


# INTELLIGENT PARADIGMS



A.I.,  
Music,  
and a Video Game

by Jay Tobin

# SOME WORDS

Artificial Intelligence occupies a growing niche in the immense field of music. Common applications include automated mastering plugins capable of saving small-time musicians entire fortunes, or algorithmic rhythm generators trained on historically significant sound samples. Not as common is the role of AI-as-composer, and for a sensible reason: programming an AI is an exhausting process requiring intense specialization, a powerful computer, and a stack of cash to burn on sky-high electric bills.

This booklet exists as an explainer for my novel attempt at integrating artificial intelligence into the greater compositional process of creating music. Built using Unity and Max/MSP, this project takes on the challenge of sonifying the processes of AI as it moves through a 3D environment. While there's a ton to unpack here, this booklet will do its best to lead you through the process in as non-jargony and picture-intensive a way as I could manage.

# CONTENTS

I.	The A.I. and Navigation	4
II.	The Biomes	6
III.	The Sound	16
IV.	The Possible	24

# THE A.I.

## Navigation and Training

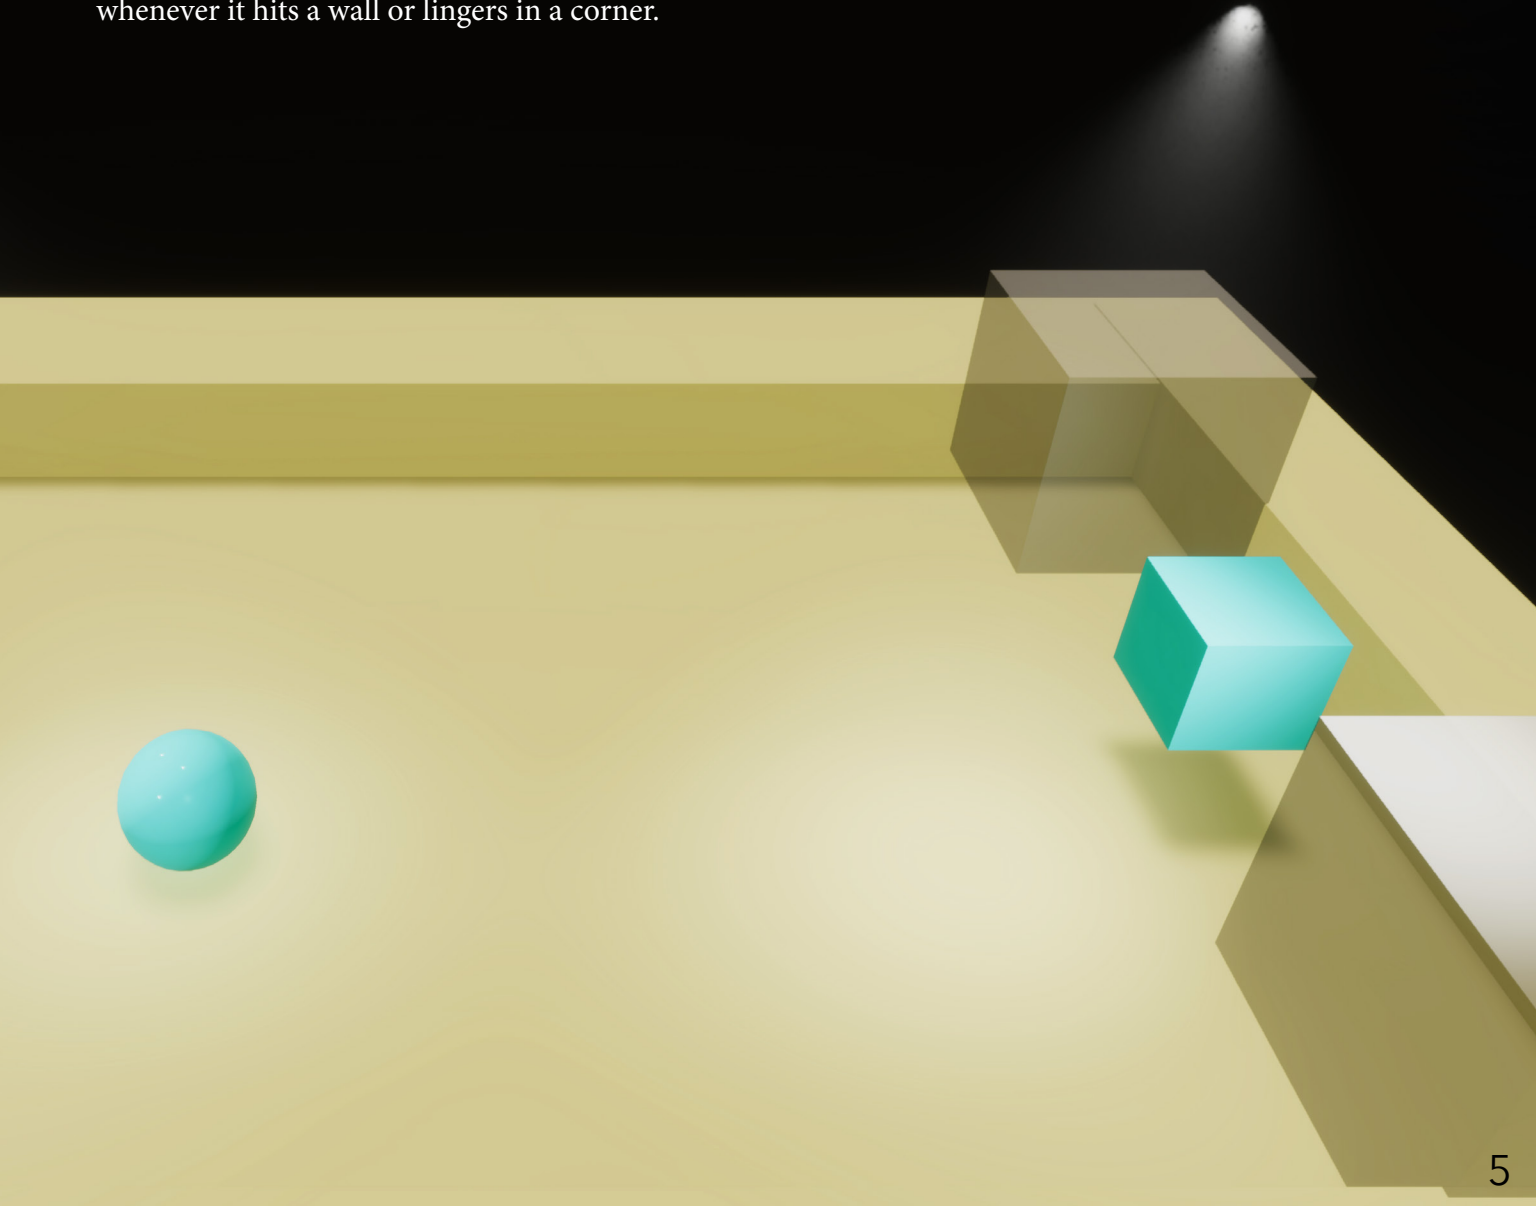
Artificial Intelligence means many things. In this case, I'm referring to an algorithm trained on a set of test data from which it can make its own conclusions without human supervision.

