Wandering the Piano: A Soundscape Installation
Lydia Krifka-Dobes

My concept: Wandering through acoustic space

As an avid hiker, I sometimes imagined walking through a landscape of sounds. The ups and downs of space we experience on a walk in nature, the smooth or rough surfaces we trod on, the gradual change from open space to thicket, from sandy soil to boulders, the occasional getting-stuck in a ditch – all this should be transformed into mostly smooth, sometimes abrupt trajectories in acoustic space, where the spatial clines should be translated into sound clines like pitch, amplitude, and timbre.

Intertwining sound and space is nothing new of course. There are the Rennaisance missas and motets of Striggio[1] and Tallis[2], to be performed by eight or more spatially arranged choires in grand church buildings. There are the "Wandelkonzerte", promenade concerts, sometimes performed in parks, with different orchestras situated distributed in a park, like Stockhausen's "Sternklang"[3] composed for five groups arranges in a wide circle. The strolling visitor can collect musical experiences by moving from one sound source to the other, and enjoy a mix of different sounds in between. I was inspired by this concept. However, I wanted the wanderers to trigger the sounds themselves when they are getting close to the sound sources, to add a moment of surprise and discovery.

A better concept for what I had in mind seemed to be provided with the audio guide technology, where visitors to a museum or another space get speech or music input via headphones when moving close to certain objects. I became familiar with this when I visited an immersive sound installation in the Botanical Garden in Berlin at the CTM Festival 2020, at the occasion of Alexander von Humboldt's 250th birthday.[4] The idea was to get the ambulating audience closer to the world of plants – palms, vines, cacti, ferns. This was technically carried out by the company Usomo[5] in Berlin, which specializes in such installations, especially for museums. However, I wanted that the wanderers realize the sounds that the other wanderers on their strolls through space, which meant no headphones. And I did not want the sounds to help explaining things or adding to the experience of things. I wanted the sounds to stand for themselves, without any relation to objects that could be seen or touched.

---

[1] https://en.wikipedia.org/wiki/Missa_sopra_Ecco_s%C3%AC_beato_giorno
[2] https://en.wikipedia.org/wiki/Spem_in_alium
[3] https://en.wikipedia.org/wiki/Sternklang
[4] https://usomo.de/en/projekt/2020-ctm-festival-at-the-botanic-garden-berlin-en/
[5] https://usomo.de/en/

An opportunity ot realize my concept

I was given the chance to realize my ideas with the Biennale Musica 2023, the 67th Festival of Contemporary Music.[6] Under the direction of composer Lucia Ronchetti, the Biennale announced the topic of “Micro Music”[7]. It was to be “dedicated to digital sound and to its production”. And it stated: “The particular acoustic imprint of each location and the act of listening, understood as reading the surrounding space through the complex, ever-changing reverberations that the ambient creates, emerges in the electronic creations of the various different sections of the Festival, which are designed to seek out the magical sounds of the architectural context of Venice’s buildings.”

Even though my concept was to create a spatial sound world in its own right, and not in relation to particular objects, buildings or places, I applied with a concept where visitors ramble through a featureless space, thereby triggering sounds played from loudspeakers arranged around this space. These sounds should change gradually in one acoustic dimension for a person walking in one direction, and change gradually in another dimension hen walking in another direction. If visitors decide to walk diagonally, the changes in these dimensions should combine, of course. I called this concept “homotopic mapping from places to sounds”. That is, when a place y is between x and z, $x < y < z$, then the sound S(y) triggered at y should have features between the sounds triggered by x and z, $S(x) < S(y) < S(z)$. The idea was that visitors should follow the clines of sound just as migratory birds follow the lines of a magnetic field. They are of course allowed to change their direction and walk back and forth to enjoy and investigate the resulting sound patterns. And they should listen to other visitors and coordinate their movements with them, thus generating even more interesting emerging patterns.

