Probabilistic Generative Music as a Data Sonification Medium

Kantan, Prithvi Ravi

Published in:
Proceedings of the 3rd Conference on Sonification of Health and Environmental Data (SoniHED 2025)

DOI (link to publication from Publisher):
10.5281/zenodo.15032599

Publication date:
2025

Document Version
Publisher's PDF, also known as Version of record

Citation for published version (APA):
Kantan, P. R. (2025). Probabilistic Generative Music as a Data Sonification Medium. In S. Pauletto (Ed.), *Proceedings of the 3rd Conference on Sonification of Health and Environmental Data (SoniHED 2025)* (pp. 77-84). KTH Royal Institute of Technology. https://doi.org/10.5281/zenodo.15032599

Probabilistic Generative Music as a Data Sonification Medium

Prithvi Ravi Kantan
Department of Architecture, Design and Media Technology
Aalborg University, Copenhagen
`prka@create.aau.dk`

ABSTRACT

Auditory displays that incorporate musical structural elements and textures have been used to increase long-term listenability and engagement. Such approaches are often met with skepticism due to the inherent tradeoff between aesthetics and information fidelity. Additionally, their typical reliance on simple deterministic mappings (e.g. MIDI pitch or note duration changes) produce sonic outputs that are either overly predictable or musically incoherent. I propose and demonstrate a framework that expands the parameter mapping space through the manipulation of determinism and nondeterminism in real-time sequenced generative music. The result is dynamic and improvisatory musical patterns that may sustain listener interest while saliently conveying data trends with adequate perceptual independence between musical dimensions. The resulting parameter space may be suited to applications targeting artistic expression or augmented relaxation, sleep, and movement, where informational fidelity and temporal precision may marginally be sacrificed in favor of musicality and its associated affective / motivational benefits. Although pending formal evaluation, this work reflects flashes of the potential that probabilistic generative music has as a data communication medium.

1. INTRODUCTION

1.1 Aesthetics and Fidelity in Sonification Design

Sonification is the systematic and reproducible scientific method for representing data (typically as nonspeech sound) in an auditory display [1]. Sonification has many application domains ranging from scientific to artistic, and the exact data-sound relationship is naturally pivotal to the utility and usability of a sonification. Depending on the application, criteria such as aesthetics, data representational fidelity, and the intuitiveness of data-sound metaphors used may be weighted differently, and research over the years has revealed that these criteria share a complicated push-and-pull relationship with many studies revealing tradeoffs between musicality / aesthetics and data fidelity [2–7].

Focusing primarily on scientific applications, a general consensus among veterans in the field has been that while conveying the intended 'message' is paramount, it is prudent to design sonifications that are aesthetically pleasing to the extent possible [3, 4, 8, 9]. This rings especially relevant when considering applications that entail extended periods of listening to a sonification, where 'poor acoustic ecology' may make listening and subsequent data interpretation difficult [10]. These use case contexts include movement rehabilitation [11], peripheral process monitoring [12], and relaxation / mindfulness [13], where it is necessary to appropriately balance communicational efficiency and aesthetic character so the rendered sound does not seem arbitrary in relation to the data underlying it [3]. To promote meaningful mappings, the concept of *conceptual metaphors* has more recently been formalized as a cornerstone of sonification design [14]. Metaphors (e.g. love is a journey), which economically transfer large chunk of information from familiar domains of thought to unfamiliar ones, can connotatively convey information about phenomena that even words cannot, and metaphors are integral to several critical domains of language and thought [14, 15]. Sounds have a natural proclivity to bind to metaphorical associations [16], and a sound-specific domain that most humans have extensive (and usually positive) embodied experience with is music, which has made music an intriguing vehicle for data communication among sonification designers over the years [2].

1.2 Musical Sonification

Notwithstanding ongoing debates regarding the definition of music [2], the term *musical sonification* has historically been used to loosely characterize sonification designs that, for either artistic or aesthetic purposes, integrate traditional elements of musical structure or texture into their design. Musicality may be realized at either the symbolic level (e.g. manipulating MIDI pitches) [17] or signal level (e.g. altering digital processing applied to recorded / synthesized music) [18]. Common reasons for imbuing sonification design with musical traits include enhancing long-term listenability and mediating affect / motivation [2, 11, 18, 19] coupled with the notion that music provides an appealing and universal 'sonic grammar' that can be leveraged in the creation of metaphors to reliably convey information about a dataset [3, 20]. Despite these positives, studies have routinely found that 'data-irrelevant' musical components may run the risk of 'cluttering the message' in the data and reducing user task performance in certain regards [5–7, 18], aside from the reality that even 'musical' sonifications may be unengaging and / or fatigu-

DOI: https://doi.org/10.5281/zenodo.15032599

ing to listen to [2]. Plausible reasons include the prevalence of simplistic data-to-MIDI pitch mappings as well as the information-dense nature of most numerical datasets, which can often lead to musically incoherent sonic outputs [2].

