XEPA - Autonomous Intelligent Light and Sound Sculptures That Improvise Group Performances

Philip Galanter
Artist, researcher, educator
Department of Visualization
Texas A&M University
C108 Langford Center
3137 TAMU
College Station, Texas 77843-3137
galanter@viz.tamu.edu

Philip Galanter

ABSTRACT

XEPA anticipates a future where machines form their own societies. Going beyond mere generative art, machines will exhibit artistic creativity with the addition of artistic judgment via computational aesthetic evaluation. In such a future our notions of aesthetics will undergo a radical translation. The XEPA intelligent sculptures create animated light and sound sequences. Each sculpture "watches" the others and modifies its own aesthetic behavior to create a collaborative, improvisational performance. No coordination information or commands are used. Each XEPA independently evaluates the aesthetics of the other sculptures, infers a theme or mood being attempted, and then modifies its own aesthetics to better reinforce that theme. Each performance is unique and widely varied. XEPA is an ever-evolving artwork, intended as a platform for ongoing experiments in computational aesthetic evaluation.

Introduction

The name XEPA is intentionally playful. First, following a tradition from the world of computer programing, XEPA is a recursive acronym. It stands for "XEPA Emerging Performance Artist." In addition, this is a reference to the emergent nature of XEPA's performances, as well as its anthropomorphized art-world status as an "emerging artist."

This paper explains the intent as well as operation of the XEPA hardware and software. Also discussed is how lessons learned from my personal experience with improvisational performance are related to psychological research in aesthetics, and how they have influenced XEPA's development. Some details regarding the use of color systems in XEPA and computer art are presented. Finally, directions for future XEPA development are discussed.

First, however, the state of the art that sets the context for XEPA requires some explanation.

The Current State of Affairs in Generative Art

In a now decade-old paper, I offered what is to date the most widely cited definition of generative art in the literature:

> Generative art refers to any art practice where the artist uses a system, such as a set of natural language rules, a computer program, a machine, or other procedural invention, which is set into motion with some degree of autonomy contributing to or resulting in a completed work of art [1].

The key element in generative art is the use of an external, possibly non-digital, system to which the artist cedes some control. Many, however, restrict their interest in generative art to computer-based systems.

In recent years the nascent field of complexity science has studied, compared, contrasted, and mathematically and computationally modeled various kinds of systems. An abstract

  Leonardo, Vol. 47, No. 4, pp. 386-393, 2014 doi:10.1162/LEON_a_00844

understanding of systems that spans the physical, biological, and social sciences is beginning to emerge. And it is these very systems that are being used as state-of-the-art generative systems by artists.

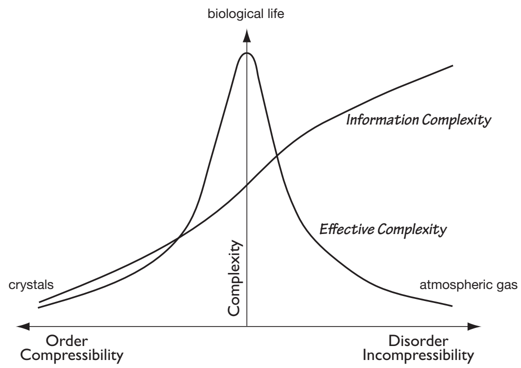

Figure 1. Effective complexity increases in systems that combine order and disorder. © 2012 Philip Galanter.

Complex systems defy simple description and easy prediction. Many would agree that the most complex systems we encounter are other living things. And life requires a mix of order and disorder. It was this kind of intuition that led physicists Murray Gell-Mann and Seth Lloyd to suggest the notion of effective complexity [2]. Unlike Shannon's information complexity that increases with disorder, effective complexity peaks where there is a mix of order and disorder (Figure 1).

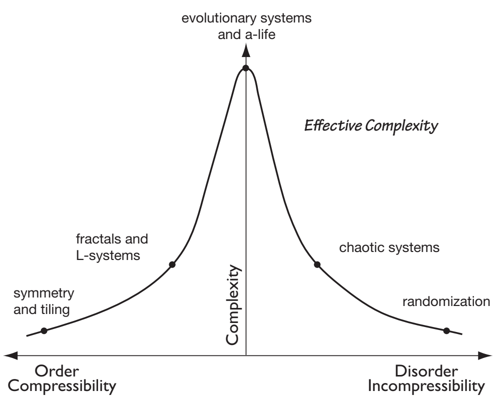

Figure 2. Generative systems organized by effective complexity. © 2012 Philip Galanter.

