training data by aggregating recordings from multiple locations into a single server for supervised model training (Nieto-Mora et al., 2023). This creates serious risks of exposure, surveillance, and misuse.

Federated learning (FL) reverses this model: each edge device trains on its own data and only shares model updates (not audio) with a central coordinator. These updates are then aggregated to improve the shared model.

Benefits include:

- Data Sovereignty: Communities retain control over their acoustic data.
- Scalability: New edge nodes can participate without disrupting the global model.
- Risk Reduction: Even if central servers are compromised, sensitive data remains secure

FL supports participatory governance, as local actors can influence training data quality and domain adaptation. It also allows for regional fine-tuning, reducing cultural or ecological bias in sound classification.

5.3 Differential Privacy: Protecting Identity in Aggregates

Differential privacy (DP) ensures that an individual's data cannot be reverse-engineered from aggregate outputs. It works by adding statistical noise to model updates or query responses—obscuring any single data point's influence without compromising overall accuracy.

In soundscape AI, DP can be applied at multiple levels:

- Model updates from devices can be privatized before aggregation.
- Analytics dashboards for researchers or policymakers can use differentially private summaries.

- Alert systems (e.g., illegal logging detection) can report patterns without exposing precise event locations.

DP reinforces contextual integrity and trust by offering mathematical guarantees of privacy—even under adversarial conditions.

5.4 Homomorphic Encryption: Secure Computation on Encrypted Data

Whereas DP protects data from inference, homomorphic encryption (HE) allows computation directly on encrypted inputs. This means that even a central server can process sound labels, anomaly scores, or environmental metrics without ever decrypting them.

While computationally intensive, recent advances make HE increasingly viable for low-dimensional or binary acoustic outputs (e.g., presence/absence of sound types). HE is especially valuable in scenarios where federated learning must operate across sensitive geopolitical or institutional boundaries.

Together with DP, HE ensures that “listening” does not imply “knowing”—preserving the rights of individuals and communities while enabling meaningful ecological analysis.

5.5 Integration and Governance

These four components are not isolated—they form a layered privacy architecture, where each mechanism compensates for others’ limitations. For example:

- Edge processing limits data exposure.
- Federated learning removes the need to centralize.
- DP prevents inference attacks on shared models.
- HE secures high-risk computations.

This layered approach is not only technically robust but also legible to communities. It allows stakeholders to make informed decisions about trade-offs between accuracy, responsiveness, and privacy.

Technical tools, however, are not a substitute for participatory governance. Communities must be consulted in defining what counts as “sensitive sound,” who has access to outputs, and how misclassifications are handled. Consent, once embedded in paper forms, must now be reimagined in acoustic protocols, model thresholds, and dashboard designs.

By embedding ethical principles directly into the technical stack, soundscape AI can evolve from a passive sensor network into an active collaborator in ecological and social care. The next section illustrates how this architecture can be applied in practice through a brief example, before concluding with broader reflections on participatory, justice-oriented design.

6 Case Scenario: Contextualizing Ethical Soundscape AI

6.1 Community Forest Monitoring: Participatory Conservation in Practice

Imagine a community-managed forest in Southeast Asia, where illegal logging and biodiversity loss are pressing concerns. A conservation NGO proposes to deploy AI-powered acoustic sensors to detect chainsaw activity and monitor avian populations. However, the forest is also home to an Indigenous community whose daily life, rituals, and oral histories are deeply embedded in sound.

Rather than treating the deployment as a technical rollout, the NGO initiates a participatory design process. Community members co-identify “sensitive” sound categories—such as ceremonial music, family conversations, or children playing—which are to be excluded from model training and output logs. Edge devices are programmed to recognize and ignore these categories in real time.

Federated learning is used so that audio data never leaves the forest; only anonymized model updates are shared. Differential privacy is added to ensure that no single update reveals culturally identifiable patterns. The resulting system successfully identifies chainsaw events and helps track endangered bird species, while upholding the community’s right to acoustic sovereignty.

The project evolves into a co-management platform, where ecological monitoring supports both conservation and cultural continuity. This scenario demonstrates how ethical soundscape AI can be relational, situated, and responsive—not merely technical infrastructure, but part of an ongoing dialogue between humans, machines, and environments.

6.2 Urban Surveillance and Acoustic Policing: Rethinking Smart City Listening

In a major metropolitan “smart city” district, municipal authorities deploy AI-enabled acoustic sensors on lamp posts to detect urban disturbances—gunshots, vehicle collisions, and noise ordinance violations. Though intended to improve public safety, the system begins recording ambient sounds in residential areas, including voices, arguments, and emotional outbursts.

Residents of a low-income, multi-ethnic neighborhood raise concerns: the system disproportionately flags their area for “excessive noise,” triggering fines and increased police patrols. Community groups note that culturally significant sound practices—such as evening music gatherings or street prayers—are misclassified as disturbances.

After public backlash, city planners initiate a redesign. Acoustic thresholds are recalibrated using input from local cultural organizations. A participatory audit identifies categories of “harmful misclassification,” leading to a retraining of the model using context-aware labeling. Edge devices are updated to filter private speech before classification. Real-time dashboards are introduced with opt-in community access to monitor flagged events.

This scenario highlights how urban deployments of soundscape AI can reinforce spatial and racial inequities if left unchecked. It also demonstrates the potential for participatory feedback loops, privacy-aware design, and community-defined

accountability to shift the system toward fairness. In dense, pluralistic urban soundscapes, ethical listening must be as dynamic and responsive as the environments it seeks to interpret.