I used reinforcement learning (RL), a process through which a trainer leads an AI through the completion of a certain process and sets rewards for doing so quickly and without fault. In this use case, I led an AI, controlling a rock able to move on its own, towards an objective, penalizing it for unnecessary steps and rewarding it for getting to its destination quickly.

The result is a rock that's able to propel itself through 3D space to get to its destination entirely on its own.

As it moves through the world I've sculpted in Unity, it passes through individual Biomes, each with distinct sonic profiles. The movement of the AI through the biomes is what produces the bulk of the score you hear.

The AI was trained in a simplified mock-up of the field it would encounter. This image of the work in its first phase shows the AI, controlling the aquamarine ball, rolling towards its target in the form of a blue square. The AI applies force in the direction of its goal, but is penalized whenever it hits a wall or lingers in a corner.



# THE BIOMES

## Central Actors and Sonic Profiles


The simulation takes place on a single map at a scale of about 10 square miles (relative to the main actor). Within the map are four individual biomes, each with their own sonic profile, filled with unique, individual objects that all produce their own distinct sounds.




DESERT




MOUNTAINS



RUNOFF



BOGS



# DESERT

The image features a stylized desert landscape. In the foreground, there are rolling, dark brown dunes. In the middle ground, several rows of red, spiky, cactus-like flora are silhouetted against a warm, orange, hazy sky. The overall color palette is dominated by warm tones of orange, brown, and red. The word "DESERT" is written in a large, white, serif font across the upper portion of the image.