Over the years, generative artists have created art using systems such as genetic algorithms, reaction diffusion systems, cellular automata, artificial life, deterministic chaos, fractals, and Lindenmayer systems. The notion of effective complexity can be used to classify the various systems used in generative art. While these systems can offer a seemingly unending stream of visuals and sound, they typically do so without discrimination, and they lack any self-critical functionality (Figure 2).

Philosophers such as Boden, speaking of creativity, emphasize that novelty is a necessary but insufficient criteria for creativity. Creativity also carries with it the implication that the results are useful or otherwise of value [3]. To fully qualify as creative artists, computers will have to at least combine generative systems with computational aesthetic evaluation. However, computational aesthetic evaluation remains an unsolved problem.

Computational Aesthetic Evaluation

It is this wish to go beyond merely generative systems and invent creative self-critical systems that led me to create XEPA. My previous light sculptures, such as *RGBCA #1* and *RGBCA #2* (2010), had no such aspect.

Artists exercise critical aesthetic judgment in all phases of their work. Aesthetic evaluation comes into play when studying other artists, when applying micro-decisions while creating a piece, in learning from a newly created piece prior to beginning the next piece, and so on. It also comes into play when trying to categorize art as to genre or movement or specific intent.

Systematic attempts to measure aesthetics have, for the most part, been unsuccessful. I've written a chapter-sized overview of these attempts for those interested in greater detail [4]. All that can be offered here are a few examples.

One of the most well-known attempts was the mathematician George David Birkhoff's aesthetic measure. However, the way he operationalized his formula was fraught with difficulties and almost immediately disproved in empirical studies.

The Golden Ratio φ, approximately equal to 1.618, and the related Fibonacci series have been said to generate proportions of optimal aesthetic value. Writers such as Livio have arguably debunked this [5].

On somewhat firmer ground is a principle commonly referred to as Zipf's law. This is related to 1/f power laws that can be observed in statistical distributions of notes in melodies and colors in images. However it is at most a suggestive partial test [6, 7].

Machado and Cardoso have adapted Birkhoff's aesthetic measure as resistance to jpeg compression divided by resistance to fractal compression. While they reported good results with a specific kind of generative image, there is no evidence that this approach generalizes [8].

Lessons from Improvisation and the Psychology of Aesthetics

The XEPA algorithms have been heavily influenced by lessons learned from my personal experience as an improvisational musician and performance artist. These lessons correspond well with some of the leading psychological research regarding aesthetics.

Rudolf Arnheim applied the principles of gestalt psychology to aesthetic perception and established the notion of aesthetic perception as cognition. The law of pragnanz in gestalt states that the process of perceptual cognition endeavors to order experience into wholes that maximize clarity of structure. From this law come the notions of closure, proximity, containment, grouping, and so on, now taught as design principles. Unfortunately Arnheim's theory of aesthetics is much more descriptive than normative, and direct application to computational aesthetic evaluation is not obvious [9].

Daniel E. Berlyne's most significant contribution to the psychology of aesthetics is the concept of arousal potential and its relationship to hedonic response. Arousal potential is a property of stimulus patterns and a measure of the capability of a stimulus to arouse the nervous system. Using concepts from Shannon's information theory, he proposes that hedonic response is the result of separate and distinct reward and aversion systems [10].

Berlyne suggests that reward and aversion systems' activation increases as a Gaussian cumulative distribution. Because the reward system requires less arousal potential exposure to activate, it creates a response function called a Wundt curve. Stimulation initially increases pleasure, but as the aversion system kicks in, pleasure decreases and then crosses over to increasing levels of displeasure. In previous writing I've suggested that this response corresponds well to a cognitive system optimized for the perceptual processing of complex systems that exhibit an order/disorder curve similar to that shown in Figure 1 [11].