7 Discussion and Conclusion

This paper has examined the ethical challenges and design imperatives for AI-powered soundscape monitoring in ecological and social contexts. By centering the principles of acoustic privacy, data minimization, transparency, participatory governance, and environmental justice, we have proposed a normative and technical framework that seeks to make machine listening not only more accurate but more accountable.

The proposed architecture—grounded in edge computing, federated learning, differential privacy, and homomorphic encryption—shows that privacy and ecological insight are not mutually exclusive goals. With thoughtful design, it is possible to build systems that serve both environmental and community interests without falling into extractive or surveillant logics.

However, technical safeguards alone are not sufficient. Ethical listening systems require ongoing negotiation: with local knowledge holders, policymakers, designers, and the communities who live within the soundscapes being monitored. Listening, as Schafer and Truax have argued, is always cultural, embodied, and situated. Machine listening systems must therefore engage not only with what is acoustically detectable, but also with what is socially meaningful.

Crucially, this also means recognizing the politics of audibility: whose voices are recorded, whose sounds are flagged, and whose environments are treated as worthy of attention. Without active efforts to redistribute power in the design and governance of these systems, soundscape AI risks reinforcing the very hierarchies it claims to monitor.

This paper contributes to the field of human auditory ecology by repositioning AI not as a neutral tool of environmental management, but as a participatory infrastructure—capable of reshaping relationships between humans, machines, and the ecologies they share. As listening becomes increasingly automated, the question is not just what we hear, but how we choose to listen—and for whom.

Future research should explore region-specific adaptations of this framework and empirical evaluations of acoustic privacy in practice, while also fostering the collaboration networks and shared datasets needed to create more equitable and robust models, a current gap in the field (Nieto-Mora et al., 2023). Importantly, the development of soundscape AI must remain attuned to the plurality of listening practices across cultures, disciplines, and species.

Ethical soundscape monitoring is not merely a technical challenge—it is a civic, ecological, and political project. Designing systems that can truly listen—without violating, extracting, or silencing—is not only possible, but necessary.

Disclosure of Interests. The author has no competing interests to declare that are relevant to the content of this article.

References

1. Benjamin, R. (2019). *Race after technology: Abolitionist tools for the new Jim code*. Polity.
2. Bullard, R. D. (1990). *Dumping in Dixie: Race, class, and environmental quality*. Westview Press.
3. Costanza-Chock, S. (2020). *Design justice: Community-led practices to build the worlds we need*. MIT Press.
4. Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.
5. Gabrys, J. (2016). *Program Earth: Environmental sensing technology and the making of a computational planet*. University of Minnesota Press. <https://doi.org/10.5749/j.ctt1b7x5gq>
6. Konczi, A.E., Bill, L. (2024). Advancing First Nations Principles of OCAP®. In: Garvey, G. (eds) *Indigenous and Tribal Peoples and Cancer*. Springer, Cham. https://doi.org/10.1007/978-3-031-56806-0_8
7. LaBelle, B. (2019). *Acoustic territories: Sound culture and everyday life*. Continuum. <https://doi.org/10.5040/9781501336225>
8. Lyon, D. (2018). *The culture of surveillance: Watching as a way of life*. Polity Press.
9. Nautsch, A., et al. (2019). Preserving privacy in speaker and speech characterisation. *Computer Speech & Language*, 58, 441-480. <https://doi.org/10.1016/j.csl.2019.06.001>
10. Nieto-Mora, D. A., Rodríguez-Buritica, S., Rodríguez-Marín, P., Martínez-Vargaz, J. D., & Isaza-Narváez, C. (2023). Systematic review of machine learning methods applied to ecoacoustics and soundscape monitoring. *Heliyon*, 9(9), e20275. <https://doi.org/10.1016/j.heliyon.2023.e20275>
11. Nissenbaum, H. (2009). *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press.
12. OCAP®. (n.d.). First Nations principles of OCAP. <https://fnigc.ca/ocap-training/>
13. Pijanowski, B. C., Farina, A., Gage, S. H., Dumyahn, S. L., & Krause, B. L. (2011). Soundscape ecology: The science of sound in the landscape. *BioScience*, 61(3), 203–216. <https://doi.org/10.1525/bio.2011.61.3.6>
14. Pijanowski, B. C., Rivas Fuenzalida, F., Banerjee, S., Minghim, R., Lima, S. L., Bowers-Sword, R., ... & Savage, D. (2024). Soundscape Analytics: A New Frontier of Knowledge Discovery in Soundscape Data. *Current Landscape Ecology Reports*, 9, 88-107. <https://doi.org/10.1007/s40823-024-00101-9>
15. Schafer, R. M. (1994). *The soundscape: Our sonic environment and the tuning of the world*. Destiny Books.
16. Schlosberg, D. (2007). *Defining environmental justice: Theories, movements, and nature*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199286294.001.0001>
17. Sethi, S. S., et al. (2020). Characterizing soundscapes across diverse ecosystems using a universal acoustic feature set. *Proceedings of the National Academy of Sciences of the United States of America*, 117(29), 17049–17055. <https://doi.org/10.1073/pnas.2004702117>
18. Smith, J. W., & Pijanowski, B. C. (2014). Human and policy dimensions of soundscape ecology. *Global Environmental Change*, 28, 63-74. <https://doi.org/10.1016/j.gloenvcha.2014.05.007>
19. Sterne, J. (2012). *MP3: The meaning of a format*. Duke University Press.
20. Stowell D. 2022. Computational bioacoustics with deep learning: a review and roadmap. *PeerJ* 10:e13152 <https://doi.org/10.7717/peerj.13152>
21. Truax, B. (2001). *Acoustic communication* (2nd ed.). Ablex Publishing.

22. Zuboff, S. (2019). The age of surveillance capitalism: The fight for a human future at the new frontier of power. PublicAffairs.