Organised Sound (2025), **30**, 375–382
doi:10.1017/S1355771825100897

CAMBRIDGE UNIVERSITY PRESS

Article

Generative soundscapes in exhibition contexts: Methods, implementation and perspectives

Polina Stolbova[1] and Gleb Rogozinski[2]

[1]Faculty of Audiovisual Arts, Russian State Institute of Performing Arts, Russian Federation and [2]Scientific and Methodological Laboratory Music and Computer Technologies, Herzen State Pedagogical University of Russia, Russian Federation

Abstract

Generative soundscapes in exhibition spaces offer new possibilities for integrating artistic practice, technological innovation and perceptual experience. Contemporary tools – including stochastic algorithms, random oscillators and diverse methods of sound synthesis – enable the construction of environments that respond dynamically to external conditions. With the integration of artificial intelligence and machine learning, such systems acquire additional flexibility: they are able to register the presence and movement of visitors, evaluate changes in audience density and adjust to the acoustic properties of the space in real time. As a result, sound layers can emerge when a participant approaches, the balance of elements may shift with fluctuations in the crowd, and potential peaks in volume can be anticipated and mitigated. In this way, a fixed soundtrack is transformed into an adaptive system, where the exhibition environment functions as an active, responsive organism. Sound ceases to serve merely as a background and becomes a structural component that directly influences how the artistic work is perceived.

Keywords: adaptive audio environments; context-aware sound design; Csound language; generative soundscapes; probabilistic modelling

(Received 21 July 2025; revised 30 August 2025; accepted 4 September 2025; First Published online 24 February 2026)

Introduction

A soundscape may be understood as the constellation of sonic events that unfold in natural or constructed settings, weaving a temporal and spatial fabric of sound. Although such environments exist independently, their perception and cultural resonance emerge only through human engagement. The concept of the soundscape highlights this dynamic link between sound and place, shaping auditory experience as well as broader ways of sensing and interpreting space (Southworth 1969; Schafer 1977).

Later research extended this view. Samuels (2010) describes soundscapes as a productive lens for cultural and social analysis, positioning sound as both a spatial and social phenomenon with distinct forms of circulation and meaning. In musicology, scholars such as Dudley (2002) and Manuel (1994) adopt the notion more broadly, connecting it to performance practice and musical culture rather than limiting it to acoustic conditions. Within audio art and sound studies, the soundscape serves as a conceptual meeting point where technology, aesthetics and cultural interpretation converge.

Soundscapes may be documented through field recordings or synthetically generated for artistic ends, appearing in installation art, film sound design, interactive media and generative music (Dunn 1996). This dual orientation establishes the term both as a category of analysis and as a medium of creative practice.

Recent advances in computer music and digital technology have greatly widened the scope for composers, sound designers and developers, pushing sonic practice beyond the constraints of traditional instruments. Extended reality (XR) – including virtual (VR) and augmented (AR) environments – adds further dimensions to artificial sound worlds by placing listeners in spaces where auditory perception is no longer tied to physical acoustics. Immersive three-dimensional audio, combined with synthesised timbres, enables the creation of complex environments shaped directly by technological innovation.

The interplay between natural and synthetic soundscapes opens new ways of understanding how organic processes intersect with constructed systems. Present models still fall short of full acoustic realism, mainly because existing algorithms lack sufficient micro-temporal precision. Yet even with these limits, they signal broader cultural and technological shifts. Within this framework, sound is not only something perceived but also a means of shaping experience and, at times, of reconfiguring how reality itself is sensed.

In current practice, soundscapes move beyond their role as cultural objects and function as domains of artistic and technological experimentation. In digital contexts, they are increasingly produced and transformed through algorithmic processes and synthesis methods, enabling real-time simulation in interactive media, audiovisual installations and generative composition. This expanding field positions sound as an adaptive and dynamic dimension at the meeting point of art, technology and perception.

A review of approaches to generating soundscapes

In their early work, Mynatt et al. (1998) presented one of the first automated approaches to soundscape generation with the *Audio Aura* project. Later, Birchfield (2005) proposed a probabilistic method of selecting audio tracks to produce environments that

Corresponding author: Polina Stolbova; Email: stolbova.2001@mail.ru

Cite this article: Stolbova P and Rogozinski G. (2025). Generative soundscapes in exhibition contexts: Methods, implementation and perspectives. *Organised Sound* **30**: 375–382, https://doi.org/10.1017/S1355771825100897

© The Author(s), 2026. Published by Cambridge University Press.

https://doi.org/10.1017/S1355771825100897 Published online by Cambridge University Press

evolve gradually over time. Finney and Janer (2010) worked with unstructured audio databases, while Shevade and Sundaram (2003) explored annotation systems and semantic models for organising sound data. Serafin and Serafin (2004) examined how sound design can influence presence and immersion in virtual environments.

Today, soundscapes are applied in diverse domains including music, film, theatre, podcasts and exhibition practice. In exhibition settings, generative sound systems often create adaptive spaces that respond to visitors' presence and movement. Despite their potential, such systems still face technical limits, especially when aiming for convincing realism.