Colin Martindale developed a (natural) neural network model of aesthetic perception dynamics that he calls prototypicality. Martindale's neurons form nodes that accept, process, and pass on stimulation from lower to higher levels of cognition. Low levels tend to be ignored, and high-level semantic nodes encoding for meaning have the greatest strength in determining preference [12, 13].

Nodes are connected in an excitatory manner to nodes corresponding to superordinate categories. But nodes at the same level have a lateral inhibitory effect. So nodes encoding for similar stimuli will be physically closer together than unrelated nodes, thus creating semantic fields. The overall nervous system is optimally activated when presented an unambiguous stimulus matching a prototypically strong path up the neural hierarchy. Preference is given to stimuli typical of their class.

Lessons from my experience as an improvisational musician and performance artist correspond well with these psychological theories. This has led directly to the approaches encapsulated in the XEPA algorithms.

One lesson is that effective improvisation can be far from technical perfection because the audience will meet an improvised performance more than halfway. Per Arnheim, our gestalt mechanisms will "fill in" and structure our perceptions.

Another lesson is that the audience wants to be surprised, but not left behind, by a performance too unpredictable to follow. This is similar to Berlyne's concept of arousal potential and my suggestion that our perceptual processing is tuned for high effective complexity.

A third lesson is that micro-aesthetic decisions by themselves don't matter nearly as much as the contribution they make to a clear high-level semantic impression. This is similar to Martindale's notion of prototypicality, where low-level sensations result in successful aesthetics when they resonate with a unified high-level abstraction.

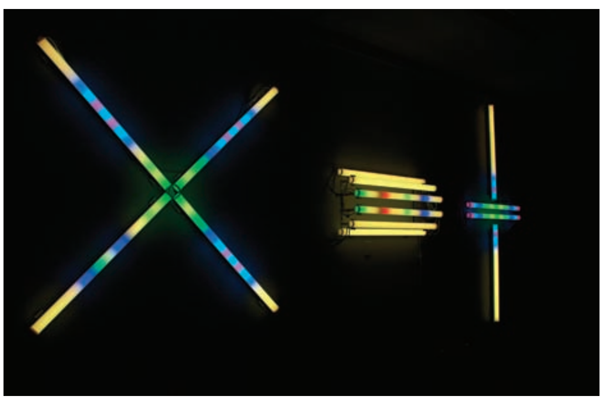

Figure 3. Three XEPAs in alpha testing during software development. © 2013 Philip Galanter.

XEPA as an Evolving Technology and Artwork

XEPA (Figure 3) is a platform for experiments in computational aesthetic evaluation. It should be kept in mind, however, that the project is fundamentally artistic in motivation. No pretense of controlled scientific research is implied. There is, however, an engineering aspect and an attempt to anticipate where these new engineering approaches may lead us aesthetically and culturally.

XEPA is intended to be reconfigurable and variable. Individual XEPAs are constructed using four to eight one-meter-long tubes. XEPAs can be wall-mounted, free-standing, or suspended sculptures. Different installations may have differing numbers of XEPAs of different designs. Each tube is a milky white diffuser with 16 RGB LED lighting elements inside acting as 16 pixels. Each pixel is individually addressable as a 24-bit color using the lighting industry DMX control protocol.

Each XEPA has a single studio-quality monitor (e.g. Genelec 1029A) with built-in amplification. A XEPA acts as a single instrumentalist, and together they create a band.

Each XEPA uses three inexpensive processors. An Arduino Mega 2560 is used for high-level observation and decision-making. The Mega 2560 is an open-source hardware platform using an ATmega2560 microcontroller chip with 256 KB of flash memory for code, 8 KB of SRAM for

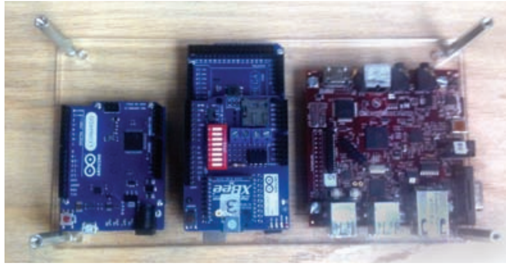

Figure 4. XEPA “Brain” without front acrylic cover and processor interconnects. © 2013 Philip Galanter.

