DOI: 10.61782/fa.2025.0145

FORUM ACUSTICUM EURONOISE 2025

# ARTIFICIAL IRS IN DYNAMIC VIRTUAL ACOUSTICS FOR MUSIC EXPERIMENTATION

**Luna Valentin**[1*] **Celeste Betancur Gutiérrez**[1]
[1] CCRMA, Stanford University, USA


## ABSTRACT

This paper presents a method for generating highly customizable artificial impulse responses (IRs) made for the CAVIAR (Cave of Augmented Virtual and Interactive Audio Realities) system. These synthetic IRs power a real-time live convolution engine, enabling on-the-fly manipulation of virtual acoustic environments. The approach paves the way for innovative musical experimentation and performance and adapts seamlessly to the requirements of both the virtual acoustics system and the performer.

Customizing virtual acoustics environments allows musicians to perform in entirely artificial yet sonically rich and responsive spaces. This freedom allows them to experiment with unconventional acoustical parameters in real time, examining how shifts in reverberation characteristics and spatial cues influence their performance and creative process. Audiences are simultaneously and similarly immersed in the listening environment, perceiving the music as unfolding within a dynamically evolving acoustic space.

From two complementary perspectives—technical/acoustical and creative/artistic—this paper explores our system's capabilities through two case studies conducted at CCRMA, Stanford University. These studies illustrate how adaptive, synthetic IRs can transform both the performer's and the audience's experience, offering new frontiers for interactive and expressive musical performance.


*__Corresponding author__: luna@ccrma.stanford.edu.


**Keywords:** *virtual acoustics, performance, impulse responses, interactive musical performance*

## 1. INTRODUCTION

In June 2024, we premiered ”Through the looking glass”, a piece Composed by Celeste Betancur for two performers - Celeste Betancur Gutiérrez and Iran Sanadzadeh - and a virtual acoustics technician - Luna Valentin - with the idea of exploring the boundaries of perception and push them in unexpected ways. For this piece, we developed a preliminary approach to on-the-fly manipulation of acoustical spaces relying on existing tools at our disposal. The piece was performed on the CCRMA stage (Center for Computer Music and Acoustics, Stanford University), that happened to be equipped with the CAVIAR system. The CAVIAR system was developed at CCRMA initially to reproduced acoustical landscapes otherwise hard to access, and allow full immersion with maximum freedom in these spaces [1]. The system is made of microphones, whose signal is sent to a Max Patch performing multichannel convolution and feedback canceling, allowing for full immersion of both the public and the performers. This very realistic system seemed to offer a nice avenue for blurring the limits between natural and artificial acoustics. We therefore created, first, an algorithm that generates artificial impulse responses based on customizable parameters, and second implemented a system allowing us to manipulate sounds in a 3D space and swap impulse responses in real-time. We describe in this paper the collaborative process to make this piece, the techniques used and how using the space as an additional instrument can transform the way we think about acoustics and composition.

## 2. CUSTOM ARTIFICIAL IMPULSE RESPONSES.

The CAVIAR system operates by utilizing sets of impulse responses (IRs, measured or generated), with the number of channels varying based on the room configuration (number of speakers and number of microphones) [2]. Traditionally, the Statistically Independent Perceptually Identical (SIPI) method has been employed to generate more channels when the number of channels of the IR wasn't matching the number of channels of the CAVIAR. The SIPI method ensures that impulse responses remain independent, minimizing phase interference and other wave artifacts. This method was particularly meant to adapt measured impulse response to the CAVIAR. However, using measured IRs remains constraining and lacks flexibility. We wanted to go further and create more contrasting spaces.

To address this, we built an algorithm that efficiently creates sets of IRs tailored for the CAVIAR system. This approach ensures control over phase coherence/incoherence as wanted and perceptual consistency/inconsistency across channels. In this section, we first detail the algorithm developed for synthetic IR generation and, secondly, describe how we used CAVIAR's features to control the acoustical space in real-time.