A number of artists have pursued related directions. Christian Marclay, for instance, challenges conventional boundaries between sound and visual art through reproduction, fragmentation and spatial rearrangement of audio within architectural contexts. His video installation *Crossfire* (2007) immerses audiences in a multichannel sonic collage that unsettles habitual relations between image and sound. By layering and displacing fragments, Marclay generates a dynamic interplay in which sound reshapes visual narratives and encourages multisensory engagement. This approach mirrors broader tendencies in contemporary sound art, where mutable sonic structures play a key role in shaping spatial and conceptual meaning.

Marclay's practice has also been examined in *Audio Culture: Readings in Modern Music*, where it is positioned within experimental sound art and postmodern strategies of appropriation (Cox and Warner 2004). His work demonstrates how sound can shift from a supportive role to a primary artistic medium, questioning distinctions between composer, performer and audience. These ideas resonate with generative and algorithmic soundscapes, where sonic form is not fixed but continuously shaped by technological processes and listener interaction. In this sense, Marclay provides a precedent for understanding sound as an adaptive and relational medium, linking early experiments with the emerging practices of new media.

Formats of sound delivery range from private listening through closed-back or in-ear headphones to public systems employing semi-directional speakers or full-range arrays. Each configuration functions not only as a means of distributing audio but as a system that shapes physical and perceptual engagement, prompting active interaction with the acoustic environment.

Generative acoustic models often rely on probabilistic algorithms combined with digital signal processing. Oscillators, noise generators and pseudorandom functions are integrated to produce evolving textures that unfold over time. Such systems may take the form of virtual orchestras, algorithmically structured scores or parametric modulations of rhythm and timing. Their aim is to sustain variation and prevent mechanical repetition while maintaining the impression of continuous, responsive sound.

These approaches seek to construct immersive environments in which listeners engage not only cognitively but also bodily, through constant interplay with their sonic surroundings. Yet their degree of realism remains constrained by current synthesis techniques, particularly in replicating the subtle micro-temporal features that define natural acoustic environments.

A synthetic soundscape model for artistic and analytical applications

This section presents a soundscape model grounded in set theory. It is intended as a universal framework, applicable both to intelligent generative systems and to algorithms that analyse and extract elements from an audio environment. Formalised models of this type remain relatively rare in soundscape research, where scholarly focus has typically emphasised musical or psychoacoustic aspects. Consequently, the structural organisation of sound phenomena has often been underexplored.

In both natural and synthetic contexts, a soundscape is made up of multiple sound objects that differ in duration, envelope, spectral profile and temporal placement. We can define the set of all unique sound objects as $S = \{s_i | i = 1, \ldots, N\}$,where each element s^i represents an individual object.

For artificially generated soundscapes, this set is finite – at least when considered within a defined time span, denoted as τ_s. To further formalise S, we divide it into three subsets corresponding to distinct time scales of sound.

Thus,

$$S = S_1 \cup S_2 \cup S_3$$

Each subset S^i corresponds to sound objects defined by a particular temporal scale. In this classification, we aimed to align our segmentation with the well-known framework proposed by *Curtis Roads in Microsound* (2001), where sonic events are divided into temporal zones based on their duration and perceptual characteristics:

1. S_1 – short-term sound objects (up to about 250 milliseconds). These include sounds with sharp attacks and brief duration, such as clicks, percussion hits or instantaneous impulses. They form the microstructure of the soundscape, adding rhythmic detail and dynamic energy;
2. S_2 – medium-term sound objects (roughly 250 milliseconds to 30 seconds). Objects in this range possess internal structure and create perceptible segments within the environment. Examples include fragments of speech, musical phrases or environmental sounds such as passing cars or footsteps;
3. S_3 – long-term sound objects (tens of seconds or more). This category covers continuous background layers, such as ventilation noise, rustling leaves or sustained synthetic textures. These elements provide an acoustic ground that frames spatial perception as a whole.

Each sound object s_k S can be represented as a quadruple of parameters:

$$s_k = (\tau, A(t), F(w), t_o),$$

where:
τ – duration,
A(t) – amplitude envelope as a function of time,
F (w) – spectral characteristic as a function of frequency,
t_0 – onset time.

A rigid classification of sound objects into fixed categories quickly proves restrictive, as their characteristics often overlap and exhibit hybrid qualities in practice. A more effective strategy is to employ fuzzy logic, which assigns degrees of membership rather than relying on simple inclusion or exclusion. This approach better reflects the continuity of sonic environments and enables more precise modelling of complex soundscapes, especially in generative and immersive audio systems.

https://doi.org/10.1017/S1355771825100897 Published online by Cambridge University Press

The proposed division by time scale should be understood as conceptual rather than strictly formal. Boundaries are often fluid: a short transient may fade into a long reverberation tail or gradual decay, approaching medium-term sounds. The guiding principle of the model treats sound objects as self-contained acoustic entities extending across their full duration with coherent structure. For instance, a recording of ocean surf may function as a long background layer while simultaneously containing shorter events, such as seagull calls. These examples illustrate that sound objects can be both continuous and cyclical, composed of nested elements that serve as intermediate structures linking micro-and macrolevels of a soundscape.

Within each time span τ_s, the resulting time layer constitutes an independent acoustic plane of the soundscape. Short-term objects typically occur more frequently than long-term ones, resulting in a denser population of elements within their layer.

Formalising sound objects and their subsets allows not only systematic description but also structured design of soundscapes, whether natural or artificial. This framework opens new avenues for generative audio systems and algorithmic analysis, bridging artistic and scientific perspectives.