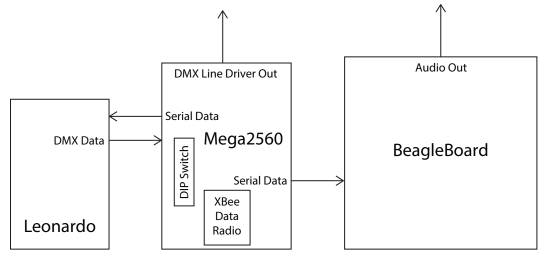

Figure 5. XEPA “Brain” interconnection design. © 2013 Philip Galanter.

variable memory, and 4 KB of EEPROM for non-volatile storage (Figure 4).

An Arduino Leonardo is used for real-time DMX communications used to control the LED tube animation. The Leonardo uses an ATmega32u4 microcontroller chip with 32 KB of flash memory for code, 2.5 KB of SRAM for variable memory, and 1 KB of EEPROM for non-volatile storage not requiring frequent updates.

The third processor is an open-source hardware single-board computer produced by Texas Instruments called the BeagleBoard. The BeagleBoard-xM used by each XEPA uses a TI DM3730 processor running at 1 GHz with an ARM Cortex-A8 core. The BeagleBoard has 512 MB of RAM for both code and data, and boots from a 4 GB microSD memory card. The BeagleBoard is designed to be a complete single-board computer. XEPA uses the BeagleBoard as a sound engine for real-time high-fidelity sound synthesis (Figure 5).

All three boards are mounted on laser-cut clear sheet acrylic enclosures that can stand freely or be wall-mounted. The enclosures are clear in order to present the XEPA “Brain” as a deconstructed, demystified element.

Each XEPA sends XBee data radio packets at the start of a phrase describing what it is doing. This includes things such as base hue, color harmony scheme, pattern type, tempo, key, scale, and so on. There is no score or conductor. No XEPA “tells” another what to do. This use of data radio makes it seem, in principle, as if the XEPAs simply watch each other and only modify their own behavior.

This is meant to suggest that future machine societies will interoperate by robust indirect observation rather than brittle low-level protocols.

My design includes an 8-bit DIP switch used to assign unique ID numbers and set debug modes. There is also a small line driver circuit used to convert the +5 volt data from the Leonardo to a balanced DMX signal. Not shown is a microSD memory card reader intended for future storage of larger data tables.

The Mega 2560 watches the other XEPAs, executes aesthetic evaluation, and decides which light animation and sound phrases will be performed. At regular intervals the Mega 2560 sends short commands to the Leonardo and BeagleBoard. The Leonardo then executes an animation phrase, and the BeagleBoard generates a sound phrase.

Current Algorithms

XEPA is designed to execute effective unique improvisations based on correspondences provided by the artist as probability tables.

A hierarchical model inspired by Martindale is used to gain leverage over what could be a combinatorial explosion of hues, color schemes, animation patterns, scales, keys, note patterns, and so on. Themes are at the highest semantic level, and each is a suggestive phrase such as "arctic zone" or "house on fire" or "spring life." For each theme, each color palette, scale, animation sequence, and so on is given a weight based on artistic intuition. For example, a palette of blues and whites would be given a large weight for the theme "arctic zone," while a palette of reds and yellows would be given a low weight. While this is a combinatorial burden, it is not at all impossible for 20 or so themes.

Because the XEPAs begin in random states and explore a complex media space, having only 20 themes does not mean that there will only be 20 kinds of performances. The current state of XEPA uses a quasi-genetic approach. The probability tables create the basis for a kind of fuzzy logic for theme membership given a genotype. And it is theme membership, akin to Martindale's prototypicality, that acts as a fitness function.

In performance, each XEPA independently executes table-driven computational aesthetic evaluation of the other XEPAs, and then adapts its own performance.