### 2.1 Parametric Approach to IR Generation

First, we observe the behavior of Recorded Impulse Responses (RIR), see Fig. 1 in terms EQ, decay time and echo density. Secondly we mathematically mimic these physical parameters. Third we assess the amount of variation that can be observed in the RIR parameters previously mentioned. Lastly, we fine-tune the algorithm to replicate or surpass the variation observed in the RIRs physical parameters. As a result, the following parameters are variables to be determined by the user in the process of creating the impulse responses:

- Reverberation Time (T60): the overall decay duration of sound energy, defines the size and character of the virtual acoustic space.
- Echo Density Profile: how the density of reflections builds up over time, affecting perceived clarity and texture of the reverberation.
- Frequential Decay Shape: frequency-dependent attenuation
- Predelay: a brief delay before early reflections begin, which influences spatial perception and the perceived distance of sound source.
- Variability Across Channels: variation in multichannel configurations to control the perceived geometry

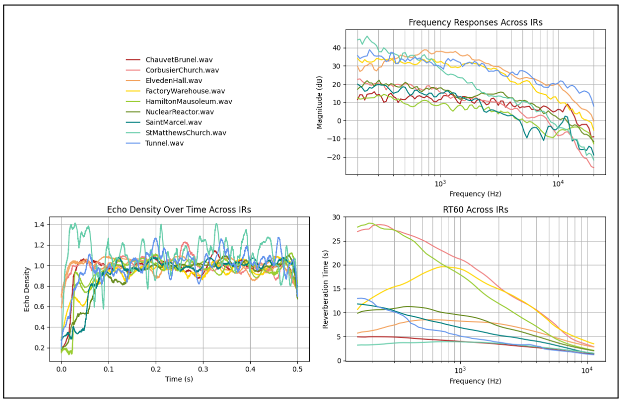


**Figure 1**. Analysis of a set of 10 very different recorded impulse responses.

The algorithm "Irreverberable" was implemented in Python 3.11 to generate impulse responses with the required characteristics for CAVIAR (customizable number of decorrelated channels), featuring customizable parameters as described above. The algorithm follows these steps:

1. Create an Echo Density Profile: A logistic growth function shapes the temporal evolution of echo density. The curve is tuned to stay within the range of previously measured impulse responses, with random variations and adjustable parameters. Two noise components are interpolated with the logistic growth function:
   (a) A low-frequency noise component introduces distinct groups of echoes, which may contribute to the perception of discrete acoustical reflections.
   (b) A high-frequency noise component replicates the unstable variations observed in real-world echo density curves.
2. Distributing Echoes: Once the echo density curve is generated, it determines when discrete echoes occur in time. These echoes are distributed based on cumulative density probability and modulated

to ensure natural scattering, producing a result that perceptually resembles white noise.

3. Frequency-Dependent Exponential Decay: The algorithm shapes the tail of the impulse response by applying frequency-dependent exponential decay, reaching -60 dB around 1000 Hz at the time specified by the user and extending slightly longer at lower frequencies. Each frequency band is assigned a random decay factor, decreasing with frequency to mimic real-world absorption phenomena.

4. EQ Shaping: While applying frequency-dependent decay, the algorithm also performs equalization (EQ) shaping by introducing random gain variations across different frequency bands.

5. Adding Pre-Delay: A pre-delay is introduced to the impulse response to simulate spatial perception, ensuring realistic localization effects.

6. Final Processing and Export: Before exporting, impulse responses are processed with a high-pass filter and normalized to prevent low-frequency feedback and ensure consistency.

The algorithm produces artificial IRs (Fig. 2) that were parameterized to physically resemble the set of 10 measured room impulse responses (RIRs) previously described (Fig. 1).

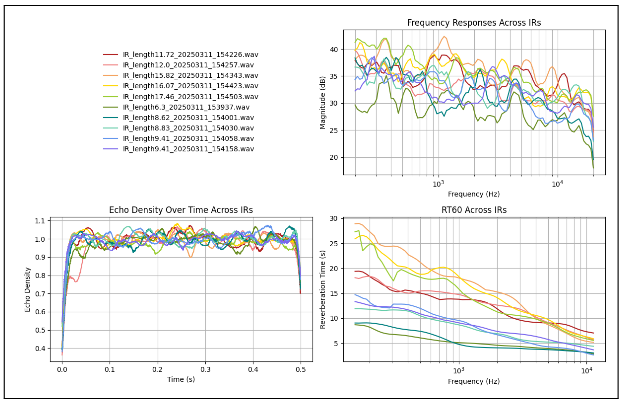


**Figure 2**. Analysis of a set of 10 generated impulse responses.

Although these synthesized IRs are designed to sound natural, they can be manipulated to emphasize artificial qualities, effectively blurring the boundary between real-world acoustics and algorithmically generated reverberation. This flexibility enables both realistic spatial simulations and creative sound design applications, allowing for exploration of spaces that do not physically exist.