To accommodate the inherent ambiguity in classifying sound objects, a fuzzy membership function can be applied:

$$\mu[S_j(s_i)]: S \rightarrow [0, 1]$$

This expresses the degree to which a sound object s_i belongs to the subset S_j, allowing smooth transitions between categories and a more nuanced description of complex structures.

Software architectures supporting generative sound environments

The process typically begins with field recordings capturing natural sounds, human-made noises or isolated elements such as water and wind. These recordings provide a foundational layer that can be processed, transformed or combined with synthesised material to create a distinctive auditory environment. The source area defines the intended set of sounds and serves as a reference model, while the target area is analysed to identify key sonic features, including unwanted noises, which may be masked or supplemented with synthetic components.

Using synthesis algorithms, sound effects and procedural generation techniques, composers and sound designers can craft immersive, dynamic soundscapes that evoke particular moods, emotions or atmospheres. Available tools range from simple one-click applications on portable devices to sophisticated programming environments and synthesis engines, providing flexible means to construct artificial auditory worlds.

Although many soundscapes can be produced within digital audio workstations, for example *Ableton Live* or *Logic Pro*, fully generative structures often demand more specialised platforms. Computer music languages including *Max/MSP*, *Pure Data*, *SuperCollider*, *ChucK* and *Csound* provide precise control over synthesis and signal processing, making it possible to design evolving, nonlinear and interactive sonic environments. These platforms expand creative potential, enabling both artists and researchers to develop highly customised and responsive audio systems.

A key advantage of such environments is their capacity for real-time interactivity, enabling soundscapes to respond dynamically to user input, environmental conditions or external data streams. This responsiveness enhances immersion and encourages listeners to engage actively with the auditory experience.

Despite these capabilities, current synthesis algorithms remain limited in fully replicating the complexity of natural acoustic environments, particularly in capturing micro-temporal details and subtle spatial cues crucial for convincing realism.

Many computer music languages employ a modular, patch-based architecture, assembling soundscapes from independent blocks or objects, each responsible for generating or processing specific sounds. This modularity allows creators to experiment freely, combining elements to explore novel sonic possibilities. Such systems excel in procedural sound generation, producing evolving textures driven by algorithms, mathematical models and stochastic processes. This approach is especially valuable in interactive installations and virtual environments, where soundscapes must adapt to user behaviour and surroundings, delivering context-sensitive, dynamic auditory experiences.

These languages also permit simultaneous control over multiple parameters, including pitch, timbre, volume, spatial position and other sonic features. Detailed manipulation of these aspects allows designers to construct rich, layered audio environments that evoke strong emotional responses and fully immerse listeners. By experimenting with parameter combinations, creators can refine both the depth and expressive qualities of the soundscape.

Artificial soundscapes are sometimes integrated with real-world environmental data. For example, an installation simulating a tropical forest might synchronise its audio with live weather information, enhancing both immediacy and realism. Spatial audio techniques further increase immersion by positioning sounds within three-dimensional space, creating the impression that they originate from specific directions and distances. Common approaches include binaural recording, which produces surround effects for headphone listening, and ambisonics, optimised for multi-speaker arrays. Once individual sounds are generated or recorded, they are mixed and mastered to form a coherent, harmonious soundscape that conveys a strong sense of place.

Nonetheless, even with these sophisticated tools, reproducing the fine-grained, micro-temporal nuances that characterise natural acoustic environments remains a significant challenge for contemporary synthesis and spatialisation techniques.

Application of Csound for artificial soundscape creation

The objective of this modelling approach is to create an immersive soundscape that also communicates meaningful information about the environment to the listener. Achieving this requires balancing two key considerations: maintaining acoustic realism while avoiding repetitive elements that could appear artificial and undermine the credibility of the auditory experience.

In this study, the model is implemented using *Csound*, a versatile platform for sound synthesis and processing. However, the underlying principles are not restricted to *Csound* alone – similar strategies can be adapted to other sound engines or programming environments capable of managing dynamic, generative sound structures.

Csound is widely recognised among computer music languages as a mature and flexible tool for both synthesis and algorithmic composition (Lazzarini 2012; Yi 2015). Its modular architecture and long development history make it popular in both academic research and artistic practice, offering precise control over synthesis parameters and supporting complex generative sound

https://doi.org/10.1017/S1355771825100897 Published online by Cambridge University Press

design. Core techniques, including granular synthesis, spectral resynthesis and physical modelling, are thoroughly documented in the *Csound Book* (Boulanger 2013), which remains a central reference for practitioners and researchers.

Game engines such as *Unity* are increasingly employed to create interactive sonic environments, providing advanced spatialisation features and real-time audio integration (Pablo et al. 2024). Despite these capabilities, algorithmic and synthesis approaches still face challenges in fully reproducing the micro-temporal intricacies and spatial richness that characterise natural acoustic environments.

Randomness plays a crucial role in constructing sound systems that incorporate elements of unpredictability. In *Csound*, stochastic processes can be applied at multiple levels – from simple noise generation opcodes to complex generative patches that produce continuously evolving and dynamic sonic structures.

Csound also supports *User-Defined Opcodes (UDOs)*, allowing users to build custom sound generators and processing modules, which further expand generative possibilities. A wealth of examples and tutorials illustrating these techniques is available, for instance, on Iain McCurdy's Csound Resources. Available at: http://iainmccurdy.org/csound.html.