1.3 Proposed Generative Paradigm

As the fundamental medium for data communication, I propose the use of music that is:

1. Generative (algorithmic) and polyphonic with evolutionary components that are both dataset-dependent and independent.
2. Amenable for the manipulation of both its deterministic and nondeterministic characteristics.
3. Probabilistically sequenced in real time.
4. Stylistically inspired by repetitive instrumental popular music genres (e.g. drums and bass).

The vision is, in essence, one where data variable variations convey information through improvisatory musical changes applied in real time to a pre-sequenced basic musical structure. Although sequencer-based sonification approaches have unsystematically been applied before [11], such a paradigm is, for all practical purposes, a novel one with several strong motivations. Generative music, with algorithmically creates a musical output in a rule-based, data-driven, or hybrid manner [21, 22], makes principal sense as a data sonification medium for certain applications. Unlike completely predefined or completely data-controlled music, it allows sonification designers to define which aspects of musical evolution are data-controlled and data-independent, enabling the generation of output sequences that are not only coherent in a musical sense but also informative about the dataset along very specific musical dimensions. Secondly, depending on the generative architecture, it may also expose a variety of musically meaningful control parameters not otherwise accessible when mapping to pre-rendered or predefined MIDI music. Finally, generative music can help avoid copyright issues, flexibly blend musical elements, and tailor the output to environmental and listener-specific aspects [21].

It is common for music performances to comprise a combination of deterministic elements (e.g. composer-specified melodic contours) and nondeterministic elements that are subject to stochastic or creative variation [23]. Although musically informed sonification designs have often mapped data variables to deterministic elements (e.g. melodic pitch, tempo, note duration) [11, 17, 20], nondeterministic elements have received little consideration as ‘mapping material’. Nondeterminism can be computationally modelled using chance processes such as random number generation [24, Sec 1]. Randomness has successfully been harnessed in interactive products to enrich user experience (e.g. iPod Shuffle, *randomURL*), and it has been argued that randomness in music has a distinctive and audible character, which merits its consideration as musical material in and of itself [24, Sec 2]. Nondeterminism has been shown to be an indispensible entailment of creative processes in music [23], and musical novelty / surprise underpins feelings of musical reward and enjoyment at a neurophysiological level [19].

Generative systems typically exhibit a degree of nondeterminism (e.g. seeded random number generation) in their working [22], which not only adds novelty and surprise to the short- and long-term evolution of their output, but also allow for entirely new musical sequences to be generated each time such systems are run. An example parameter mapping to a nondeterministic attribute would be when a continuous data variable is mapped to the *probability* of a musical event being triggered at a particular point in time, as opposed to simply triggering the event when the variable crosses a threshold (the deterministic counterpart thereof). Despite their intrinsic ties to musical creativity and meaning, mappings to nondeterministic attributes have been exceedingly rare in musical sonification design; Hermann’s taxonomy even explicitly dictates that sonifications must be reproducible at the audio signal level [1]. However, I argue that *deterministic mappings to nondeterministic properties* can convey non-audio information similarly to how the colour / timbre of a noise waveform conveys information about the nondeterministic process that generates it; different bursts of noise sound the same to us even if the exact wave sequences vary. Real-time sequencing and time-continuous manipulation of various dimensions of *musical nondeterminism* can similarly help convey non-audio information through a variety of musical metaphors.

The choice of repetitive instrumental popular music as stylistic inspiration has several motivating factors. Popular music is universal and has ubiquitous appeal, which means that most prospective users will likely have lifelong familiarity with effortlessly decoding its characteristic sonic grammar and symbolic syntax irrespective of musical training [2]. Popular music is commonly used as background music during various non-musical activities such as exercise, learning, studying, driving, etc. to improve concentration, mood, and task performance [19,25,26]. Moreover, repetition in music is shown to be an important mediator of anticipatory mechanisms and, in turn, enjoyment during music listening [27]. The simple metrical structure of most popular music also makes it conducive to training (and synchronously providing auditory feedback on) cyclical bodily movements (e.g. walking) [11]. Although musical preferences vary and task performance depends on arousal level [25, 28], popular music seems to readily lend itself to pleasurable long-term listening, a characteristic worthy of leveraging when designing sonifications for process monitoring or real-time biofeedback applications.

1.4 Caveats

Regarding information content, the proposed paradigm is still vulnerable to the possibility of generating a musically incoherent output in certain data mapping configurations, but the in-depth parametric control it offers can allow designers to realize appropriate metaphors for a variety of situations. Regarding fidelity, a nonstationary musical signal with data-independent variation will most likely introduce considerable perceptual ‘noise’ in terms of data value resolution (Y-axis) as well as temporal resolution (X-axis). At the outset, it therefore seems plausible to only make use of this paradigm when conveying data variables that (a) take a relatively small number of discrete values, and (b) do not rapidly change in value. These conditions are broadly satisfied in several instances of process monitoring / biofeedback that convey information about discrete states that evolve relatively slowly.