- Whenever a new packet is received from another XEPA
  - Time-stamp the packet for possible later synchronization
  - Compare the packet (genotype) to the weights for each theme generating an error score (fitness score) for each
- At the end of a phrase, compare your error score to the error scores of the other XEPAs
  - If there is a XEPA with a lower error score, use a Monte Carlo technique to copy part of its genotype, i.e. use crossover, and synchronize with that XEPA
  - Otherwise see if a random change will produce an even lower error score, i.e. use mutation

XEPAs initialized in random states will execute this quasi-evolutionary system in a loosely coupled manner. Over time the performing XEPAs will converge on a coherent theme.

The XEPA Painterly Color System

In my previous light sculptures, and indeed in the work of others, I've noted a general dominance of cool colors and a lack of strong yellows and oranges. This is primarily due to the use of the additive RGB color system and the resultant spacing of colors around the color wheel. The typical RYB subtractive system used by painters spreads the warm colors further around the color wheel (Figure 6). Using the HSV color mode in Adobe Photoshop, or in Processing or other digital color applications, will yield the cool color-dominant spacing seen in the RGB color wheel. The result is that uniformly random hues, or color harmonies based on equal color wheel spacing, yields cool color-dominated palettes.

In order to encourage a balance of warm and cool colors and painterly color palettes, XEPA does all calculations using a modified RYB color system I've called RYB Plus. The RYB Plus system smooths the RYB transition points from cool to warm colors (Figure 7).

Future Directions

One can think of each XEPA as the embodiment of a large multi-dimensional media space. Each theme, such as "arctic zone" or "house on fire," and its corresponding probabilities create an aesthetic attractor in that space. Computational aesthetic evaluation for a given XEPA consists of measuring its distance from each aesthetic attractor. Those that are close to a given

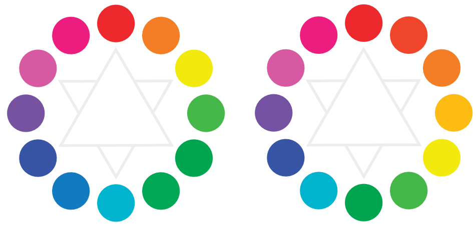

Figure 6. The additive RGB system (left) versus the subtractive RYB system (right). © 2013 Philip Galanter.

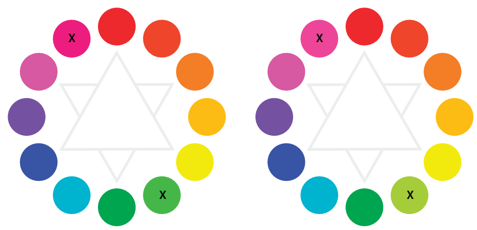

Figure 7. The RYB system (left) versus the RYB Plus system (right). © 2013 Philip Galanter.

attractor try to self-modify to get even closer. Those that are far from all the aesthetic attractors try to get closer to those XEPAs that are in turn closer to given aesthetic attractors.

From a quasi-genetic point of view, these distances from aesthetic attractors in media space are akin to fitness function results. Self-modification on the part of each XEPA to minimize distance is akin to genetic mutation. Self-modification of a "faraway" XEPA by copying genes from a "closer" XEPA is akin to genetic crossover.

A problem one immediately discovers with this approach is that some aesthetic attractors can be overly strong. In the degenerate case, all random initializations result in performances that lead to the same theme every time. To help prevent this, the probability tables are maintained as spreadsheets that translate artist-assigned weights to normalized values.

In a future variation the probability tables will be modified at the end of every performance to relatively weaken the theme/attractor that dominated. In a mild form, this will help further dampen overly strong aesthetic attractors.

In a stronger form, dynamic probability tables could result in the XEPAs as a group "tiring" or becoming "bored" with a given theme and turning their attention to others. This is akin to what some have termed "curious agents" [14]. It might also lead to the development of trends and styles as suggested by Martindale [15].

Another direction would be to add layers of semantic abstraction between the themes and the very specific probability tables. For example, a given theme might be tagged as requiring cool colors, and cool colors would then be probabilistically defined. Going further, some tags could then be defined by other tags through the construction of a semantic network.

Finally, rather than manually constructing themes, it might be possible to create themes based on the statistical analysis of a database of preexisting artworks. For example, color palette probabilities could result from the image analysis of a database of paintings.

