Measuring Infinity: Autonomy in David Dunn's *Thresholds and Fragile States*

David Kant

University of California, Santa Cruz
david.kant@gmail.com

Abstract. Composer David Dunn's 2010 piece *Thresholds and Fragile States* is an autonomous network of analog circuitry designed to generate an endless and constantly evolving variety of musical sound. Like many electronic works in the American experimental tradition, the music is difficult to analyze because compositional decisions are embedded within layers of circuit design, and musical features are explained at the level of individual electrical components. In this paper, I analyze Dunn's piece using tools from circuit analysis, digital modeling, and machine learning, and present two computational experiments designed to study how autonomy arises from it. I find that the endless variety of constantly evolving music is the result of hysteresis, or time lag of the vactrol, a circuit component and common "kludge" in electronic music, and I find that the extreme unpredictability that Dunn describes occurs within a small region of carefully tuned network settings that Dunn has discovered empirically through musical performance.

Keywords: Autonomy, Generative systems, Machine learning, AI, Music analysis, Digital modeling

1 Introduction

Composer David Dunn's 2010 piece *Thresholds and Fragile States* (TFS) (Dunn, 2012) is a sprawling tangle of wires, integrated circuits, potentiometers, and capacitors that, when switched on, produces electronic soundscapes that are complex, evolving, and musical. Motivated by his experience listening to the patterns of insects on a swamp in the Atchafalaya Basin of Louisiana in addition to his study of cybernetics, the concept of autonomy is central to the piece. Dunn has explored the phenomena of autonomy throughout his career, although often under different names or metaphors, from large-scale, outdoor site-specific works of the 1970s, which sought to connect with the spirit of a location, to later explorations of chaotic and nonlinear dynamical systems, began in the late 1980s, in which Dunn networked hybrid systems of digital computers and analog circuitry. Running throughout these inquiries is an infatuation with life and mind, concepts that are closely related, perhaps even equivalent, for Dunn, with autonomy, the central object of my analysis project.[1]

[1] For an overview of Dunn's work and the metaphors of mind that run throughout it, see (Heying & Kant, 2018).

Dunn's approach to electronic musical instrument design is characteristic of a tradition of experimentalism that developed in the United States and abroad starting in the mid-twentieth century. Artists and instrument makers, such as David Tudor, Don Buchla, Eliane Radigue, and Bebe and Louis Barron, built circuit systems capable of producing complex and unpredictable sonic behaviors, systems that were less akin to traditional instruments but rather exhibited intrinsic tendencies, character, and agency with which the performer had to contend, systems that were more akin to life. Like many electronic works in this aesthetic tradition, Dunn's creative decisions are embedded within layers of circuit design and musical features are explained at the level of individual electrical components.

How can we reason about the creative decisions involved in designing such autonomous generative systems, especially those that are complex and unpredictable? In this article, I present an analysis of Dunn's circuit network, combining tools from circuit analysis, digital modeling, and machine learning to study how autonomy arises from it. I provide an overview of Dunn's system, describe the implementation of a digital realization of it, and detail two computational experiments. In the experiments, I analyze the scope of the emergent music, mapping the parameter ranges that produce it, and identify the source of long-term change and variation. This article points to a large, multi-faceted project that I conducted in close communication with Dunn. My work is based Dunn's published writings, personal and published interviews, score diagrams, schematics, and measurements of Dunn's circuitry.

2 The Nonlinear Chaotic Oscillator

In the score for TFS, Dunn describes the design and implementation of a custom-built, autonomous circuit network that is inspired by the self-organizing behavior of acoustic biological ecosystems. Dunn explains that the circuitry—a modular synthesis style feedback network of coupled nonlinear chaotic oscillators—is based on the underlying design philosophy of *autopoiesis*, a theory proposed by Chilean biologists and neuroscientists Humberto Maturana and Francisco Varela in the 1970s (Varela, 1979) to explain the self-organizing phenomena of cellular organisms, and produces a seemingly endless and constantly evolving variety of ever-changing electronic soundscapes.