### 2.2 Real-Time Adaptability, Implementation in the CAVIAR System

The CAVIAR system, was recently augmented by Hassan Estakhrian, who's long been working on virtual acoustics systems at CCRMA [3], with a function to enable real-time switching between impulse responses (IRs). The CAVIAR Max/MSP patch now allows seamless crossfading between two IRs within a user-defined time frame. This functionality enabled dynamic transitions between multiple IRs during the performance. These transitions could be pre-programmed as part of a timed script or executed live by the sound engineer, who could respond to the performance in real time—either following a score or improvising.

In addition to CAVIAR, a spatialization patch was employed to manipulate sonic objects, both pre-recorded and generated in real time by the computer music performer. The sound engineer had the ability to spatially position these sounds in response to the performers' movements, either mirroring their choreography or creating contrasting spatial counterpoints, thereby reinforcing or challenging the perceived relationship between sound and gesture.

## 3. *THROUGH THE LOOKING GLASS*, OR HOW TO PUSH THE BOUNDARIES OF PERCEPTION.

*Deep into the rabbit hole, is there a way out? Some doors will take you in, some others out, and others will loop you around.*

"Through the Looking Glass" is a piece composed for two performers or two groups of performers, positioned to face each other rather than the audience. One group initiates the movement and music, but this is the only instance in the piece where hierarchy is present. Beyond that moment, both performers possess equal agency to change, act, react, or propose, fostering a dynamic and evolving interaction.

The piece is built around the concept of illusion, making prior material preparation essential for the performance. The score includes pre-composed elements that must be rehearsed and recorded in advance. These recordings serve as the training data for a concatenative synthesis algorithm and a Variational Autoencoder (VAE), enabling the system to generate and manipulate musical textures that blur the boundaries between the performers' live actions and their processed signals. The sound material -in terms of timbre and melodic/rhythmic contour is shared by both performers.

These two models are performed using a movement sensor called MUGIC, operated by the other performer. This sensor enables full-range motion tracking, and since it is wireless, it can be easily concealed [4], further enhancing the illusion at the heart of the piece.

All the signals generated by the performers are played through localized speakers positioned next to them, ensuring a direct and tangible sonic presence. Simultaneously, these signals are duplicated and spatialized throughout the room using an octophonic setup, creating a layered, immersive environment. This spatial interplay between the real and the synthetic not only reinforces the illusion but also expands the perceived dimensionality of the performance, blurring distinctions between the original sound sources and their artificial counterparts. The result is a constantly shifting auditory experience that challenges perception and redefines the relationship between gesture, sound, and space.

Finally, all the signals get processed by the CAVIAR and sent back into the room.

### 3.1 Thoughts

This piece was only made possible through the Artificial IR system, which provided a clean and transparent method for seamlessly blending the two performers into a singular character with unified timbres. The elegance of the setup lies in its minimalism: there are almost no visible cables, lights, or machinery. Instead, the room itself transforms into a highly reactive instrument, shaping and reshaping the sonic environment in real time.

As a result, the audience gradually loses track of who is performing and where the sounds originate. The interplay between live and processed signals creates a sense of spatial ambiguity, making it difficult to discern the true acoustic characteristics of the venue. This deliberate confusion reinforces the piece's illusionary nature, immersing the listeners in a constantly shifting sonic landscape where perception is continuously challenged and redefined.

Finally, the constant reshape of the reverberation triggers a constant felling of instability for both, performers and audience, shaping the listening and sculpting a constant question: What is going on?

## 4. *A BLANK PAGE*, OR HOW TO IMMERSE BOTH AUDIENCE AND PERFORMERS INTO A ( 13-MINUTES-LONG-17-SECONDS REVERBERATION) WALL OF SOUND

*She looked in the mirror; the other one wasn't there. She found a blank page to start from scratch and build the universe again.*

A Blank Page is an exploration of creation, absence, and reinvention. Written for N performers and live electronics, the piece unfolds in real time, shaped by both human performance and algorithmic intervention. At its core is Pandora's Dream, a system that dynamically generates and distributes notation while simultaneously processing the performers' sound, weaving it into an evolving electronic landscape [5] using all sort of techniques from computer vision, old school machine learning and newer AI approaches.

The score is never fixed. Instead, it emerges moment by moment, dynamically altered by machine listening, artificial intelligence, and real-time transformations of the acoustic space. Each musician's sound feeds into the system, influencing the next set of instructions, erasing the distinction between composition and improvisation. The performers are both creators and interpreters, navigating a constantly shifting sonic terrain.

Spatialization plays a key role. Electronic processing—including granular synthesis and fluid, dynamically shifting reverberation (ranging from dry intimacy to vast, cavernous 17-second decay) reshapes the perception of space itself. The orchestra, whether treated as a singular collective entity or as a network of independent voices, exists within this ever-changing sonic architecture.