I was very happy that the jury selected my concept and invited me as one of ten sound artists, especially as this happened in a time where I had to go through a tragedy in my familiy. The selection included joining the Music Biennale in Venice four times for periods of two weeks each, to develop my concept in discussions with the other invitees and to work out possible realizations with the musicians and technicians of the Biennale. The advisers with whom we could collaborate included Thierry Coduys, Brigitta Muntendorff, Ali Nikrang, the sound artist Kyoka, Miller Puckette and Gerfried Stocker. It turned out that Thierry Coduys, Ali Nikrang and in particular Miller Puckette were of invaluable help for me to make my initial ideas work.

Selecting Soundscapes

It just so happens that I already had a particular idea about where I should draw the sounds that the visitor should wander through, and that this idea was convincing to the Biennale jury. But I considered a number of other possibilities.

[6] https://www.labiennale.org/en/music/2023
[7] https://www.labiennale.org/en/music/2023/biennale-musica-2023-micro-music

I was truly fascinated by the collection and classification of birdsongs by Kyle McDonald and others at Google Creative Lab, using data from the Cornell Lab of Ornithology[8]. They used thousands of bird sounds that where available as short acoustic files and as sonograms and classified them with a clustering technique, mapping them onto a two-dimensional space. They then created a web-interface that allows users to click on a sonogram, hear the sound and see a photograph and the name of the associated bird. It is fascinating to wander with one's cursor through the world of bird sound, and it would perhaps be even more fascinating to actually walk through an imaginary forest representing the avian universe. I communicated with the creators and got valuable support for my project but put this aside for another occasion.

Another idea would be to have visitors traverse the sounds recorded at particular places in a city, like Venice. This would have been close to the suggestive text by which the Biennale wanted to attract sound artists. However, this would have required a relatively large space and many trigger points to work. Yet another idea was to have the initial parts of national hymns of the the more than 200 nations on earth, triggered by points that correspond to their geographic location on the globe. However, this would have been more of an educational tool, of limited aesthetic use. Yet another idea was to map the vowels or even other sounds that we find in human speech, following the vowel square of the International Phonetic Alphabet that represents the position of the tongue and the rounding of the lips when pronouncing 33 distinct vowels[9] – with nasals and other modulations, like whispering and vowel fry, this would lead to more than 100 distinct sounds. Again, after discussion, I put this idea aside for possible future realizations.

Recording Piano Sounds

My original idea, on which I settled on the end, were piano sounds that I had recorded in March 2023 together with composer Samir TimajChi at the Carl-Maria-von-Weber Hochschule für Musik in Dresden with a Grand Bösendorfer Piano with an extended key range of 97 keys. This was done with a spatialized presentation in mind on my side, but also for possible other purposes on Samir's side. We could work with a Bösendorfer piano for three days and have now both very hiqh-quality audio recordings as well as video recordings with the player dressed in a green suit for potential future video manipulations.

We systematically recorded normal productions of piano sounds by hitting the keys using the pedals in various ways, but also by manipulating the strings directly with many different methods. This included:

- Production of narrow accords with neighbouring keys
- Rapid oscillation between keys for a tremolo effect
- Exploring sustained sonorities with pedal usage
- Keys tapping silent and key backward flip
- Creating sounds on strings with percussion mallet and steel, glass and rubber balls

[8] https://experiments.withgoogle.com/bird-sounds
[9] https://en.wikipedia.org/wiki/IPA_vowel_chart_with_audio

- Creating sounds on strings by moving a vibrator along the string
- Placing a screwdriver between strings to manipulate tension and contact points
- Opening and closing the piano, with and without pedal

These recordings formed the raw material of my installation, which consequently should was named “Wandering the Piano”. The advantage of this sound material was the gradual changes that they afforded, the high quality of the recording, and the ease of combining sounds that would result from a person traversing space quickly, or from several persons walking through the space simultaneously. Even though I was worried about what would happen when too many persons – more than five, say – would generate sounds simultaneously. But a lot had to be done before I could make use of these recordings in the way I wanted.