Beyond basic opcode usage, higher-level frameworks such as *CsoundAC* provide more structured environments for algorithmic composition and adaptive soundscape creation. These tools enable the design of dynamic, evolving sonic systems that strive to maintain organic variability while minimising perceptible repetition (Boulanger 2013).

At the same time, it is important to acknowledge that, although randomness introduces valuable variation, accurately capturing the fine temporal details and spatial complexity of real acoustic environments remains a significant challenge for current generative systems.

Generation of random events and noise

The generation of random processes and noise is fundamental to the construction of generative sound systems. In *Csound*, this is achieved through specialised operators – opcodes – that model both natural stochastic phenomena and formal mathematical representations of uncertainty.

Csound provides a wide range of pseudorandom generators and noise signals that form the basis for stochastic modulation and the creation of evolving sound textures. These include uniform white noise, frequency-limited noise, stepwise or smoothly interpolated noise, and Gaussian-distributed generators, which often produce more natural, organic effects. Advanced opcodes, such as *jspline* and *rspline*, enable gradual parameter changes, offering precise control over variability. Such tools can be used to modulate filter parameters, amplitude, pitch and rhythm, generate random musical phrases or simulate natural acoustic processes.

Random number and event generators produce discrete or stepped values with various distributions, such as uniform or Gaussian. Examples include *rand, randh, randi* and *gauss*, which are commonly employed for stochastic modulation of synthesis parameters to introduce unpredictable variation over time. By contrast, noise generators produce continuous signals with defined spectral characteristics, including white, pink and brown noise. Opcodes such as noise, *pinknoise* and *brownnoise* generate these signals, which may function either as independent sources or as textural elements within a composition.

This distinction underscores the different roles that stochastic processes play in synthesis. Noise and random sequences are commonly employed in subtractive synthesis as rich spectral material and to model natural textures such as wind, rain or steam. The capacity for gradual, fine-grained control – achieved through opcodes like *jspline* and *rspline* – enables sound designers to craft nuanced, organic variations, thereby enhancing both realism and expressiveness in generative soundscapes.

In digital sound synthesis, sources are generally classified into several categories. Deterministic generators, such as oscillators, produce periodic signals based on predefined waveforms. Random sources generate noise or stochastic sequences, while samplers reproduce pre-recorded audio fragments. Sound design often involves combining multiple sources – oscillators, filters and amplitude or frequency modulators – to create complex and rich timbres. This framework accommodates diverse methods of sound production without prescribing specific implementations.

Within generative composition, oscillators can assume multiple functions: providing melodic or harmonic foundations, forming textural layers that contrast with noise or modulating the frequency and amplitude of other components. The integration of random number and noise generators in *Csound* further expands creative possibilities. By combining deterministic and stochastic sources, sound designers can simulate natural phenomena, generate evolving textures and construct flexible, algorithmic soundscapes.

Soundscape design using three approaches

Several approaches can be employed to construct soundscapes. A common method utilises a database of pre-recorded materials drawn from publicly available libraries. From this collection, sound objects contributing to the intended acoustic atmosphere are selected, while the system manages automatic playback and layering. In *Csound*, for example, the *diskin2* opcode is typically used, with files specified either within the instrument definition or directly in the score to control playback duration. Expressiveness can be enhanced through amplitude and frequency processing, introducing variability that renders reused source material unique. Signals may be mixed and accumulated using global variables, forming complex audio layers with overlapping elements.

The workflow generally involves preparing multiple *WAV* files within the project folder, creating instruments in the orchestra for each sound and assigning relevant parameters such as amplitude, frequency and onset times. A base instrument typically plays continuously to establish the foundational layer, while melodic lines and occasional or rare sounds are layered above to enrich the overall atmosphere. This structured approach supports both controlled composition and adaptive, generative audio behaviour.

Incorporating stochastic parameters further enhances the sense of liveliness and organic motion within the soundscape. Modulating frequency and amplitude with random elements introduces natural movement and unpredictability, a feature particularly important in generative or interactive systems, where continuous variation and dense textures are essential to shaping the auditory experience.

To implement a soundscape based on randomisation, the sound material must first be prepared. This may include *WAV* files stored in the *Csound* project folder or generative sounds produced using synthesis opcodes such as *oscil, noise* or *pinkish*. Instruments are then defined in the orchestra section, each responsible for a specific sound element. For audio files, accurate naming is essential, while generative sounds require correct synthesis commands.

https://doi.org/10.1017/S1355771825100897 Published online by Cambridge University Press