Conclusion

Contemporary generative art practice encompasses a wide array of systems for the creation of visual and auditory form. What is lacking, and what seems likely to be the next frontier, is a computational self-critical function that provides feedback and guidance towards increasingly satisfying aesthetic results.

XEPA light and sound sculptures are autonomous units that observe each other and adaptively modify their behavior to converge on an aesthetically coherent real-time group performance. As such, XEPA is not only an art installation that exhibits emergent behavior, it is also a platform for experiments in computational aesthetic evaluation.

My subjective experience as an improvisational performer has been shown to correspond in several ways to specific psychological theories of aesthetic perception. These correspondences have strongly influenced the hardware, software, and artistic design of XEPA. This includes the three-processor XEPA "brain" and the algorithms used in the first generation of software. In particular, a quasi-genetic approach to variation is used to converge on aesthetic attractors in the media-space. In addition, XEPA uses a modified RYB color system to overcome some of the disadvantages exhibited by the typical use of RGB color in generative art.

Finally, because XEPA is an experimental platform rather than a singular artwork with a predefined end-design, there are a number of possible future directions. These might include algorithms for curiosity and boredom, semantic tagging for hierarchical abstraction, and the use of real-world data to establish aesthetic attractors.

Acknowledgement

The XEPA project is possible in part thanks to an Arts Research Enhancement Grant from The Academy for the Visual and Performing Arts at Texas A&M University.

References

1. Galanter, Philip, "What is Generative Art? Complexity Theory as a Context for Art Theory," *International Conference on Generative Art* (Milan: Generative Design Lab, Milan Polytechnic University, 2003).
2. Gell-Mann, Murray, and Seth Lloyd, "Information Measures, Effective Complexity, and Total Information," *Complexity*, Vol. 2, No. 1, 44–52 (1996).
3. Boden, M.A., *The Creative Mind: Myths and Mechanisms*, 2nd ed. (New York: Routledge, 2004), xiii.
4. Galanter, Philip, "Computational Aesthetic Evaluation: Past and Future," *Computers and Creativity*, ed. J. McCormack and M. D'Inverno (Berlin: Springer, 2012).
5. Livio, Mario, *The Golden Ratio: The Story of Phi, the World's Most Astonishing Number*, 1st trade paperback ed. (New York: Broadway Books, 2003), viii.
6. Zipf, G.K., *Human Behavior and the Principle of Least Effort: An Introduction to Human Ecology* (Cambridge, MA: Addison-Wesley Press, 1949), xi.
7. Manaris, Bill, et al., "Developing Fitness Functions for Pleasant Music: Zipf's Law and Interactive Evolution Systems," *Applications of Evolutionary Computing Proceedings* Vol. 3449, 498–507 (2005).
8. Machado, Penousal, and Amílcar Cardoso, "All the Truth About NevAr," *Applied Intelligence* Vol. 16, No. 2, 101–118 (2002).
9. Arnheim, Rudolf, *Art and Visual Perception: A Psychology of the Creative Eye*, new, expanded, and revised ed. (Berkeley: University of California Press, 1974), x.
10. Berlyne, D.E., *Aesthetics and Psychobiology* (New York: Appleton-Century-Crofts, 1971), xiv.
11. Galanter, Philip, "What is Generative Art? Complexity Theory as a Context for Art Theory," *International Conference on Generative Art* (Milan: Generative Design Lab, Milan Polytechnic University, 2003).
12. Martindale, Colin, "Aesthetics, Psychobiology, and Cognition," *The Foundations of Aesthetics, Art, and Art Education*, ed. Frank Farley and Ronald Neperud (New York: Praeger Publishers, 1988), 7–42.
13. Martindale, Colin, et al., "The Effect of Extraneous Stimulation on Aesthetic Preference," *Empirical Studies of the Arts* Vol. 23, No. 2, 83–91 (2005).
14. Saunders, Rob, and John S. Gero, "Curious Agents and Situated Design Evaluations," *AIEDAM: Artificial Intelligence for Engineering Design Analysis and Manufacturing* Vol. 18, No. 2, 153–161 (2004).
15. Martindale, Colin, *The Clockwork Muse: The Predictability of Artistic Change* (New York: Basic Books, 1990), xiv.