The performing space

During my first preparation phase in Venice in spring 2023, I was told that I would have my installation in a hall at the Ca’ Giustinian in the very center of Venice close to San Marco, and very close to the headquarters of the Biennale (that is, not at the somewhat remote area of the Arsenale). I was excited to learn that Richard Wagner stayed at this palace for seven months (1858/59), composing parts of “Tristan and Isolde”, and that Marcel Proust and Giuseppe Verdi were guests there when it turned into a hotel. And I also learned that the Duomo of San Marco itself was used for spatially arranged music at least since late medieval times.

The room itself could be accessed directly from the Calle de Ridotto, a narrow pedestrian walkway. The usable space measured about 12mx10m, it had an access ramp in the middle, and, unfortunately, seven narrow columns distributed in the room (Figures 1 and 2).

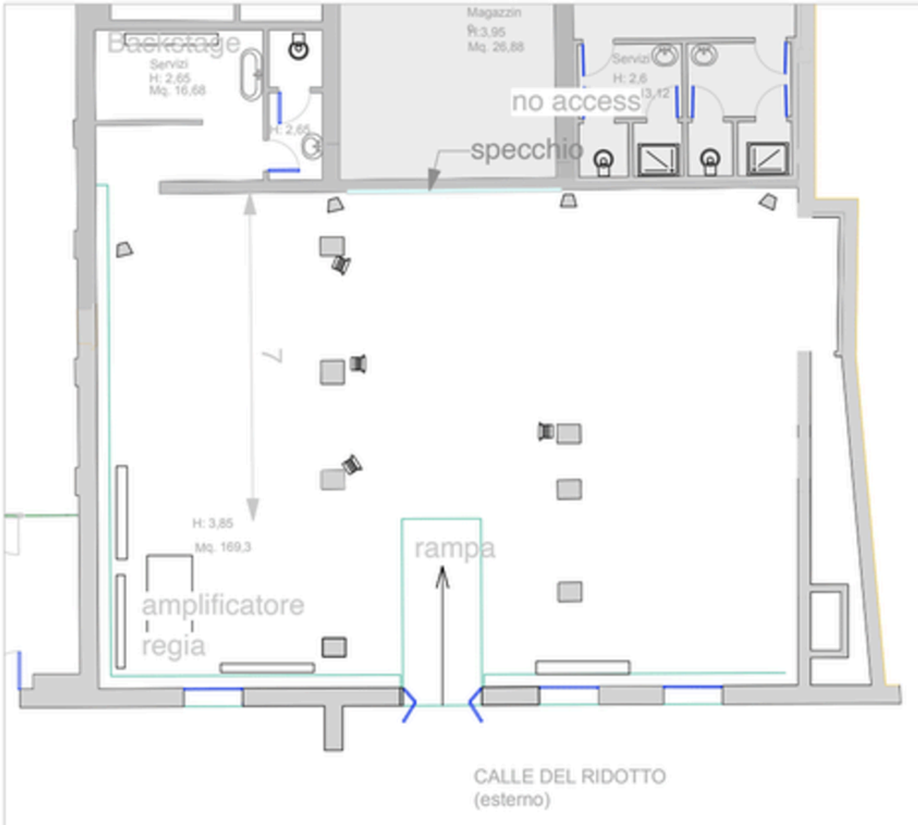

Figure 1: Ground Plan of Laboratorio delle Arte, Ca’ Giustiani

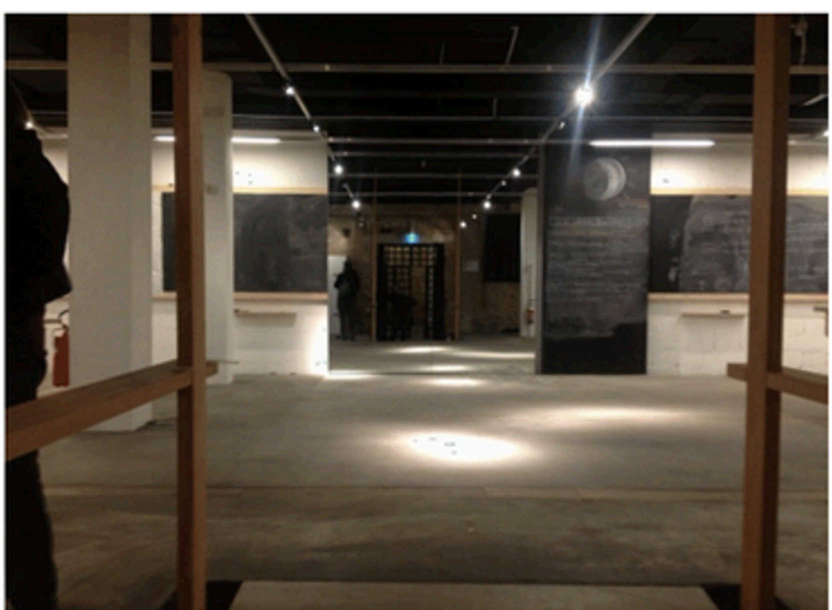

Figure 2: Laboratorio delle Arte, Ca’ Giustani, before installation

Technical Choices

I was assigned a team of technicians, who also had to work with the other applicants. First, we had to determine the best ways of realizing the trigger signals when a visitor would walk by a determined point. There are various options, including pressure sensors on the floor or RFID chips with a range that is used for, for example, securing clothes at shopping centers or books at libraries. This would have required that the visitors carry with them a chip, which would have allowed to restrict the number of visitors but would be inconvenient for the visitors and difficult for the organisation. I learned that another option would be quite inexpensive radar modules (type ld2410). Yet another method would have been tracking of moving bodies by cameras, but this appeared problematic for various reasons. In particular, I planned to have the space relatively dimly lit, which makes regular cameras a poor option; furthermore, there would be potential problems with this technique with several people walking in the space at the same time.

We ultimately settled on relatively inexpensive infrared sensors which are used, for example, to detect human bodies at elevator doors. Negotations with technicians, and considering the available space, brought the number of such trigger sensors down to 64 (initialla I had hoped for 128 trigger points). They could be installed with the necessary cables on a grid frame of wooden bars at the ceiling of the room, at a hight of about 3m. First tests showed that the passing of persons on the floor would be recognized by the sensors reliably and swiftly.