1.5 Current Scope Delimitation

This work aims to showcase a real-time technical architecture for data sonification through the systematic manipulation of deterministic and nondeterministic elements in probabilistically sequenced generative music. The overall system structure will briefly be explained, and an initial set of mappable parameters are described and demonstrated. Two musical instrumentation styles have been tested thus far - acoustic drums and major scale piano - and a small set of real-time demos with multivariate mappings is also provided. A concise initial appraisal of the potential and implications of the generative paradigm as well as its current implementation will ensue in the discussion.

2. INTERACTIVE SEQUENCER ARCHITECTURE

In order to realize, showcase, and test the concept of conveying information through the manipulation of nondeterministic traits of generative music, a real-time Windows-based generative music sequencing platform was built using the JUCE[1] audio programming environment. It comprises the following key component blocks:

- 6-Channel OSC Input
- Audio Sample Architecture
- Musical Event Chart
- Sequencing and Playback
- Parameter Mapping

2.1 6-Channel OSC Input

The platform takes in six discrete channels of numeric OSC data normalized between 0 – 1. The interface provides a real-time visualization of incoming data on each channel, and it is also possible to simulate real-time input using a slider. Each channel can be linearly mapped to no more than one mapping parameter.

2.2 Audio Sample Architecture

The interface allows for choice between various *instrumentation styles*, which are essentially sample packs that are dynamically loaded from a local folder. For each individual sample file, information about note number and velocity range is obtained from the file name, which follows a strict convention. It is hence possible to have different audio clips for different MIDI velocity ranges of the same note number (e.g. a soft drum hit for velocities 0–64, a regular hit for 65–96, and a hard rim shot from 97–128). There is no restriction on the number of possible note numbers or the number of velocity-distinct sample files per note number.

2.3 Musical Event Chart

Prior to playback, a grid-based chart of musical events over a short musical duration (e.g. 1-2 bars) must be defined using the visual interface shown in Fig. 1. As in any traditional sequencer, the grid rows indicate note number, and the columns indicate grid position. The grid resolution (e.g. quarter / eighth / sixteenth note) depends on the instrumentation style. The contents of each grid location can be manipulated by the user at any time (also during playback), and these manual changes are registered after pattern playback is complete and is preparing to loop. As shown, there are buttons to instantly clear every row (note number) and column (musical time position). Easy-access buttons (see Fig. 2 enable the user to clear the entire chart or fill each grid position with pseudorandom velocity and probability values that are programmatically weighted to accentuate beat intervals. Events may hence be certain (probability = 1) or uncertain (probability $<$ 1). When an instrumentation style is loaded, the default musical chart is an empty grid, i.e. no musical events at any chart position for any note number. A seemingly effective approach is for the user to define a small set of certain events as the *base pattern* (a simple rhythm or melody) and then randomly populate the remaining grid positions with uncertain events. This ensures a balance between identifiable musical structure and short-term variation to maintain listener interest.

2.4 Sequencing and Playback

When playback is started, the musical event chart is queued based on the user-configured grid event velocities and probabilities for each note number. Hence, each musical event is captured by its note number, probability, velocity, and temporal position in MIDI ticks. During playback, the play position is incremented by a specific number of ticks at a 1-ms interval based on the configured tempo, and the sequencer checks for new events at each interval. If one or more events are found, the sequencer decides whether to trigger them based on their event probability by making use of a JUCE *random number generator (RNG)*. When the sequencer reaches the end of the pattern, it loops back to the start. Musical structure variations are brought about through both probabilistic variations over time and the explicit randomization of specific attributes as described in the next section.

2.5 Parameter Mapping

The music sequencing architecture exposes several parameters for real-time manipulation (see next section), and each of the six OSC inputs may either remain unmapped or be mapped to one of them from a drop-down list. Only linear one-to-one mappings are possible at this stage. Each