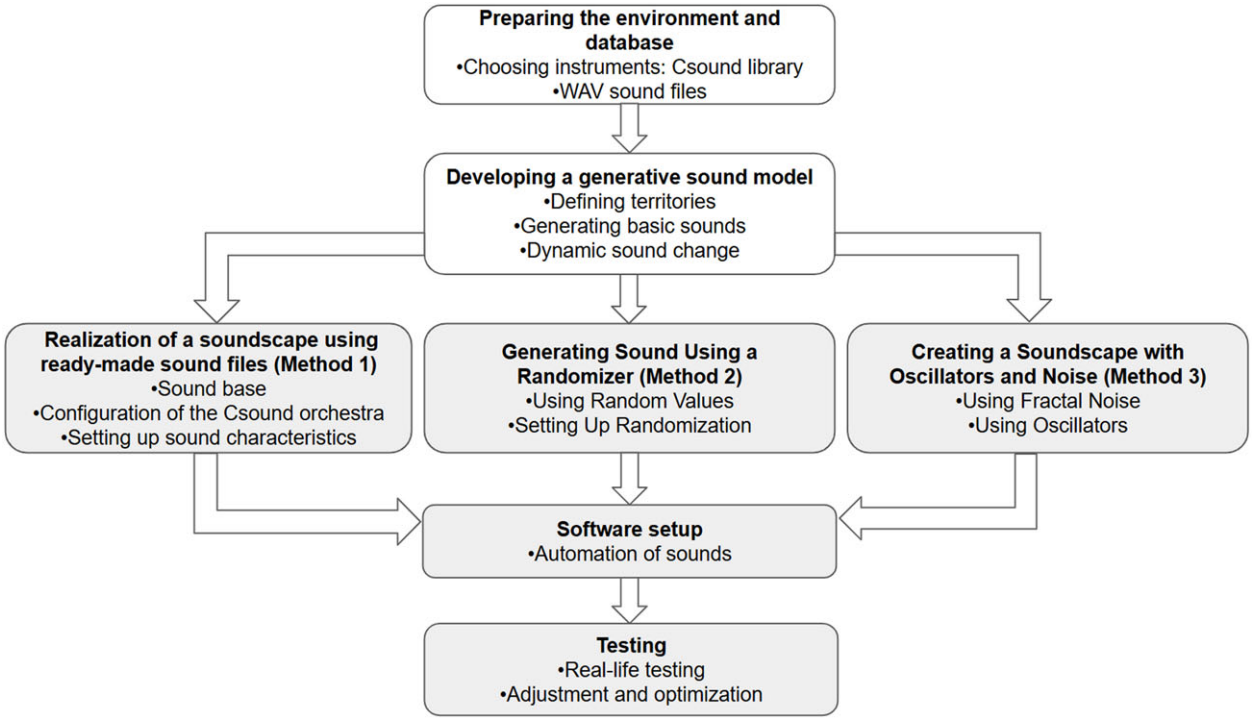


Figure 1. Complex of methods for generating soundscapes for an exhibition.

Randomisation parameters from the *rand* family control variations in aspects such as frequency, amplitude and stereo position, allowing subtle movements and nuanced changes. Within the orchestra, opcodes like rand, randi and randh produce time-based modifications – for example, random panning of noise or irregular volume shifts. The score specifies timing, duration and density of sounds, enabling either fixed or generative temporal structures.

A typical soundscape comprises a continuous base layer that establishes the general acoustic field, an active melodic line woven throughout, and sporadic details – accents or signals – that enrich the atmosphere and provide contrast. Noise-based structures often serve as foundational elements due to their broad spectral and temporal characteristics, which naturally resemble phenomena such as wind, rain or waterfalls. Digital processing allows these noises to be shaped into more abstract or stylised forms, tailored to artistic intentions or specific acoustic contexts. Combining multiple types of noise further expands expressive potential, producing complex, layered textures within the soundscape.

When building a soundscape from oscillators or noise, the process begins by preparing code for sound generation and creating the corresponding instruments in the *Csound* orchestra. Key opcodes, including oscil, noise and pinkish, are applied either by adapting documented examples or developing custom patches. Instrument parameters – such as amplitude, frequency and volume – are adjusted, and random elements introduced via rand opcodes add organic, dynamic variation, for instance, in the spatial movement of noise or changes in sound density.

The score defines the timing and duration of each instrument, structuring a soundscape that typically comprises a stable base layer as the background, a flowing melodic line and sparse details – accents or signals – that enrich the overall character and atmosphere.

Analysis of soundscape creation technologies highlights several key considerations. Three generator types were examined, each offering distinct advantages, though their suitability for mobile platforms differs. Pre-recorded audio tracks are straightforward to implement but complicate project transfer between devices, as both code and audio files must be moved together. This method preserves the project's structure and parameters while allowing substitution of individual sounds – for example, replacing surf recordings with wind or rustling leaves.

Noise-and oscillator-based generators provide greater flexibility. They can operate continuously, producing a living, ever-changing soundscape. Unlike looped playback, these approaches avoid repetition, and the inclusion of randomisation ensures virtually endless generation of unique, evolving structures. Such unpredictability brings the audio closer to natural acoustic experiences, making the environment feel dynamic and alive.

Within the project, a generative soundscape model was demonstrated using multiple implementation methods. A concept for a sound exhibition was proposed, where the audio in each room dynamically adapts to the visitor's location, fostering interactive engagement. The three main approaches – pre-recorded track playback, parameter randomisation and oscillator-and noise-based generation – were evaluated. Noise-and oscillator-based methods proved most effective for long-term exhibitions and interactive spaces, offering continuous, non-repetitive sound, with randomisation maintaining ongoing novelty and conveying the impression of a living, evolving acoustic environment.