As for the lightning of the room, I wanted to have very low light with as few features as possible (where of course the columns could not be removed!) In discussions with the technicians I was convinced that LED lamps on the ceilings that identified the trigger points on the floor by circular lights would be helpful, and add to the aesthetic experience. I was very satisfied about the resulting overall impression (Figure 3, with visitors).

Figure 3: Laboratorio delle Arte with trigger points and visitors

Creating the Bösendorfer Corpus of Piano Sounds

It turned out that a considerable amount of work still had to be done to prepare the piano recordings so that they could be used to create the soundscapes that I intended. The acoustic recordings had to be cut into short files that contained a single acoustic event. This could be done automatically with transient detection and silence cutting, ensuring that each file was trimmed to contain only relevant sound material. Furthermore, the sound files had to be named in a way that revealed the sounds that they contained, in a way that helped to determine how similar to each other they were.

These tasks were achieved with a python program for which I made use of tools provided by the librosa library.[10] The program cut the Bösendorfer recordings into small wav files that contained single sound events, and categorized these sound events according to five features:

(a)	Frequency: spectral_centroid	values 0…9	n_1
(b)	Amplitude: np.max(y)	values 0…9	n_2
(c)	Melspectra: melspectrogram	values 0…9	n_3
(d)	Chroma: chroma_stft	values 0…9	n_4
(e)	AudioLength: audio_length,	values 0…9	n_5

The sound snippets were assigned an integer value from zero to nine according to these five features, resulting in file names of the form “nnnnn.wav”, where n stands for a digit (0 to 9). In this way, the closeness or distance of sound snippets to each other could be computed; the sound of the file 12345.wav was categorized as closer to the sound of

[10] https://github.com/librosa/librosa

22345.wav than to the sound of 32345.wav, for example. The distance between two snippets could be computed easily by the sum of the numerical distances between the digits (for example, resulting in a distance of 3 between 12345.wav and 32445.wav). This distance between snippets corresponds quite well to the distance that human listeners would assign to the sounds, as the employed features (a) – (e) correspond to features that are relevant for human sound recognition as well.[11]

As a result I ended up with more than 10,000 short wav files that were numbered in a way that contained essential information about the sounds they contained, and their similarity. If the program generated files with the same name, which occurred rarely, existing files were overwritten and only one file survived.