---

[1] `https://juce.com/`

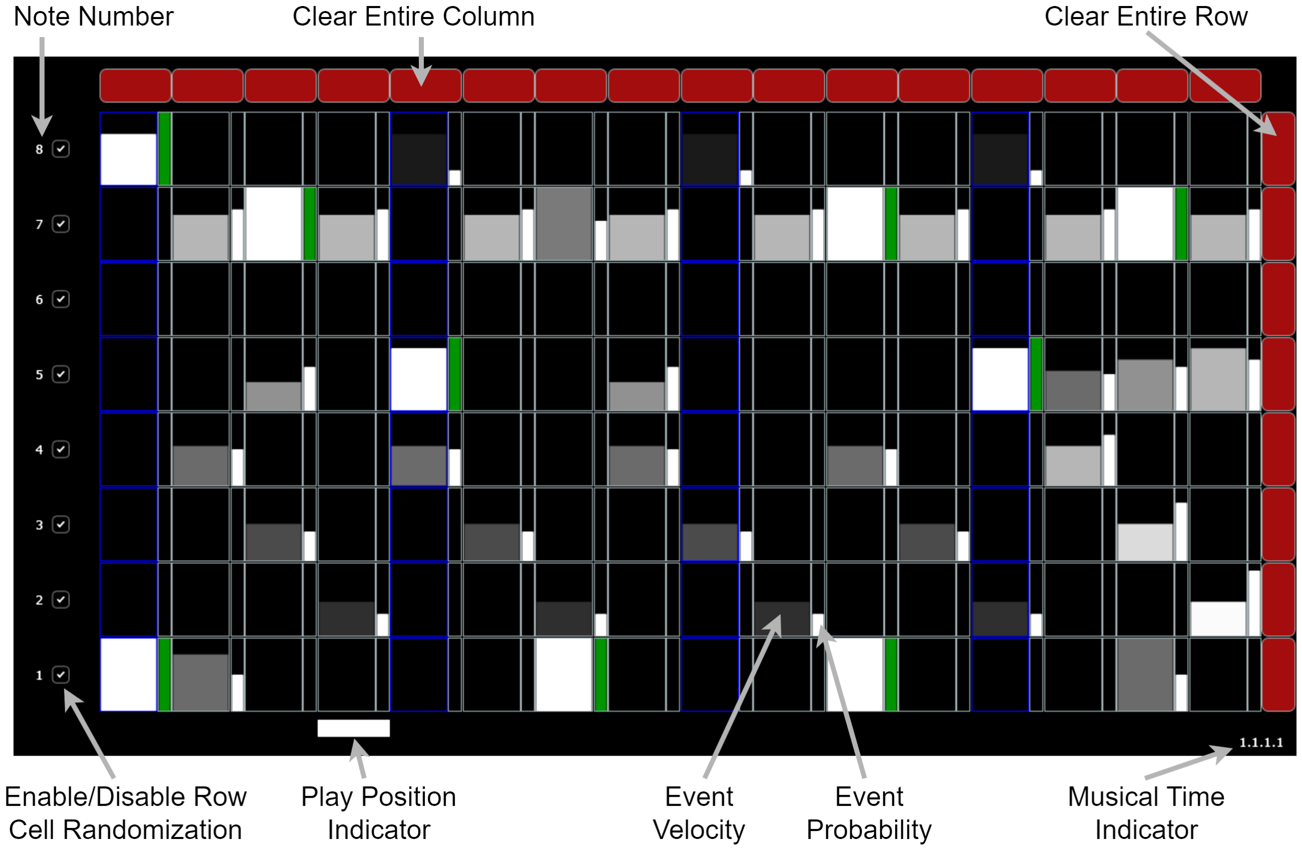

Figure 1. The interactive sequencer interface with key components labelled for a 4-beat grid with sixteenth note resolution (beats marked in blue). For each grid position, the thicker rectangle shows event velocity and the thinner one shows probability (green for certain events). All velocities and probabilities can be manipulated directly from the UI.

input signal (0 – 1) is linearly scaled to the full range of the parameter it is mapped to. The sequencer is updated with new parameter values every millisecond.

3. MAPPABLE SEQUENCER PARAMETERS

Please refer to the supplementary folder[2] for A/V demonstrations of each parameter (as well as some bivariate combinations) described in the following subsections. In each video, input data variable changes can be seen in the mapper section to the **left** of the screen (indicated by downward blue arrows).

3.1 Linear Multiplier Parameters

These parameters work by applying simple multiplicative factors to specific sequencer- or playback-related variables.

1. *Velocity Multiplier:* All triggered event velocities are uniformly multiplied by this factor (default value 1.0, range 0.0 – 1.8). This manifests as changes in volume and event envelope shape that are modulated by changes in the mapped data variable. If different audio clips / samples are used for different velocity ranges (e.g. with acoustic drums), the velocity multiplier alters articulation and timbre as well.
2. *Probability Multiplier:* All uncertain event probabilities are uniformly multiplied by this factor (default value 1.0, range 0.0 – 3.0) prior to RNG check prior to triggering. When it is zero, only the base pattern (certain events) plays. As it is raised, the uncertain events on the chart play with increasing probability, making the output pattern musically busier. As more and more uncertain events become certain, the pattern is at its busiest and most predictable.
3. *Tempo Multiplier:* This applies a multiplier to the tick increment applied to the sequencer every millisecond (default value 1.0, range 0.0 – 2.0). Altering the multiplier value thus alters how quickly the sequencer ‘moves through’ the pattern, changing the output tempo on a moment-to-moment basis without any extraneous alteration to the sound.
4. *Playrate Multiplier:* This applies a uniform multiplier (default value 1.0, range 0.2 – 10) to the read index increment of all loaded audio sample files, which speeds up or slows down all triggered events equally. This manifests as changes in the pitch and envelope shape of individual triggered events, while pattern tempo remains constant.