The building block of Dunn's piece is the nonlinear chaotic oscillator. In Dunn's system, multiple chaotic oscillators (up to eight) are networked together in a complex and hierarchical modulation system, including various nonlinearities, filters, and gainstages with additional low-frequency oscillators. Individually, each chaotic oscillator is capable of producing an array of sounds, from periodic sinusoidal waveforms to quasi-periodic oscillations and band-limited noise. Networked together, the system produces sounds that are not exhibited by any single oscillator in isolation but rather are emergent from the complex, audio-rate modulations.

Dunn's oscillator has a single control parameter, a variable resistor `R` which, when adjusted, brings the oscillator through a complex and varied sequence of sounds—illustrated in Figure 1.[2] Changes in pitch, amplitude, and spectra are modulated simultaneously, mapped together along this single modulation parameter. Dunn's design philosophy runs counter to conventional engineering epistemologies, in which electronic instruments are carefully designed to have semantically interpretable controls that respond in meaningful and predictable ways. Contrary to conventional oscillators, which typically produce simple, repeating waveforms, and are generally engineered to respond proportionally to changes in input, Dunn's oscillator, by comparison, can produce disproportionately large changes in output for even the smallest adjustment to a control parameter. This makes it difficult to control, let alone reason about the design decisions involved in constructing the entire circuit system.

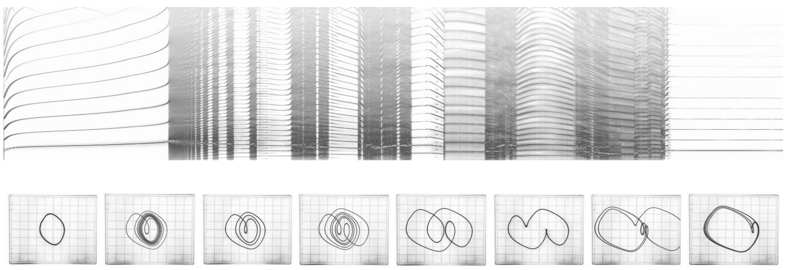

Fig. 1. Spectrogram (above) of Dunn's chaotic oscillator as the control parameter `R` is swept and the sequence (below) of fixed, single, and double-scroll attractors produced.

3 The Feedback Network

Central to Varela and Maturana's theory of autopoiesis is the idea of "structural coupling," which describes how a system can be both operationally closed, or distinct from its environment, internally maintaining the processes necessary to sustain itself and constitute a whole, while at the same time coupled to its environment, part of a dynamic system and sensitive to interaction. Dunn's circuitry reflects this in the design of the feedback pathways between chaotic oscillators, which are coupled together hierarchically in pairs. Within each coupled pair, the oscillators do not share voltage signals directly, but rather interact through the control parameter `R`. The modulation signals are not mapped directly to resistance values in a linear fashion, but are passed through a complex signal path, including filtering, amplification, and other nonlinear operations necessary

[2] Additional materials, including audio and video recordings of the chaotic oscillator and digital model can be found at `http://davidkantportfolio.com/aimc-2021`.

to implement the connection in analog form, producing a modulation map that is nonlinear and time-variant.

A critical component in the feedback path is the vactrol, or opto-isolator. Consisting of a light-emitting diode (LED) taped end-to-end with a photosensitive resistor (photoresistor), the vactrol is a common "kludge" in electronic music for allowing voltage control of resistive circuit components. The vactrol, however, introduces additional complexities into the system, which can be desirable (such asymmetric attack and decay) and, in Dunn's circuitry, are crucial to the emergent behavior. Specifically, the photoresistor exhibits *hysteresis*, or time-lag. The sensitivity of the photoresistor slowly adjusts over time, dependent not just on the current signal but on its past exposure as well, which provides a very low-frequency modulation to the entire system.

Another interesting feature of the feedback path is the inclusion of low-frequency oscillators. The LFO signals are combined with the modulation signals, one for each chaotic oscillator, immediately before the amplification stages. Implemented using a 555-based square wave circuit, the LFOs saturate the mixer, temporarily nullifying the modulation signal. In Dunn's electronic microcosm, the LFOs serve the role of perturbation, a concept critical to the emergence of autonomy in Varela and Maturana's theory. These LFO perturbations can cause the oscillators to radically reorganize or have little-to-no audible effect, depending on the current state of the system.