Creating soundscapes

I wanted that the creation of individual soundscapes should be done by automatic clustering of the snippets, or rather the names of the files that were created in the previous steps. After talking to Ali Nikrang, we employed the technique of k-means clustering[12] of the data into distinct cells on the basis of the similarity of their file names. Such cells formed the basis for constructing soundscapes, which contained 64 sounds each.

I found that k-means clustering can be particularly useful for managing large datasets, allowing for easier navigation, retrieval, and understanding of the data. The technique can reveal patterns or relationships within the audio data that might not be immediately apparent. For example, clustering might group together all files with high percussiveness or certain melodic properties, revealing commonalities that were not explicitly tagged or known in advance.[13]

For the task of mapping the snippets to trigger points, I created a map of the trigger points and numbered them in the way shown in Figure 4. This does not show the adjustments that we had to make for the columns in the room but identifies the location of the access map. To test out possible arrangements before the installation was available I created a program with a simple GUI interface where I could walk through virtual trigger points with the cursor. Then I found a central "pillar" which I usually located in the center of the space or at the sides. I mapped the soundfiles of a soundscape to one of the 64 trigger points, mostly respecting the homotopic mapping explained above. I did not solely rely on the numeric characteristics of the sound given by the sound file but also checked many individual transitions whether they represented the gradual changes that I intended to produce.

[11] I thank Ali Nikrang (Ars Electronica Linz,) for very important assistance in creating the programs for trimming, analyzing, and naming the sound files.
[12] https://en.wikipedia.org/wiki/K-means_clustering
[13] In music technology, sound design, or other audio-related fields, clustering can facilitate novel insights or applications such as automatic genre classification, mood-based playlist generation, or even guiding sound synthesis and manipulation by identifying and leveraging the natural groupings of sounds

01	02	03	04	05	06	07	08	09	10
11	12	13	14	15	16	17	18	19	20
21	22	23	24	25	26	27	28	29	30
31	32	33	34	35	36	37	38	39	40
41	42	43					44	45	46
47	48	49		Access			50	51	52
53	54	55		ramp			56	57	58
59	60	61					62	63	64

Figure 4: Assignment of trigger points

Altogether I created 16 different soundscapes. In one of the soundscapes, I varied from this and I created an imaginary division of the room at 6, 16, 26, 36 where I split the room in two imaginary sections in order to create two mirror sides of the space, where one side triggered the sound normally, while the other side had the exact same sounds played backwards (this was the only hint of manipulation to the original recordings to be heard). This setup allowed two people entering the space simultaneously to perceive a kind of correspondence and explore it further, although this turned out to be difficult to recognize.In other soundscapes I focused on non-predictability, employing very different organizations, such as dividing the space into five zones with clear borders, and considering the relationships between these zones to each other. I identified the spatial central point, or the pillar, and built everything around that feature. This gave structure to the space, with complex sounds clustered in the corners and the quality of sounds deliberately chosen to relate to different zones based on different features, such as brightness, color, highness, or lowness. In the middle of this process, like in soundscape number 12, I discarded the dome structure to create chaos,. My goal was not to find more possibilities to enlarge the field but to focus and see more possibilities within, driven by both artistic and analytical intentions. I find it important to follow both goals, avoiding the creation of music that is too “analytical” or “musical”. Being clear about my own intentions was crucial. Exploring both structure and chaos was particularly intriguing to me when putting together, or “composing”, the soundscapes.