---

[2] `https://doi.org/10.5281/zenodo.14568976`

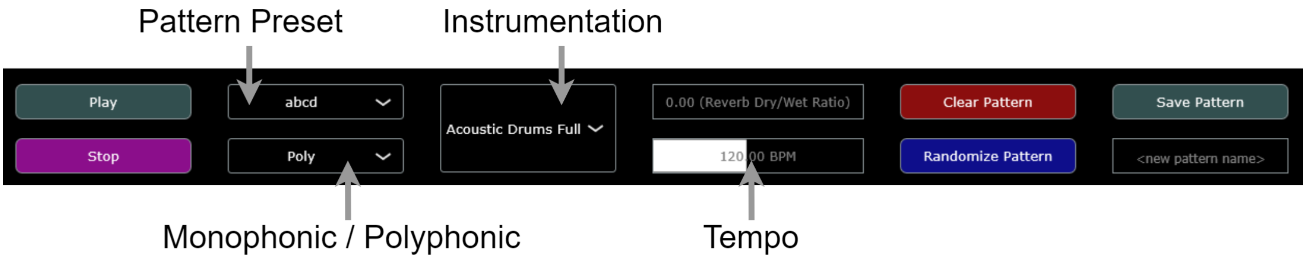

Figure 2. Sequencer playback controls with options relating to tempo, polyphony, instrumentation, and pattern preset choice. Buttons are also provided to instantly clear the entire pattern grid and randomize the probability and velocity of grid positions not currently holding certain events.

3.2 Note Operation Parameters

This class of parameters works at the level of individual musical events by manipulating note numbers or including / excluding events based on note number.

1. *Note Transpose:* This parameter applies a simple additive factor to the note number of all events (default value 0.0, range -10 – 10), which manifests as a simple transposition of the entire pattern assuming that note numbers are arranged in increasing pitch. For non-pitched instrumentation styles (e.g. drums only), a similar principle may be applied but with spectral centroid (bass drum = lowest note, crash cymbal = highest).

2. *Note Randomization Probability:* For uncertain events only, this parameter alters the probability of the event note number being changed to a random other number. Increasing this parameter results in the uncertain component of the pattern becoming increasingly unpredictable in terms of its note progression while keeping articulation and rhythmic content the same. When applied in moderation or in combination with other forms of randomization, this parameter can contribute to adding a degree of short-term musical variation to a pattern, virtually ensuring that the exact same pattern never plays twice even if the higher-order structure of the pattern stays uniform.

3. *Note Filter Cutoff:* This comprises two discrete parameter controls that allow an upper and lower note number cutoff to be applied to a pattern, much like a high- and low-pass filter applied to an audio signal. By default, all available notes have the possibility of being played, but raising (lowering) the lower (upper) cutoff eliminates lower (higher) notes, which responsively changes both the musical complexity and spectral richness of the pattern.

Raising / lowering velocity, probability, tempo, overall pitch, pitch range, and playrate (= pitch) have a relatively clear connotation of an increase / decrease in some quantity related to amount or activity level. As all multipliers are applied with millisecond precision, the temporal resolution of these mappings is limited only by the temporal distance between successive musical events (grid unit duration). Coupling this with the dynamic nature of the output, it is likely that temporal resolution is most likely poorer, and just-noticeable differences are considerably larger than when mapping to more stationary audio signals, while potentially still being sufficiently small for certain applications.

3.3 Randomization Amount Parameters

This class of mapping parameters aims to affect the output audio signal in musically meaningful ways by manipulating the *amount* of random variation applied to various sequencer-related attributes (e.g. event triggering probability) and signal-specific attributes (e.g. envelope shape). For a given attribute a varying between 0 and a_{max} with default value a_d, mapping parameter p_m controls the amount of randomness (random value r in 0–1 range) added to a_d to yield the final attribute value a_f as:

$$a_f = (1-p_m)\cdot a_d + p_m \cdot r, \quad \text{where } a_{\text{min}} \leq r \leq a_{\text{max}} \quad (1)$$

Random value r is generated using the JUCE *Random* class in real time prior to the triggering of all certain and uncertain events, and its ‘contribution’ to the attributes of said event is governed by Eq. 1. The temporal responsiveness of these parameters during real-time application depends on the temporal density of sequenced events, which is a function of the sequencer tempo and the presence of metrical subdivisions in the sequenced pattern. To ensure appropriate attribute-specific perceptual scaling of each mapping parameter, an exponential or logarithmic scaling factor may be applied to p_m.:

1. *Velocity Randomization:* This parameter increasingly randomizes the triggering velocity of events, which affects their loudness, envelope shape, and the triggered sound sample. A small amount of velocity randomization can often serve to *humanize* the articulation of most musical patterns, but an excessive amount can cause the rhythmic stability and overall predictability of patterns to break down as rhythmic subdivisions are randomly accentuated or de-accentuated.