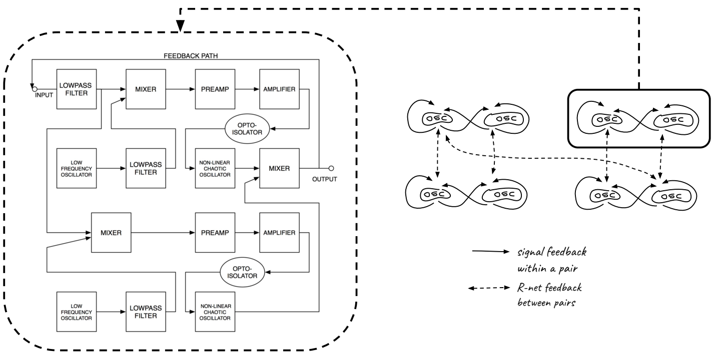

Fig. 2. Network topology (right) and signal flow within a coupled pair (left).

4 Modeling *Thresholds and Fragile States*

One of the central questions Dunn asks in the score is whether or not the emergent behaviors produced by the oscillator network are specific to its instantiation in analog form or endemic to the abstraction of coupled chaos. In response to this question—as well as in effort to better understand the instrument, the design decisions involved, and how it works—I built a digital realization.

I created two implementations of the digital realization. The first, built in SuperCollider, is intended for real-time performance and contains a collection of GUIs to visualize system state, organize control parameters, and inspect the signal flow through each node of the network.[3] The second, implemented in Rust, runs offline and is used to generate the large datasets necessary for machine-learning analysis.[4]

While the details of the digital model can be found in a github repository, two components in particular are worth mentioning here because they are critical to the model's behavior: the nonlinear chaotic oscillator and vactrol. Dunn's nonlinear chaotic oscillator belongs to a family of jerk systems published by physicist J.C. Sprott (Sprott, 2000). The circuit can be described by a third order differential equation having three integration stages and a feedback nonlinearity, $sgn(x)$, with choice of circuit component values determining the values of the equation coefficients A, B, and C:

$$w = -Az - y - Bx + C\,\mathrm{sgn}\,(x)\ . \tag{1}$$

When modeling the circuitry, however, Dunn's oscillator deviates from the equations given. In order to model the behaviors produced by Dunn's analog oscillator, it is necessary to modify the equations in a number of ways. The deviations are due to non-ideal properties of Dunn's op-amps, which introduce additional nonlinearities such as saturation and slew limiting, as well as the choice of circuit components and topologies that realize key mathematical operations—the use of a passive RC filter to implement the middle integration stage (perhaps in order to fit the circuit on a quad op-amp chip) and lack of buffering which yields significant self-modulation.

The implementation of the vactrol is also critical to the digital realization. The vactrol is modeled using an asymmetric exponential filter, with independent parameters for attack and decay rates, and the filter response curves are fit to measurements taken from RadioShack components similar to those used by Dunn. I found it necessary to fit the vactrol response to measurement data in order to produce behaviors similar to Dunn's system. The vactrol hysteresis is implemented with an additional exponential filter, tuned to a much lower frequency range, which modulates the vactrol sensitivity.

5 Two Studies in Autonomy

Dunn's fascination with autonomy, his interest in the chaotic and unpredictable, and his desire for a system capable of producing an ever-changing sonic output led me to wonder about the design decisions involved. Is there a way to more precisely analyze such a piece? A few questions came to mind: How unpredictable is Dunn's system? In what ways? What causes the variety of sounds produced? In this final section, I detail the design of two computation experiments intended to study the source and character of Dunn's autonomy.

[3] Available at: `https://github.com/davidkant/thresholds`.
[4] Built in collaboration with Andrew Smith.

5.1 Study 1: Auditory Neotaxis

Dunn's instrument exhibits a tremendous capacity for continued change and variation over very long timescales, on the order of many minutes to many hours. The instrument seems to be in a state of constant flux, slowly meandering through a diverse possibility space without human intervention. What is the source of this long-term change? In the first experiment, *Study 1: Auditory Neotaxis*, I use multivariate regression analysis, a common statistical technique for studying relationships between variables, to find correlation between network parameters and long-term change.

