

UMWELTNOISE: Audiovisualization Through Non-Human Perception

Yana Mikhailava

OCAD University, Toronto, ON, CA

myana@ocadu.ca

ABSTRACT

Audio visualization is a practice shaped heavily by human perception where the artist chooses elements which highlight their own subjective reading of a piece of sound. UMWELTNOISE proposes an alternative: a system in which sound is analyzed and rendered visually not through human interpretation, but through the behavioral logic of *Physarum polycephalum*, a brainless slime mould observed in unconventional computing research to solve complex spatial problems through nutrient-seeking growth. This project draws on Jakob von Uexküll's idea that every organism inhabits a distinct, species-specific perceptual world in order to create an unconventional GLSL shader based on *Physarum polycephalum*'s foraging and growth logic, treating certain frequencies as nutrients and others as deterrents [1]. The result is a form of audiovisualization that has the potential to expand human perception and musical interpretation through machine vision guided by non-human biological intelligence, raising questions about what might be discovered through this three-way collaboration between organism, algorithm, and human co-author.

1. INTRODUCTION

Audio visualization is conventionally a human-authored practice. The elements a visualizer selects to foreground, whether beat tempo, melodic contour, or timbral texture, reflect an inherently subjective reading of the source material. This raises a question that is rarely made explicit: what is being missed? Different listeners attend to different qualities in the same piece of music, and the visualization produced by one person may foreground elements that another would not have noticed, or overlook dimensions that another finds central.

UMWELTNOISE approaches this problem through the theoretical framework of the *umwelt*, introduced by biologist and philosopher Jakob von Uexküll [1]. Uexküll proposed that every organism inhabits a distinct, species-specific perceptual world shaped entirely by its own sensory and biological apparatus. No two organisms experience the same environment in the same way. Building on this, UMWELTNOISE extends this logic to a machine-made algorithm, introducing an additional layer of machine vision into this layered translation between sound, biological logic, and visual output. What might audiovisualization look like if the perceptual logic

driving it were a collaboration between non-human biological intelligence, machine vision, and a human co-author?

The organism selected for this purpose is *Physarum polycephalum*, a single-celled slime mould that has been a point of interest within unconventional computing research for its capacity to solve complex pathfinding and optimization problems without a brain or nervous system [3, 4]. By emulating *Physarum*'s nutrient-seeking growth behaviour within a GLSL shader in TouchDesigner, this project attempts to produce audiovisualization that is structured by biological rather than cultural or aesthetic logic.

Figure 1: *Physarum Polycephalum*




Image Source: Katia Moskvitch. 2018. Slime molds remember — but do they learn? *Quanta Magazine*.
<https://www.quantamagazine.org/slime-molds-remember-but-do-they-learn-20180709/>

2. RELATED WORK

2.1 Umwelt Theory and Biosemiotics

The theoretical foundation of UMWELTNOISE rests on Jakob von Uexküll's concept of the *umwelt* [1], which proposes that every organism constructs its own subjective world from the environmental stimuli it is biologically equipped to perceive. Carlo Brentari's biosemiotic analysis of Uexküll's work [2] extends this framework by examining how organisms produce meaning from environmental signals, a process that UMWELTNOISE seeks to simulate through *Physarum*'s differential response to audio frequencies.

2.2 Unconventional Computing and *Physarum polycephalum*

Jeff Jones's research into agent-based approximations of *Physarum* transport networks [3, 4] provides the computational model underpinning the project's shader logic. Jones demonstrates that the organism's complex emergent growth patterns can be reproduced through populations of simple agents following local chemotaxis rules, and catalogues the specific environmental mechanisms, including attractants, repellents, and obstacles, that can be used to shape network formation. These mechanisms directly inform how audio data is mapped onto the shader's behaviour in UMWELTNOISE.

2.3 Technical Images and Machine Vision

Vilém Flusser's writing on technical images [5] argues that images produced by apparatus are not neutral representations but carry the encoded logic of the systems that generate them. This resonates with UMWELTNOISE's central concern: that conventional audiovisualization encodes human perceptual bias into its output. Jill Walker Rettberg's work on machine vision [7] extends this further, examining how algorithmic systems see and interpret the world according to their own logic, frequently diverging from human perception in ways that are both revealing and consequential.

3. CONCEPT AND DESIGN

UMWELTNOISE's central design proposition is that certain audio frequencies function as nutrients within the simulation, producing smooth, generative growth patterns, while others act as deterrents, introducing distortion and disruption into the visual field. This mapping is drawn directly from *Physarum*'s documented behaviour: in foraging, the organism expands toward nutrient sources while avoiding inhospitable conditions, leaving behind a network that encodes the density and distribution of sustenance in its environment.