Case studies in implementation: sound design for an exhibition of electronic synthesisers

The algorithm can be divided into several successive stages, each utilising different methods of sound generation – from utilising ready-made audio files to procedural and random algorithms. As an example, the *Csound* environment is considered, which allows combining various approaches within a single project.

https://doi.org/10.1017/S1355771825100897 Published online by Cambridge University Press

Based on observations and feedback, the sound generation parameters can be fine-tuned. This allows for greater expressiveness and adaptability of the sound landscape (Figure 1).

Thus, the exhibition's sound environment becomes dynamic and responsive to visitors. The atmosphere shifts from a passive backdrop to an active participant, interacting with the space and responding to its spatial and temporal rhythms.

Among the methods explored, the most expressive and technically adaptable proved to be purely synthetic sound generation via software code, without reliance on pre-recorded audio. This approach eliminates dependence on external media, simplifies transfer across different hardware setups and allows precise calibration of sonic parameters to the characteristics of the exhibition space.

Soundscapes constructed from oscillators and noise generators offer a high degree of variability. They evolve organically, producing fresh textures during each listening experience. This self-renewing quality is particularly valuable for long-term exhibitions, sustaining an open-ended, vibrant auditory environment.

Practical applications of fractal organisation in sound design have been documented – for instance, Rhys (2016) demonstrates its use in granular synthesis to create sound clouds with rich, evolving structures. Building on this principle of fractal self-similarity, the present project applies fractal modulation alongside oscillators and noise sources to generate an adaptive, continuously changing sonic landscape. Unlike sample-based granular approaches, this soundscape emerges entirely from continuous signal generators – such as oscil, noise and pinkish modules – providing a flexible foundation ranging from low, sustained bass drones to airy, high-frequency textures.

A key advantage of this approach is its ability to model fractal structures within sound. By applying filters, modulation and feedback loops to noise sources, it becomes possible to emulate natural acoustic phenomena – such as wind, rainfall or flowing water – without direct sampling. The resulting textures are layered and dynamically shaped in real time, producing a non-repetitive, evolving sonic environment with an organic complexity akin to that of natural soundscapes.

This method offers high adaptability, sonic richness and long-term stability. Constructing the soundscape entirely through real-time synthesis removes the constraints of fixed audio material, allowing continuous evolution without perceptible repetition. Combining oscillators, noise generators and fractal modulation not only expands the palette of textures but also permits precise control over their temporal and spatial characteristics. (Figure 2).

Integrating audience interaction transforms the soundscape from a static backdrop into an active element within the exhibition, capable of responding in real time to visitor movements and environmental changes. This responsiveness sustains engagement, while the generative system ensures that each visitor experiences a unique auditory environment.

Moreover, the absence of dependence on pre-recorded samples enhances portability, reproducibility and scalability across different venues and technological setups. As a result, this approach is particularly well-suited for long-term and site-specific installations, preserving freshness, adaptability and the conceptual coherence of the sound environment.

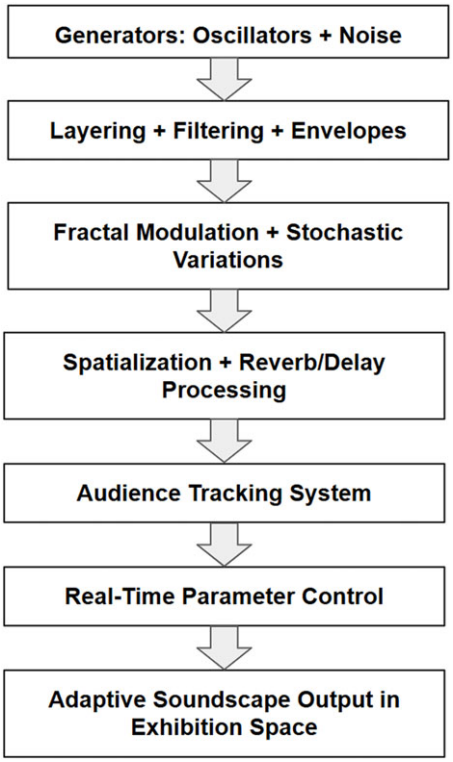


Figure 2. Conceptual flow of the system.

Examples of implementation. Sound design of the exhibition about musical synthesisers

For an exhibition devoted to music synthesisers, generative soundscapes offer an effective way to showcase the evolution of sound technologies while enhancing visitor engagement. The most productive strategy involves constructing the entire auditory environment through software code, without relying on pre-recorded samples. This approach provides flexibility, eliminates dependence on fixed media and allows the sound to adapt organically to both the acoustics of the space and the movements of the audience.

Soundscapes built from oscillators and noise generators are inherently dynamic, continuously evolving without exact repetition. They can run indefinitely, creating a rich, shifting atmosphere that embodies the spirit of experimentation and innovation central to electronic sound synthesis. Generative methods highlight not only the technical sophistication of synthesisers but also their artistic potential, turning each visit into a unique auditory journey that bridges the history and future of electronic music.

One of the earliest generative soundscapes was designed for a hall dedicated to the Hammond Organ. Its structure comprises several layers: a continuous pad forming the background, a melodic organ line, rhythmic percussion accents and spatial effects that enhance the perception of depth and immersion.

The background pad is produced by two smoothly enveloped oscillators, generating a stable and warm sonic field that sets the overall ambience. The melodic organ line is constructed through sequential frequency readings from a note table, replicating the instrument's characteristic timbre. This line is enriched by six oscillators generating complex overtones, further enhanced with reverberation to deepen the spatial impression.