2. *Timing Randomization:* This parameter applies a random delay (in sequencer ticks) to every event in

the operational musical event chart. Specifically, increasing the parameter increases the absolute spread of the delay amounts, which adds random timing variability to the musical output. Like with velocity randomization, a small amount of timing variability can humanize most patterns (often desirable), but excessive variability can cause patterns to sound sloppy and ultimately arrhythmic.

3. *Playrate Randomization:* This parameter randomizes the playback rate of all triggered events. The default playback rate is 1x (a sample increment of 1 sample per audio sample interval), and increasing the parameter to its maximum value introduces a spread of 0.99x centered around 1x (resulting in a 0.01x - 1.99x range of possible sample playback rates). This manifests as increasingly random pitch and duration variations that gives percussive patterns an 'other-worldly' quality and detunes melodic patterns.

4. *Envelope Shape Randomization:* This parameter randomizes the shape of an AR (attack-release) envelope applied to each sample when it is triggered. The attack time is 1 ms, and the release time is the remaining length of the sample. By default, the envelope is rectangular (i.e. no envelope). When the parameter value is raised, the envelope shape for triggered events is randomly configured to be anything between the default rectangular envelope and an exponential envelope with increasingly sharp attack and decay. The resulting output is hence made up of sounds that randomly vary in envelope shape. A small amount of envelope shape may serve to humanize a performance, but excessively exponential envelopes can cause musical events to sound overly impulsive and unnatural.

5. *Bandwidth Randomization:* This parameter randomly varies the bandwidth of each triggered sample by manipulating the cutoff frequencies of 2nd order Butterworth low-pass and high-pass filters that operate on the sample (every sample has its own dedicated filters). By default, each sample has a bandwidth of 20 Hz – 20 kHz, but as the parameter value increases, the low-pass cutoff frequency spread increases (maximum range 500 Hz – 5 kHz) as does the high-pass cutoff spread (maximum range 20 Hz – 320 Hz). This results in a very distinct sonic special effect comparable to a traditional auto-filter.

6. *Probability Randomization:* This parameter applies an increasing degree of random variability to the probability of all uncertain chart events, which randomly alters their likelihood of triggering. As such, the audible result is a marginally less predictable musical pattern, although the effect of the parameter on the resulting output is indirect in nature as it works by manipulating triggering probability and not sound attributes directly.

7. *Sample Startpoint Randomization:* This parameter increasingly randomizes the exact audio sample index from which triggered audio events play (anywhere between the sound clip start and 40% of its length). The resulting output is a stream of auditory events that sound randomly chopped at the start, which manifests as an unusual and unnatural special effect.

The musical changes elicited by the randomization amount parameters can have various connotations that influence listener interpretations of the phenomenon under observation, and should therefore be factored into design choices. In some cases, the connotation is very clearly negative (e.g. delay / startpoint / envelope / playrate of melodic samples) that causes the music to sound out-of-tune, out-of-time or 'chopped'). In others (playrate of percussive samples, bandwidth), it is ambiguous as due to the *special effect*-like nature of the resulting musical changes that are neither clearly positive / negative nor have any clear embodied connection with the constructs of size or amount. They can potentially be used to signal specific anomalies during process monitoring tasks, but listener training will likely be necessary. Regarding perceptual resolution, although all of the above parameters are technically continuous from 0–1, the non-stationary nature of the musical output most likely makes it impossible for continuous-valued data trajectories to reliably be deduced from listening to the output even if the changes themselves are salient. A conservative estimate is that one may deduce binary or ternary states, and it therefore appears appropriate to quantize the input data variable to 2-3 discrete levels prior to mapping.

Regarding orthogonality in the perceptual domain, as the mapping parameters pertain to sequencer playback and not directly auditory perceptual attributes, the question of their perceptual independence is an interesting one. In particular, the linear multiplier parameters are conceptually distinct and may afford considerable perceptual independence from each other *(e.g. playrate multiplier – probability multiplier)* as well as several randomization amount parameters *(e.g. probability multiplier – delay randomization)*, which may be suited to representing independent variables independently. At the same time, there are several combinations that may interfere with each other *(e.g. note filter cutoff – probability multiplier, sample startpoint randomization – velocity multiplier)*. As with traditional mapping parameters, it is important to consider the implications perceptual interactions as part of the design process.