The project connects to broader areas of inquiry including unconventional computing, machine vision, biomimicry, and the emerging field of non-neural intelligence. It also situates itself within a lineage of technically driven visual art practice, particularly the work of Kat the Poet Engineer, an artist whose TouchDesigner-based practice incorporates multiple programming languages to

explore unconventional forms of representation. As suggested by her self-identification as a "poet," her work, while highly technical, operates primarily within visual exploration, where function may emerge as a result but is not the primary goal. UMWELTNOISE shares this orientation: the process of exploration is central to the project rather than any fixed outcome.

4. IMPLEMENTATION

4.1 Process Timeline

Stage 1: Conceptual Research

Engagement with Jakob von Uexküll's concept of the *umwelt* [1] and Carlo Brentari's biosemiotic framework [2] introduced the idea of species-specific perception as a lens for rethinking audiovisualization. Alongside this, Jeff Jones's research into *Physarum polycephalum* as an unconventional computing substrate [3, 4] proposed the organism as a viable non-human model for data processing. Vilém Flusser's writing on technical images [5] and Jill Walker Rettberg's work on machine vision [7] further grounded the project's interest in non-anthropocentric modes of seeing and representation.

Stage 2: Initial Approach in TouchDesigner

With a conceptual framework in place, the project moved into TouchDesigner. The initial instinct was to simulate *Physarum* behaviour through instancing, a familiar method from prior audiovisualization work. This approach was set aside after it became clear that instancing would not support the fluid, real-time audio-reactive behaviour the project required.

Stage 3: Discovery of GLSL Shaders

Research into alternative methods surfaced the possibility of writing custom GLSL shaders via TouchDesigner TOPs. A GLSL shader is a program written in the OpenGL Shading Language that runs directly on the GPU to control how pixels are rendered on screen. This opened a more direct path to simulating *Physarum*'s growth patterns at the pixel level, with greater control over how audio parameters could influence the shader's behaviour in real time. GLSL was a particularly suitable choice given its cross-platform compatibility and native support within TouchDesigner.

Stage 4: Learning GLSL

With no prior coding experience, structured learning became a necessary step. Lake Heckaman's shading course provided a foundational understanding of the OpenGL Shading Language and GPU-based rendering logic, forming the technical basis for the shader development that followed.

Stage 5: Development with Cursor AI and Jascha Süss


Prior experience with TouchDesigner had been largely confined to instancing and particle-based geometry through POPs, which formed the basis of earlier audiovisualization work. Although the OpenGL Shading Language was entirely unfamiliar at the outset, the integration of Cursor AI into the development workflow made it possible to experiment with shader coding iteratively. This process was further guided by Jascha Süss's work on shader development within TouchDesigner, which informed the creative and technical direction of the implementation.

4.2 Version 1

The first version of the shader prioritised expressing the visual appearance of *Physarum polycephalum* as observed in a petri dish environment. While the result produced a sound-reactive simulation aesthetically aligned with the goal, the relationship between the visual output and the underlying audio analysis remained simple. This version was developed without external code references, relying entirely on iterative prompting within Cursor AI, systematic bug resolution, and further adjustment within TouchDesigner. Through this process, significant progress was made in understanding how to assign vector uniforms and determine which shader parameters could be driven by audio qualities extracted in TouchDesigner.

A key conceptual limitation became apparent at this stage. Despite the successful adaptation of an unconventional visual construction method, the approach remained constrained by human-authored audio compositions and the built-in analysis tools available within TouchDesigner. This raised a fundamental question about the project's alignment with the goals of UMWELTNOISE: if the audio interpretation layer is still structured by human-made tools and human compositional logic, the resulting visualization may not differ meaningfully from conventional audiovisualization in its underlying epistemology. It became clear that a closer alignment with the project's intent would require developing a new approach to audio interpretation itself, not only to its visual representation.


Figure 2: TouchDesigner Structure (Version 1)



In this screenshot of the Version 1 patch, the uniforms which govern the reactive parameters on the main GLSL shader are visible on the right, where:

```
uniform float uAnimSpeed; // 0 = frozen pattern; 1 = normal time  
uniform float uScale; // Overall vein density  
uniform float uEvolve; // How fast the field drifts in time  
uniform float uWarp; // Warp strength (organic bending)  
uniform float uOctaves; // depth (integer steps ranging 1-6)  
uniform float uRidge; // Thin vs thick filaments (higher = thinner highlights)
```

Figure 3: Visual Experimentation




This series of five rendered outputs shows variations in shader parameters during early development, including changes to vein density, colour, and warp strength.