Percussive elements are created from shaped bursts of random noise, introducing movement and rhythmic vitality. A rotating effect, implemented through frequency modulation controlled by a low-frequency oscillator, adds a sense of spatial motion and dimensionality.

https://doi.org/10.1017/S1355771825100897 Published online by Cambridge University Press

Together, these components form an immersive, multidimensional soundscape that reflects the conceptual focus of the exhibition and encourages deeper audience engagement through a carefully crafted sonic environment.

You can find the prototype code on GitHub[1] and listen to the soundscape online[2].

The soundscape creates an atmospheric and melodic environment that reflects the aesthetics of Soviet electronics and experimental synthesisers, such as the ANS. It is constructed from three instruments operating in parallel, producing a rich, multilayered audio texture.

The first instrument generates a smooth, evolving background by blending pink noise with low-frequency oscillators in the infrasonic and subsonic ranges. Carefully shaped amplitude envelopes provide gradual fade-ins and fade-outs, conveying the impression of a continuous, living sonic field.

The second instrument conveys the melodic electronic timbre characteristic of the Soviet synthesiser tradition. A sine oscillator slowly shifts between 220 and 330 Hz, with its tone softened through amplitude modulation and frequency filtering. Subtle panning toward the left channel adds spatial depth and a gentle sense of movement.

The third instrument produces a soft drone shaped by smooth frequency modulation, evoking the 'photo-optical beam' – a symbolic reference to the light-driven synthesis central to the ANS. Its restrained dynamics and delicate amplitude envelope contribute subtlety and sonic depth.

Together, these layers form an immersive, dynamic soundscape that highlights both the technical and artistic qualities of the ANS, providing a fluid, ever-changing accompaniment to the exhibition.

Including the ANS in this project is particularly significant, as it remains an innovative instrument often overlooked in Western academic and musical discourse. Developed by Soviet engineer Evgeny Murzin in the mid-20th century, the ANS employed a unique photo-optical synthesis method that converted visual images directly into audio signals. This approach, unparalleled in Western electronic music traditions, opens new avenues for exploring alternative models of sound generation.

By situating the ANS within this study, the project introduces valuable geographical and historical breadth, linking contemporary generative practices with lesser-known traditions of electronic experimentation. Its visual-to-audio synthesis expands the artistic vocabulary of generative sound beyond conventional algorithmic methods, inspiring novel approaches based on image transformation and stochastic processes.

Today, the ANS remains highly relevant to sound art and computer music research, particularly where visual data intersects with sonic creation. Its underlying principles continue to inspire algorithmic methods capable of generating unique, continuously evolving soundscapes, expanding the expressive possibilities of generative sound.

This project therefore demonstrates not only the technical achievements of the ANS but also its broader cultural and artistic significance. By connecting algorithmic sound generation with sensory perception, it situates the instrument within contemporary practice, opening pathways toward interactive and adaptive sonic environments that act as active co-creators in museum and exhibition spaces.

You can find the prototype code on GitHub[3] and listen to the soundscape online[4].

The creation of soundscapes for exhibition halls dedicated to classic electronic instruments, such as the Hammond Organ and the ANS, demonstrates how algorithmic techniques can shape distinctive auditory environments. Each soundscape reflects the character and technical features of its instrument: layered oscillatory structures with reverberation for the Hammond and textured noise combined with slow, evolving amplitude envelopes for the ANS.

This approach provides flexibility and adaptability, generating a continuously unfolding atmosphere that avoids repetition and does not rely on pre-recorded material. Programmability and parametric modulation allow the environment to be fine-tuned in real time, adapting to the specific acoustics of the space and the movement of visitors.

Algorithmically generated soundscapes thus not only enhance audience engagement with the exhibits but also act as autonomous artistic agents, adding a unique sensory dimension to museums and galleries.

Prospects

An essential factor in the effectiveness of a soundscape lies in how it is perceived by the listener. In this project, attention was directed not only to the technical mechanisms of sound generation but also to the qualities shaping emotional and cognitive responses. The variability and continuous evolution of the sonic environment create the impression of a living, breathing space, enhancing immersion and sustaining engagement throughout the visit.

Analysis of the sound material shows that combining stochastic processes with carefully tuned synthesis parameters conveys natural unpredictability while maintaining artistic clarity. Spatial processing and dynamic adaptation – through modulation, filtering and panning – further strengthen this effect, ensuring that each listener experiences a unique auditory environment.

Thus, the project not only demonstrates the technical capabilities of generative methods but also extends the artistic vocabulary of sound art. It links algorithmic approaches with sensory perception and opens new avenues for designing interactive, adaptive environments that act as active collaborators within the exhibition space.

Conclusion

Today, soundscapes – both natural and artificially constructed – occupy a space at the intersection of aesthetics, technology and cultural analysis. This project examines how stochastic algorithms, noise generation, oscillator models and subtractive synthesis contribute to the formation of flexible and expressive sonic environments. While the implementation was carried out in *Csound*, the underlying principles are applicable to any platform capable of detailed algorithmic sound design.