Rows of dunes and rugged red flora sprawl across the map, spilling forth from the heart of the Desert. All things begin here.





**SONIC PROFILE:**

Grainy, Dry

**UNIQUE ACTORS:**

Old Pillars, Wandering Pebbles

# MOUNTAINS

A landscape of rolling sand dunes under a soft, hazy sky. The dunes are illuminated from the side, creating a gradient of light and shadow. The sky is a pale, uniform color, suggesting a clear or slightly overcast day. The overall mood is serene and quiet.

Home to many rocks, three gods rule the peaks,  
carefully watching over their realm. Their collective  
song rings off the rocky mountainsides.

A hand is shown holding three gemstones. The largest gemstone is a bright, glowing yellow-orange, held in the palm. A smaller, dark grey gemstone is held in the index finger. A small, reddish-orange gemstone is held in the thumb. The background is dark and blurry.

**SONIC PROFILE:**


Bright, Glistening

**UNIQUE ACTORS:**

Three Gods, Wandering Pebbles

# RUNOFF


A pool of decay where old gods lay and many rocks find an oily end. Those still able to sing lament their fate in collective elegy. Few who enter ever leave.





**SONIC PROFILE:**  
Low, Foreboding

**UNIQUE ACTORS:**  
Dead Rocks



# BOGS

Wetlands teeming with life, the Bogs are beacons of nature with a thrumming soul at their center.



## SONIC PROFILE:

Dense, Mystical

## UNIQUE ACTORS:

Living Tree



# THE SOUND

Sonification, Synthesis, Actors and Space



This piece is foremost an exercise in algorithmic music making: while a robust set of visuals detail the exploits of the central AI moving through its world, the behind-the-scenes sound engine is what constitutes the backbone of the piece.

The AI, the map, and its actors all work in concert to produce a flowing work of music that changes as time progresses and as the AI itself throws itself through its world. The music itself consists of a single algorithmic synthesizer, which creates a bulk of the melodic content, and dozens of smaller subsystems attached to the 3D actors within each biome. In addition, all actors are spatialized in real-space relative to their location on screen, facilitated by an 8-speaker surround-sound system.





# SOUND IN SPACE

The music of the piece changes as the central AI rolls around its map in pursuit of its endless goal. To do so, each biome is assigned a specific sonic profile, complete with a distinct set of chords and a special set of identifying factors that make each biome sonically distinguishable.

The sonic profiles correspond to space within the Unity world, meaning that the profiles can smoothly blend into one another wherever they overlap. The diagram to the right is the actual system used to interpolate between sonic profiles: the yellow dot corresponds to the AI's location within the Unity world, while the colored circles illustrate the range of each sonic profile and where they smoothly coalesce.



Each of the presets, numbered and in their corners, with the rock in the center.




The spherical interface of the spatialization engine.

The piece's audio spatialization system is a multi-channel system, modular in construction, made to immerse listeners in a full field of audio. With 8 operating speakers, listeners can hear the AI moving through space and interacting with the many actors it encounters.

What's seen on the screen is what's heard in the installation: if an actor is at the far left corner of the projection, so too will its sound occupy the far left corner of the installation room.

# THE SYNTHESIS



This piece makes extensive use of a sound synthesis technique called granular synthesis. Granular synthesis uses tiny blips of audio, often between a few milliseconds to half-a-second in length, and plays them back rapidly. The effect is that of a cloud made of millions of grains, able to flow amorously between sonic shapes.

In this particular instance of granular synthesis, I use a 2 minute long sample made specifically for the piece. The combination of a tailor-made sample and an extensive granular synthesis engine combine to create dense musical textures that never repeat, no matter how long the piece has been running.

# PROBABILITY AND CHANCE

current biome: 0

ruleset:  
1 1 0, 1 2 1, 1 3 1, 1 4 0,  
2 1 1, 2 2 0, 2 3 1, 2 4 1,  
3 1 1, 3 2 1, 3 3 0, 3 4 1,  
4 1 1, 4 2 0, 4 3 0, 4 4 0,

determined  
progression:



The actual chord system, with determining rules, set to run calculations for the desert biome (labeled '0').