3.4 Interactive Demos

In the attached supplementary folder, several examples of real-time interactive sonification are provided. In all cases, angular velocity / orientation data from forearm-mounted inertial sensors are used to control various combinations of the aforementioned parameters. Although lacking clear use case scenarios at present, these demos serve to demonstrate the level of real-time responsive musical control afforded by the platform and parameters.

4. DISCUSSION

In this work, I have demonstrated how a probabilistic generative musical platform can serve as a medium for sonifying multivariate data. In this paradigm, data variables may be mapped to deterministic and nondeterministic properties of the generative process, ultimately conveying information about the data through variations in musical articulation, dynamics, timbre, and rhythmic patterning.

Although the platform is yet to be formally evaluated, there are several aspects of the overall approach worth discussing. Relative to the variety of generative music paradigms devised of the years [21, 22], the notion of a probabilistic step sequencer operating based on limited user input and random number generation is not altogether novel or complex. But based on initial experimentation, my current appraisal is that when configured suitably with a well-designed set of audio sample clips and initial rhythmic pattern, the platform is capable of producing musical outputs that expressively convey information in multivariate datasets. Future user studies will investigate combinations of parameters (base randomization amount values for timing, velocity, probability, note) for ensuring a balance between a predictable repetitive base structure and musically interesting short-term variations. In addition, it will also be necessary to experimentally ascertain optimal base values for the velocity and probability multipliers so as to balance task performance and musical arousal for various use cases [28]. A larger palette of instrumentation styles must be designed and rendered to suit user variability in musical taste as well as the stylistic requirements of different application domains. Finally, it will be necessary to design and evaluate discrete combinations of instrumentation styles, mapping parameters, and sequenced musical arrangements in well-defined sonification use cases.

Concerning the set of available mappable parameters, key aspects such as perceptual salience, meaningfulness, and orthogonality have briefly been remarked upon previously. Overall, it appears that depending on the mapping choices made, the generative paradigm has considerable potential for representing data changes in a salient and musically interesting manner, although the perceptual resolution of many affected musical attributes may be limited both in time and in terms of exact value. In order to quantify the informative potential of each parameter in terms of task accuracy and response time, experimental comparisons such as [12, 18] with univariate and multivariate mappings will be necessary to conduct. The current palette mappable parameters allows deterministic control of deterministic and nondeterministic musical attributes alike, opening up an interesting parameter space for musical metaphors for communication. Ultimately, application-specific user-centered studies will need to be carried out so as to ascertain meaningful combination of metaphors for different target groups. In addition, clear distinctions will need to be made empirically between parameters that should contribute to adding musical interest independently of the data, and parameters that data variables should be mapped to. To facilitate meaningful exploration, future versions of the platform will include an expanded mappable parameter set to allow further real-time manipulation of symbolic and signal-related properties alike.

5. CONCLUSION

The core concept of mapping data variables to nondeterministic aspects of generative music synthesis seems worthy of investigation in sonification applications that (a) involve long-term listening, (b) may rely on musical appeal to support the goals of the user, (c) convey a relatively small number of data variables that evolve slowly over time and can tolerate discretization to 2-3 values. Although the current work made use of a relatively basic sequencer architecture and generative paradigm, it appears so far that the biggest potential advantage of such mappings over traditional ones is that they elicit sonic changes that are both perceptually salient and musically compelling. I attribute this to the ‘predictable unpredictability’ brought about by deterministic mappings to nondeterministic components of the music generation process. Using a probabilistic multitrack sequencing paradigm affords symbol-level manipulation possibilities simply not achievable when mapping to pre-recorded stereo mixes / stems. This facilitates the creation of multiple perceptually orthogonal mappings that are clearly separable not only from each other but also from the short-term evolution and variability of the music itself. In conclusion, I believe that skilful data-driven manipulation of determinism and nondeterminism can effectively unleash the latent potential of generative music as a data communication medium.

Acknowledgments

I would like to thank my supervisors Sofia Dahl and Erika G. Spaich for continuing to support risky creative exploration in times of significant uncertainty.

6. REFERENCES

[1] T. Hermann, “Taxonomy and definitions for sonification and auditory display,” in *Proceedings of the 14th International Conference on Auditory Display (ICAD), Paris, France, June 24 - 27*, 2008, pp. 1–8.

[2] P. Vickers, “Sonification and music, music and sonification,” in *The Routledge Companion to Sounding Art*. Routledge, 2016, pp. 135–144.

[3] P. Vickers, B. Hogg, and D. Worrall, “Aesthetics of sonification: Taking the subject-position,” in *Body, sound and space in music and beyond: multimodal explorations*. Routledge, 2017, pp. 89–109.