Programming the installation

The final step turned out to be quite complex – to develop a program for uploading the files of a soundscape into an array of 64 fields, to relate these fields to the signals captured from the infrared sensors at their trigger points, and to play these sound fields on the sixteen spatially arranged loudspeakers, or on subsets of these loud speakers. Several principled decisions had to be made. For example: How should the transitions between individual files be handled? I decided that the sounds of individual files should overlap; in fact, each file should be fully played to its end. How should the sounds be played back on the 16-channel loudspeaker system, which could be accessed individually? We compared methods of using the three closest speakers to a trigger point to send the maximum sound energy to that point, we tried wave field synthesis, and other techniques for localisation of sound. For example, we integrated a check if a sound file exceeded a specified duration, in case the sound was spatializied, creating a path from the trigger point to the opposite side of a room. Each had certain qualities I liked.

One further issue to decide was: How long should the individual soundscapes last so that visitors could explore them? I decided for 3 minutes, as this would give visitors a chance to realize that the soundscapes were changing, and to experience several distinct soundscapes and develop an interest to listen to all of them. But how should the change from one soundscape to the next be indicated? I decided on 15 seconds silence and a dimming of the LED lights, but never reaching darkness (which was also required for security reasons). The change of soundscapes was also a good signal for visitors to leave the room.

The programming turned out to be considerable more complicated than I initially envisioned, and would not have been possible without the help of experienced sound engineers, in particular Thierry Coduys and his team. The tool that was used because Thierry was most comfortable with was Panoramix,[14] developed by IRCAM as a standalone application dedicated to spatial audio mixing and post-production. In the development of the final product, which was going on right before the inauguration, we tried out quite a few other interesting manipulations that showed the potential of this setup, which ultimately did not make it to the final presentation.

.

[14] https://forum.ircam.fr/projects/detail/panoramix/

Wandering the Piano, in Laboratorio delle Arti at Ca' Giustinian reimagines the grand piano, as the most iconic instrument in Western music. It can be used to produce a wealth of sounds, from the intended classical methods by hitting the keys and applying the pedal to manipulations of the strings themselves, as in avant-garde music.
Visitors can trigger the contrast and borders of topologies of sound by just walking around, thus exploring the acoustic universe of the piano. The layout invites visitors to listen to each other and react to each other, creating their own improvised passages. Visitors are also invited to record their experiences with their mobile phones.

Figure 5: Announcement by the Biennale di Venezia

My own impression and the experience of visitors

I was present at the installation (or near to it) every day during its run, sometimes for several hours. This was also necessary because the system sometimes went down and had to be reset, which did not take long. The installation was well attended, sometimes with more visitors than we wished. Access was free, but an employee of the Biennale would have visitors wait if too many of them were in the room.

It was exciting to see that quite a few visitors liked the installation, some spending quite some time in it. There was one family who came back three times because their kid liked it to much. Others moved and just looked at it from the entrance, or just spend a few seconds and came out because they could not make sense of it, but this was clearly the exception.

Many people thought that the sounds were synthesized and were surprised to hear that they were piano sound recordings, in spite of the title and information on the billboard at the entrance (Figure 5). One way to avoid this would have been to have one or a few sound scapes that are closer to regular piano sounds. There was also the suggestion, by Gerfried stocker, to create a holographic representation of a piano in the middle as a central point of attraction. I found this attractive but we discarded it, among other reasons because it would distract from the acoustic universe that I wanted to create.

I had envisioned "swarms" of sound with more densely populated sounds to explore which didn't materialize as expected due to the sounds' equidistant placement. This led to a reflection on the balance between guiding the audience's experience and allowing for open-ended exploration, underscoring the nuanced dance between artistic intent and participant engagement.

Possible future implementations

The Installation “Wandering the Piano” lasted only for about one week. Where can I get to from there? For one thing, the hardware and the corresponding software could be installed in other places. This could be an exhibition space with another temporary exhibit, but also a museum or art or music school with a permanent exhibit. One could run the same installation there but also other installations with an artistic or educational purpose. I have mentioned some options above.