Granular synthesis is nothing without an interesting set of musical conditions to operate within. The synthesis engine in this piece is driven by a massive web of chance operations controlled by probability and a bit of statistics.

Within each individual biome, there are 4 possible chords that can be sonified by the synthesis engine: 3 of them are continuous and harmonious, and the last one discordant with a much lower chance of playing. This systems insures that the chord progression remains as fresh as the sonic textures produced by the granular synthesizer, and adds to the aural character of each biome.

66 69 73  F#m

65 69 74  A#m

62 65 69  Dm

61 65 68  C# maj

The chord row for the Bogs.

# THE ACTORS


The background of the page is a surreal, monochromatic landscape. The sky is a gradient of dark brown and orange. In the distance, there are dark, silhouetted mountains. On the right side, a large, abstract tree with geometric, crystalline branches stands on a small hill. In the foreground, there is a body of water with a shimmering, iridescent surface that reflects the colors of the sky and the landscape.

The work makes use of its spatialization system through the implementation of unique actors, some specific to individual biomes, and other roaming between them. The sounds produced by the actors are not part of the granular synthesis engine - instead, each sound is specifically generated for the individual actor.

The generated sounds follow a web of chance-based progressions to maintain melodic cohesion between each actor and the granular synthesizer at the piece's center. Most actors produce some form of synthesized sound distinct in its digital nature, save for one with a human choir at its soul.

# The Pillars

Mysterious, towering entities that hum along in deep bassy voices. All four of them occupy the desert, entirely unmoving, casting their shadows across the rest of the world.




# The Old Ones

Deposed despots that once ruled from the mountains, the Old Ones all sit in the deep muck of the Runoff, unable to do anything other than sing their own elegy. How they rolled into the inescapable mire is lost to history.

# The New Ones

## and their Wanderers

The current leaders of rock-kind. Three remain supreme on their mountain, each clad in gold, silver, and copper. Their subjects, the Wanderers, roam the world below them and occasionally climb the mountain to pay their dues.




# The Living Tree

A glowing, crystalline tree with harmonies known to overwhelm. Some say its voices change their song depending on who they sing to.

# THE AI AND THE MUSIC

## A Careful Dance between Algorithms



The relationship between the AI at the heart of the piece and the synthesis engines that make up its musical content is all in the algorithm. The AI, acting on its own to meet its objective as it best sees fit, shifts weights and throws switches within the chance-based systems that make up the composition.

As the AI moves and its algorithm makes its calculations, so too does the extensive web of musical probabilities, controlling synthesizer and actor alike, react in real-time. While the AI might not be the orchestra producing the music, its the conductor at its heart, waving its baton by moving through its world, guiding the music.

# THE POSSIBLE

## Potentials and the Future


While this piece is extensive in its implementation of AI and generative systems, its only a brief peak into the endless possibilities provided by artificially intelligent music making processes.

The novel approach taken by this piece employs an AI entirely devoid of musical knowledge - that is, untrained in musicality and only possessing enough knowledge to move itself through its found environment. While this was intentional to encourage interplay between composer and machine, there exist possibilities for infinitely more complex AIs to operate in significantly more thorough frameworks, if not build their own.

The limitations of supervised Reinforcement Learning give this particular piece its naive character, with the simplicity of a rock rolling through space sonified by a dense probability-based system that thrives from a sense of randomness. AI can be used to do significantly more in significantly more sophisticated ways, but the use of AI as an actor within an artistic framework, particularly as a collaborator, is a field that remains burgeoning yet sparse.

This early iteration of the piece, simple in visuals and execution, acts as a catalyst for further works into developing creative AI systems, inviting conversation and experimentation.





The next steps for this project rely on interaction with the audience: immersive VR is a very real possibility, as is an interactive video gaming experience that likewise imbues players with a sense of compositional autonomy. The structures built to support this work are entirely modular and reusable: the probability-based chord progressions can be used ad-infinitum with very little programming knowledge, and the many synthesis engines will no doubt be employed in my future work.

While I'm reluctant to call this work future-proof, it's absolutely built with the future in mind. The systems at work here, and the artificial intelligence at the hub of the piece, is an attempt to see what the future may hold for human-machine interaction. Should the future create experiences not-too-dissimilar from this one, there could be AI behaviors within the work of interests to artists and scientists alike.

If not, at least I'll have made something I like.

...thank you.