[4] J. G. Neuhoff, “Is sonification doomed to fail,” in *Proceedings of the 25th International Conference on Auditory Display (ICAD)*, 2019, pp. 327–330.

[5] J. Middleton, J. Hakulinen, K. Tiitinen, J. Hella, T. Keskinen, P. Huuskonen, J. Culver, J. Linna, M. Turunen, M. Ziat *et al.*, “Data-to-music sonification and user engagement,” *Frontiers in big Data*, vol. 6, p. 1206081, 2023.

[6] N. Rönnberg, "Sonification for conveying data and emotion," in *Proceedings of the 16th International Audio Mostly Conference*, 2021, pp. 56–63.

[7] G. Weinberg and T. Thatcher, "Interactive sonification: Aesthetics, functionality and performance," *Leonardo Music Journal*, vol. 16, pp. 9–12, 2006.

[8] B. N. Walker and M. A. Nees, "Theory of sonification," *The sonification handbook*, vol. 1, pp. 9–39, 2011.

[9] F. Grond and T. Hermann, "Interactive sonification for data exploration: How listening modes and display purposes define design guidelines," *Organised Sound*, vol. 19, no. 1, pp. 41–51, 2014.

[10] P. Vickers, "Ars informatica – ars electronica: Improving sonification aesthetics," in *Understanding and Designing for Aesthetic Experience (Workshop at HCI 2005: The 19th British HCI Group Annual Conference)*, Edinburgh, UK, September 2005.

[11] P. Kantan, E. G. Spaich, and S. Dahl, "A technical framework for musical biofeedback in stroke rehabilitation," *IEEE Transactions on Human-Machine Systems*, vol. 52, no. 2, pp. 220–231, 2022.

[12] T. Hildebrandt, T. Hermann, and S. Rinderle-Ma, "Continuous sonification enhances adequacy of interactions in peripheral process monitoring," *International Journal of Human-Computer Studies*, vol. 95, pp. 54–65, 2016.

[13] B. Van Kerrebroeck and P.-J. Maes, "A breathing sonification system to reduce stress during the covid-19 pandemic," *Frontiers in Psychology*, vol. 12, p. 623110, 2021.

[14] S. Roddy and B. Bridges, "Addressing the Mapping Problem in Sonic Information Design through Embodied Image Schemata, Conceptual Metaphors, and Conceptual Blending," *Journal of Sonic Studies*, no. 17, 2018.

[15] H. Feinstein, "Meaning and visual metaphor," *Studies in art education*, vol. 23, no. 2, pp. 45–55, 1982.

[16] D. Worrall, "Sonification: A prehistory," in *Proc. of the 24th Int. Conf. on Auditory Display (ICAD)*, 2018, pp. 177–182.

[17] L. Stockmann, A. Berndt, and N. Röber, "A musical instrument based on 3d data and volume sonification techniques," in *Proc. of the Audio Mostly Conference*, 2008, pp. 72–79.

[18] P. R. Kantan, "Comparing sonification strategies applied to musical and non-musical signals for auditory guidance purposes," in *19th Sound and Music Computing Conference, SMC 2022*. Sound and Music Computing Network, 2022, pp. 279–286.

[19] P.-J. Maes, J. Buhmann, and M. Leman, "3MO: A Model for Music-Based Biofeedback," *Frontiers in Neuroscience*, vol. 1, 12 2016.

[20] P. R. Kantan, E. G. Spaich, and S. Dahl, "A metaphor-based technical framework for musical sonification in movement rehabilitation," in *The 26th International Conference on Auditory Display (ICAD)*, 2021.

[21] A. Dash and K. Agres, "Ai-based affective music generation systems: A review of methods and challenges," *ACM Computing Surveys*, vol. 56, no. 11, pp. 1–34, 2024.

[22] S. Hunt, C. Nash, and T. Mitchell, "Thoughts on interactive generative music composition," in *2nd Conference on Computer Simulation of Musical Creativity*, 2017.

[23] P. N. Johnson-Laird, "Freedom and constraint in creativity," *The nature of creativity: Contemporary psychological perspectives*, vol. 202, 1988.

[24] P. Sloan, "Listening to possibility: Randomness as musical material," Master's thesis, Mills College, 2014.

[25] L. Kiss and K. J. Linnell, "Making sense of background music listening habits: An arousal and task-complexity account," *Psychology of Music*, vol. 51, no. 1, pp. 89–106, 2023.

[26] A. T. Rödel, "Listening to background music while studying-emotional drive or cognitive overload?" Master's thesis, University of Twente, 2021.

[27] E. H. Margulis, *On repeat: How music plays the mind.* Oxford University Press, 2013.

[28] A. B. Ünal, D. de Waard, K. Epstude, and L. Steg, "Driving with music: Effects on arousal and performance," *Transportation research part F: traffic psychology and behaviour*, vol. 21, pp. 52–65, 2013.