One can also vary the installation itself in certain ways. One example is to enlarge the number of triggering points, and correspondingly, the space of the area. This might be more easily possible with another technique for trigger points, especially for large-scale outdoors installations. Of course, this would be an expensive option. Another possibility is to use a computer screen move virtual “persons” around. The disadvantage of this is that it lacks the sensory feeling of moving around in space with one’s body. Something in between real bodily movement in space and moving an avatar on the screen would be a board game, with a board in which pressure-sensitive sensors are installed. The movement of figures (e.g. chess pawns) would trigger certain acoustic events when a figure reaches a new field. Yet another implementation would be virtual reality, where users experience a world through VR spectacles and headphones. Finally, one could imagine an implementation with augmented reality, where people trigger noises on their headphones when walking by specified points. The trigger then would be closeness to particular fine-grained geo-locations picked up by their smartphones.

There is a range of purposes that soundscapes could be employed for. Beyond the creation of aesthetical value, it can be used for educational purposes, like exploring spatially distributed phenomena that are of acoustic nature, like music styles, or which could be sonified, like mathematical functions. After experiencing the installation in Venice I can imagine that walking through soundspaces can have therapeutic purposes or bring people together socially. One particular intriguing idea is that visitors could listen for trigger points that lift them from one soundscape into the next, similar to levels in computer games. These rather general ideas do not exhaust the possibilities that triggering sounds at particular places can offer.

A Markov Walk through the Grand Bösendorfer Corpus

To conclude, I would like to present one idea that is not really a soundscape in the above sense but a ways of showcasing a larger set of sounds, or other events, namely by Markov chains. They are ideal for simulating walks through a system of states. The nature of the walks can be fine-tuned according to certain criteria. Markov Chains with no memory (oder 1) depend solely on the current state to determine the next state, disregarding the path taken to arrive at the current state. For random walks with equal probabilities of moving to any neighboring point, the transition matrix would be uniformly distributed among the neighbors. For instance, if a point has eight neighbors, each neighbor would have a one-eighth probability of being the next step. Such a state would be completely random, a bit resembling the Brownian motion of a particle pushed around by water molecules. To model walkers that tend to move in straight lines, we adjust the transition probabilities to favor the direction of the previous step. This requires a more complex Markov chain with memory that considers both the current position and the direction of the last step or steps, leading to a higher-dimensional state space.

I developed a program that walks through the sound snippets of the Grand Bösendorfer Corpus. The files with their names with the format nnnnn.wav form the system of states. The chain is set in motion at a particular file, which is played. The algorithm chooses as the next state a file that is close to that file, and so on. The selection of the next files follows aleatoric principles, using random numbers, but the algorithm allows to weigh the transition points differently.

In the realm of creating an immersive and interactive soundscape, especially one derived from a rich collection like the Grand Bösendorfer Corpus, the use of Markov chains could be involved to incorporate into a composition or a game with a user. This composition uses the prev. Played sound like its index, sonic features, and any other relevant metadata) to calculate the probabilities for the next selection. This method significantly enriches the listener's journey through the soundscape, making the experience not just a series of random steps but a guided exploration where certain paths are more favored.

The weighting system plays a crucial role here: By assigning higher weights to specific portal points within the audio landscape, the algorithm is tuned to induce jumps to new areas when certain conditions are met. These conditions in the future for a game could be based on a variety of factors, including the duration of time spent in a particular sound area, listener interaction, or reaching a predefined threshold of audio feature analysis. Such weighted transitions ensure that the exploration remains fresh and engaging, constantly unfolding new layers of the auditory experience but also letting the user to explore enough. Also regarding spatialisation: The "8ch_clusters.maxpat" patch allows for the intricate placement of sounds in the listening space, where the weighted Markov chain algorithm decides not just what sounds are played but also how they are spatially arranged around the listener. This approach affords a high degree of control over the immersive qualities of the soundscape, amplifying the impact of the audio journey.