Inspired by the idea of staring into a mirror and seeing nothing, A Blank Page invites both performers and listeners to experience creation from absence to embrace the unknown, start from scratch, and build an ephemeral universe anew with every performance.

### 4.1 Thoughts

This piece pushes the exploration of space as an instrument even further, focusing on a gradual but constant crescendo that originates from the natural resonances of the room, much like in Alvin Lucier's seminal work "I am sitting in a room" [6]. From this foundation, the piece evolves using machine listening to generate scores, drawing significant inspiration from Collins' Roomtone Variations [7].

At this stage, the piece becomes obsessed with generating audio and digital feedback loops, pushing the system to its very limits. These recursive layers of sound continuously feed into each other, forming an intricate and ever-intensifying sonic structure supporting itself in granular synthesizers, and the ever-growing reverberation.

As a result, a massive wall of sound emerges, expanding dynamically yet almost imperceptibly, growing in density and presence until it reaches a critical threshold. Then, just as subtly, it dissolves, fading away to leave the page blank once again, returning to the void from which it began, closing the loop.

## 5. TECHNICAL ASPECTS OF THE TWO PIECES

Each of the pieces employed different sets of IRs, each sets of IRs being composed of IRs with different lenght, and made of 28 channels. We summarize below the technical characterisctics associated with the IRs, for each piece, as seen in Table 1

**Table 1**. *Comparison between the two pieces.*

| Measure | TLG | BP |
|---|---|---|
| Number of IRs | 3 | 8 |
| Lengths | 1–5s | 0.4–17s |
| Structure | Random | Crescendo + cuts |
| Spatialization | Live | Fixed, then random |

In *Through the Looking Glass*, three distinct IRs were used, ranging from 1 to 5 seconds in duration. The shortest reverb was soft and homogeneous, while the longest revealed significant spatial variation. Changes in IRs occurred both gradually and suddenly, following the performers' movements, with live spatialization of prerecorded sounds.

By contrast, *Blank Page* employed eight IRs with durations up to 17 seconds, pushing acoustic parameters to their extremes. Shorter IRs had regular frequency responses, but longer ones were highly irregular and varied across channels—up to ±2 seconds for the longest. The piece unfolds in two parts: a crescendo from short to long reverbs, then a final section dominated by the 17s IR, interspersed with abrupt spatial and temporal shifts.

## 6. CONCLUSION

This paper introduced a method for generating and deploying artificial impulse responses within a dynamic virtual acoustics system, enabling a new mode of musical performance where space itself is an expressive medium. The "Irreverberable" algorithm and its integration into the CAVIAR system demonstrates a way to design, control, and evolve reverberant environments in real time.

By examining two contrasting case studies we highlighted how artificial IRs empower both performers and composers to explore new sonic territories. These experiments suggest a broader potential for artificial IRs beyond technical convenience: they offer a fluid and malleable sonic architecture for reimagining performance, composition, and listening. As the distinction between real and virtual acoustics continues to blur, artificial IRs invite us to question not only what we hear, but where—and how—we perceive sound to exist.

## 7. ACKNOWLEDGMENTS

Iran Sanadzadeh, CCRMA Stage and community, Jonathan Abel and Hassan Estakhrian.

## 8. REFERENCES


[1] V. Välimäki, J. Parker, L. Savioja, J. O. Smith, and J. Abel, "More than 50 years of artificial reverberation," in *Audio engineering society conference: 60th international conference: dreams (dereverberation and reverberation of audio, music, and speech)*, Audio Engineering Society, 2016.

[2] J. S. Abel, E. F. Callery, and E. K. Canfield-Dafilou, "A feedback canceling reverberator," in *Proceedings of the Digital Audio Effects Conference*, 2018.

[3] H. Estakhrian and J. Abel, "Production approaches, loudspeaker configurations, and signal processing architectures for live virtual acoustic performance," in *Audio Engineering Society Convention 152*, Audio Engineering Society, 2022.

[4] M. Kimura, "Rossby waving for violin, mugic® sensor, and video," in *NIME 2021*, PubPub, 2021.

[5] J. Armitage, N. Privato, V. Shepardson, and C. B. Gutierrez, "Explainable ai in music performance: Case studies from live coding and sound spatialisation," in *XAI in Action: Past, Present, and Future Applications*, 2023.

[6] S. Hasse, "I am sitting in a room," *Body, Space & Technology*, vol. 11, no. 1, 2012.

[7] N. Collins, "Roomtone variations," 2013.