The introduction of randomness and noise enables the reproduction of the inherent unpredictability of natural soundscapes, imparting digital compositions with vitality and continuous flow. When combined with oscillators, filters and modulation,

[1]GitHub. The Hammond Organ Prototype Code. Available at: https://github.com/polinastolbova/The-Magic-of-Sound-The-History-of-Musical-Synthesizers-I-Soundscapes/blob/3ad2c1de42097d0f7971de81e0e6de48fcf0dc5a/The%20Hammond%20organ

[2]YouTube. The Hammond Organ Soundscape. Available at: https://youtu.be/TXBxWEXgMLM?si=QKelh3R_HI_65B5i

[3]GitHub. Soviet Electronic Prototype Code. Available at: https://github.com/polinastolbova/The-Magic-of-Sound-The-History-of-Musical-Synthesizers-I-Soundscapes/blob/3ad2c1de42097d0f7971de81e0e6de48fcf0dc5a/Soviet%20electronic

[4]YouTube. Soviet Electronic Soundscape. Available at: https://youtu.be/Shj3Bbl6B1s?si=vqTP2uRDC1hYE_Xm

https://doi.org/10.1017/S1355771825100897 Published online by Cambridge University Press

these techniques form a generative framework relevant not only to artistic practice but also to research areas such as sound modelling and audiovisual data interpretation.

Future directions include applying machine learning for adaptive real-time sound generation, employing spatial audio formats such as ambisonics and binaural reproduction to enhance immersion and developing interdisciplinary approaches that connect acoustics, urban studies, cultural geography and digital art.

Artificial intelligence and machine learning also open new possibilities for exhibition soundscapes. Modern algorithms can track visitor behaviour and acoustic conditions in real time – including movement, crowd density and variations in sound dynamics – allowing the audio environment to adapt continuously and produce atmospheres that are both responsive and unpredictable.

For example, a system may activate specific sonic layers as visitors enter particular zones, adjust overall loudness based on crowd size or anticipate movement to subtly modify the soundscape in advance. In this way, fixed audio tracks are replaced by flexible, reactive structures, ensuring that each visitor experiences a unique auditory environment.

Integrating AI into exhibition sound design enhances engagement by positioning sound as an active mediator between space, technology and audience. It enables self-renewing acoustic environments in which generative algorithms act as co-authors of the artistic process rather than passive tools. Consequently, algorithmic synthesis combined with AI allows the creation of dynamic, adaptive soundscapes that elevate exhibitions, establishing sound as a central medium of interaction and creative expression.

References

Birchfield, D., Mattar, N. and Sundaram, H. 2005. Design of a Generative Model for Soundscape Creation. *Paper presented at the International Computer Music Conference (ICMC)*, Barcelona, Spain.

Boulanger, R. (Ed.). 2013. *The Csound Book: Perspectives in Software Synthesis, Sound Design, Signal Processing, and Programming*. Second Edition. MIT Press.

Cox, C. and Warner, D. (Eds.). 2004. *Audio Culture: Readings in Modern Music*. New York: Continuum. https://archive.org/details/audioculturereadd0000unse

Dudley, S. 2002. Dropping the Bomb: Steelband Performance and Meaning in 1960s Trinidad. *Ethnomusicology* **46**(1): 135–64.

Dunn, D. 1996. *Music Language and Environment: Environmental Sound Works 1973–1985*. Minneapolis, MN: Innova.

Finney, N. and Janer, J. 2010. Soundscape Generation for Virtual Environments Using Community-Provided Audio Databases.

Lazzarini, V. 2012. *The Audio Programming Book*. Cambridge, MA: MIT Press.

Manuel, P. L. 1994. Puerto Rican Music and Cultural Identity: Creative Appropriation of Cuban Sources from Danza to Salsa. *Ethnomusicology* **38**(2): 249–80.

Mynatt, E. D., Back, M., Want, R., Baer, M. and Ellis, J. B. 1998. Designing Audio Aura. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*.

Pablo, P., Bravo, L., Fasciani, S. and Glette, K. 2024. Csound versus ChucK: Sound Generation for XR Multi-Agent Audio Systems in the Meta Quest 3 Using the Unity Game Engine. In M. Conceição, G. Campos, I. Salselas & J. Vieira (Eds.), *Proceedings of the Sound and Music Computing Conference 2024*.

Rhys, P. 2016. Smart Interfaces for Granular Synthesis of Sound by Fractal Organization. *Computer Music Journal* **40**(3): 58–67. MIT Press.

Samuels, D. W., Meintjes L., Ochoa A. M. and Porcello T. 2010. Soundscapes: Toward a Sounded Anthropology. *Annual Review of Anthropology* **39**(1): 329–345.

Schafer, R. M. 1977. The Soundscape: Our Sonic Environment and the Tuning of the World.

Serafin, S. and Serafin, G. 2004. Sound Design to Enhance Presence in Photorealistic Virtual Reality. *Proceedings of the International Conference on Auditory Display (ICAD 04)*, Sydney, Australia.

Shevade, B. and Sundaram, H. 2003. Vidya: An Experiential Annotation System. *1st ACM Workshop on Experiential Telepresence, in Conjunction with ACM Multimedia*, Berkeley, CA.

Southworth, M. 1969. The Sonic Environment of Cities. *Environment and Behavior* **1**(1): 49–70.

Yi, S. 2015. Composing with Csound: Thoughts on Workflow and Musical Practice. *Proceedings of the International Csound Conference*, St. Petersburg, Russia.

https://doi.org/10.1017/S1355771825100897 Published online by Cambridge University Press