4.3 Version 2

The second version moves away from direct visual reference to *Physarum polycephalum*, achieving instead a closer alignment with the conceptual spirit of the project. In this iteration, the density of lines within the shader is driven by sound amplitude, while frequency range is used to distinguish between constructive and destructive sonic qualities. Lower-mid range frequencies result in smooth, rounded forms functioning as a visual analogue for nutritive stimuli, while higher frequencies introduce distortion and disruption into the visual field. The generated paths also act as attractors for one another. This mapping more directly reflects the biological logic of *Physarum*'s nutrient-seeking behaviour. Version 2 more readily omits certain aspects of sound while focusing more strongly on others, rather than seeking to represent every amplitude change in the input audio. While it represents a meaningful step toward an audiovisualization system governed by unconventional perceptual criteria, it remains a collaboration with the human author, as designing an entirely new algorithm for sound analysis fell outside the scope of this project.

Figure 4: TouchDesigner Structure (Version 2)



In this screenshot of the Version 2 patch, the uniforms which govern the reactive parameters on the main GLSL shader are visible on the right, where:

uniform float uAnimSpeed; // 0 = frozen pattern; 1 = normal time

uniform float uScale; // Overall vein density


uniform float uEvolve; // How fast the field drifts in time

uniform float uWarp; // Warp strength (organic bending)

uniform float uOctaves; // depth (integer steps ranging 1-6)

uniform float uRidge; // Thin vs thick filaments (higher = thinner highlights)

Figure 5: Version 2 Renders



Comparison depicting the "Disruption" of the attractors under high-frequency conditions (left: low-frequency input; right: high-frequency input).

5. DISCUSSION

The greatest limitation encountered through this project was an oversight in the initial framing: human-made musical composition, recording technology, and audio software are designed almost entirely around what is perceptible to the human ear. Microphones are built to capture frequencies within the range of human hearing, and the analysis tools available within software environments operate along the same constraints. The system, from recording to visualization, is anthropocentric by design.

Despite this limitation, the project yielded a meaningful discovery. Audiovisualization, whether intentionally or not, appears capable of communicating qualities of sound to a human observer without requiring them to hear it directly. As Rettberg observes, "machine vision changes what humans can see" [7], and this project gestures toward one concrete application of that idea. If an audiovisualization algorithm such as the one developed here were applied to sound wavelengths that fall outside the range of human perception, it could function as a bridge between human awareness and a sensory world that would otherwise remain entirely inaccessible. While tools such as oscilloscopes already allow for the technical visualization of frequencies beyond human hearing, they represent those sounds only in terms of their measurable properties. An approach informed by the logic of UMWELTNOISE could instead attempt to represent the qualitative character of those sounds, bringing human perception into closer contact with dimensions of the world that biology has placed beyond its reach.

6. CONCLUSION

UMWELTNOISE set out to explore whether the perceptual logic of a non-human organism could serve as a generative framework for audiovisualization, producing images of sound that are not structured by human musical interpretation. Through the development of a *Physarum*-inspired GLSL shader in TouchDesigner, the project demonstrated that such an approach is technically feasible and conceptually productive, even as it encountered the fundamental constraint that the tools available for audio capture and analysis are themselves built around human perception. The primary outcome is the shader itself, alongside the process and logic used to develop it. This project represents an early-stage investigation, and future development would focus on designing custom audio analysis methods that operate outside of humanly perceptible frequency ranges, extending the project's reach toward a more genuinely non-anthropocentric audiovisualization system.

REFERENCES

- [1] Jakob von Uexküll. 2010. *A Foray into the Worlds of Animals and Humans: with A Theory of Meaning* (Posthumanities, Vol. 12). Trans. Joseph D. O'Neil. University of Minnesota Press, Minneapolis. (Originally published 1934.)
- [2] Carlo Brentari. 2015. *Jakob von Uexküll: The Discovery of the Umwelt between Biosemiotics and Theoretical Biology* (Biosemiotics, Vol. 9). Springer, Dordrecht. DOI: <https://doi.org/10.1007/978-94-017-9688-0>
- [3] Jeff Jones. 2010. Characteristics of pattern formation and evolution in approximations of *Physarum* transport networks. *Artificial Life* 16, 2, 127–153. DOI: <https://doi.org/10.1162/artl.2010.16.2.16202>
- [4] Jeff Jones. 2015. Mechanisms inducing parallel computation in a model of *Physarum polycephalum* transport networks. *Parallel Processing Letters* 25, 1, 1540004. DOI: <https://doi.org/10.1142/S0129626415400046>
- [5] Vilém Flusser. 2011. *Into the Universe of Technical Images* (NED-New edition). University of Minnesota Press, Minneapolis.
- [6] Katia Moskvitch. 2018. Slime molds remember — but do they learn? *Quanta Magazine*. Retrieved April 13, 2026 from <https://www.quantamagazine.org/slime-molds-remember-but-do-they-learn-20180709/>
- [7] Jill Walker Rettberg. 2023. *Machine Vision: How Algorithms Are Changing the Way We See the World*. Polity Press, Cambridge.