The experiment dataset consists of 100,000 examples, each rendered for 60-minutes of audio, with parameter values chosen uniformly at random. The randomized parameters contain both *playable parameters*, which Dunn has exposed as potentiometers and switches on the top plate of the instrument for control in performance, as well as *design parameters*, which determine aspects of the circuit design not made variable in performance, such as the vactrol attack and decay response tines—fixed physical propertiess of the electronic components.

Long-term change is measured using a feature I refer to as the *trend spectrogram*. An application of trend analysis to audio, the trend spectrogram measures raw change in audio magnitude spectra over time. It is computed by first transforming a time-domain audio signal to the frequency-domain and then reducing the spectrogram to a single dimension by Principal Component Analysis (PCA). The trend spectrogram is given by performing Fourier analysis on the PCA-reduced audio spectrogram and aggregating bins into perceptually spaced time intervals, giving timescales, in this case, from 1 second to 60 minutes.

The results of the regression analysis are given in Figure 3. The 48 network parameters comprise the independent variables and the energy of the 24 trend spectrogram bins the dependent variables. Two variables in particular were found to have significant correlation with long-term change, as indicated by a low p-scores (values below 1×10^{-9}) and high correlation coefficients. At short timescales, up to a few minutes in duration, the LFO is the primary determining factor. However, as the analysis timescale increases, starting at a few minutes and continuing all the way up to 60 minutes, the vactrol hysteresis becomes the primary determining factor, showing increasingly larger coefficients and decreasingly smaller p-scores as the length of analysis timescale grows. This suggests the vactrol hysteresis is the primary source of long-term change.

These results confirm Dunn's intuition that the vactrol hysteresis is a major cause of long-term change, as it provides a very low frequency modulation to the entire system, slowly driving the oscillator network through novel regions of behavior. Long-term change, a critical emergent feature of Dunn's instrument and one that is closely related to the sense of autonomy, can be produced by the digital implementation and is localized in a single circuit component, the vactrol.

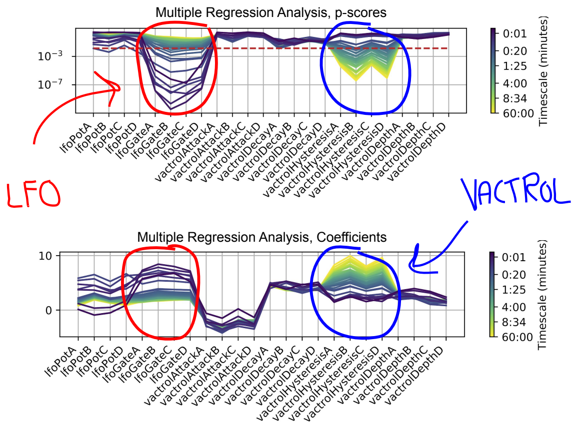

Fig. 3. Multivariate regression analysis for LFO and vactrol parameters.

5.2 Study 2: Mapping the Possibility Space

Dunn's instrument also produces tremendous variation in sound behavior, as reflected in both the timbres present as well as rhythmic patterning and greater temporal structures. However, even within this potentially limitless variation, there is an identifiable character and quality to the music. How large and how varied is this possibility space? How do we describe the character of an infinite—or near infinite—variation? The second experiment, *Study 2: Mapping the Possibility Space*, is designed to answer these questions, through the use of manifold learning to analyze the structure of emergent possibility spaces.

The experiment consists of two datasets, the first, similar to *Study 1*, contains 100,000 examples, sampled uniformly at random from the entire parameter space of the model. This dataset represents the full possibility space of the system. The second dataset, however, is sampled from a constrained parameter distribution to reflect Dunn's performance practice. In performance, Dunn adjusts certain parameters more frequently than others, leaving some untouched, either constrained to narrow ranges or left entirely fixed. Each example is rendered for three minutes, and a collection of audio features are extracted using the *Essentia Freesound Extractor* (Bogdanov et al., 2013). Audio features include a variety of time-varying descriptors and their statistical summaries (e.g., min pitch, max pitch, and pitch mean), together with first and second order differences. Ultimately, each example is described by a list of 550 values.

The datasets are visualized by Uniform Manifold Approximation and Projection (UMAP) (McInnes, Healy, & Melville, 2020), a machine-learning technique commonly used to reduce dimensionality of and visualize structure in large datasets. UMAP finds a low-dimensional embedding (two dimensions) that has similar topological structure to the original data, in this case having 550 dimensions—far too many to visualize directly. Embeddings help illustrate similarity relationships between data points and can reveal topological structure, such as whether or not the points are grouped into one big connected component or scattered across many small disconnected clumps.

Figure 4 shows UMAP embeddings for both datasets. To aid the visualization, examples are grouped into 20 sound categories, using k-means clustering, and color-coded. Comparing the plots, the constrained dataset exhibits a greater degree of fine structure (many smaller clumps) as well as a greater diversity of sound categories (more colors). The major components of the unconstrained dataset can even be seen within the constrained dataset.

These observations suggest that Dunn has found—empirically, through hands-on interaction with the system—a small region of carefully tuned network settings that exhibit greater variety and unpredictability. Dunn's parameter settings can be seen as an expression of aesthetic preference and intention. The embedding analysis shows that, while the circuitry has a vast possibility space, Dunn has carved out a subset with drastically different character, which reflects back on and is informative of his notion of autonomy.

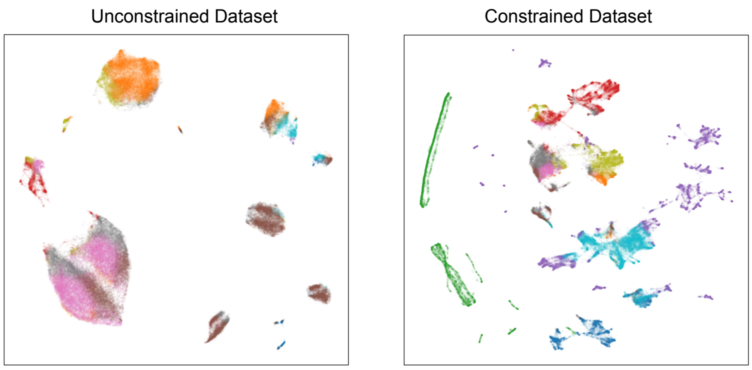

Fig. 4. Embedding plots for the unconstrained and constrained datasets.

6 Conclusions

Machine learning has tremendous potential to help us understand one of the most important paradigm shifts in creative practice of the past century: a shift towards generativity. There is a rich history of music, from the chance procedures of John Cage and Christian Wolff to the electronic feedback circuitry of David Tudor and Bebe and Louis Barron to the present day proliferation of generativity in mainstream creative culture, that can benefit from a better understanding of what it means to create with generative systems. Beyond its predictive power, machine learning can help elucidate artistic process, reason about cause and effect, identify the musical implications of design decision, and better understand artistic intuition and intent when creating with generative systems.

References

Bogdanov, D., Wack, N., Gómez, E., Gulati, S., Herrera, P., Mayor, O., ... Serra, X. (2013). Essentia: An open-source library for sound and music analysis. In *Proceedings of the 21st acm international conference on multimedia* (p. 855–858). New York, NY, USA: Association for Computing Machinery. doi: 10.1145/2502081.2502229

Dunn, D. (2012). Thresholds and fragile states: A descriptive score for a generative sound system using networked autonomous chaotic oscillators. *Moebius*, *1*.

Heying, M., & Kant, D. (2018). The emergent magician: Metaphors of mind in the work of david dunn. *Sound American*, *18: The David Dunn Issue*. Retrieved from https://soundamerican.org/issues/place/emergent-magician-metaphors-mind-work-david-dunn

McInnes, L., Healy, J., & Melville, J. (2020). *Umap: Uniform manifold approximation and projection for dimension reduction.*

Sprott, J. (2000, August). Simple chaotic systems and circuits. *American Journal of Physics*, *68*(8), 758-763.

Varela, F. (1979). *Principles of biological autonomy.* New York, New York: Elsevier North Holland.