Towards a User's Manual for Building Worlds

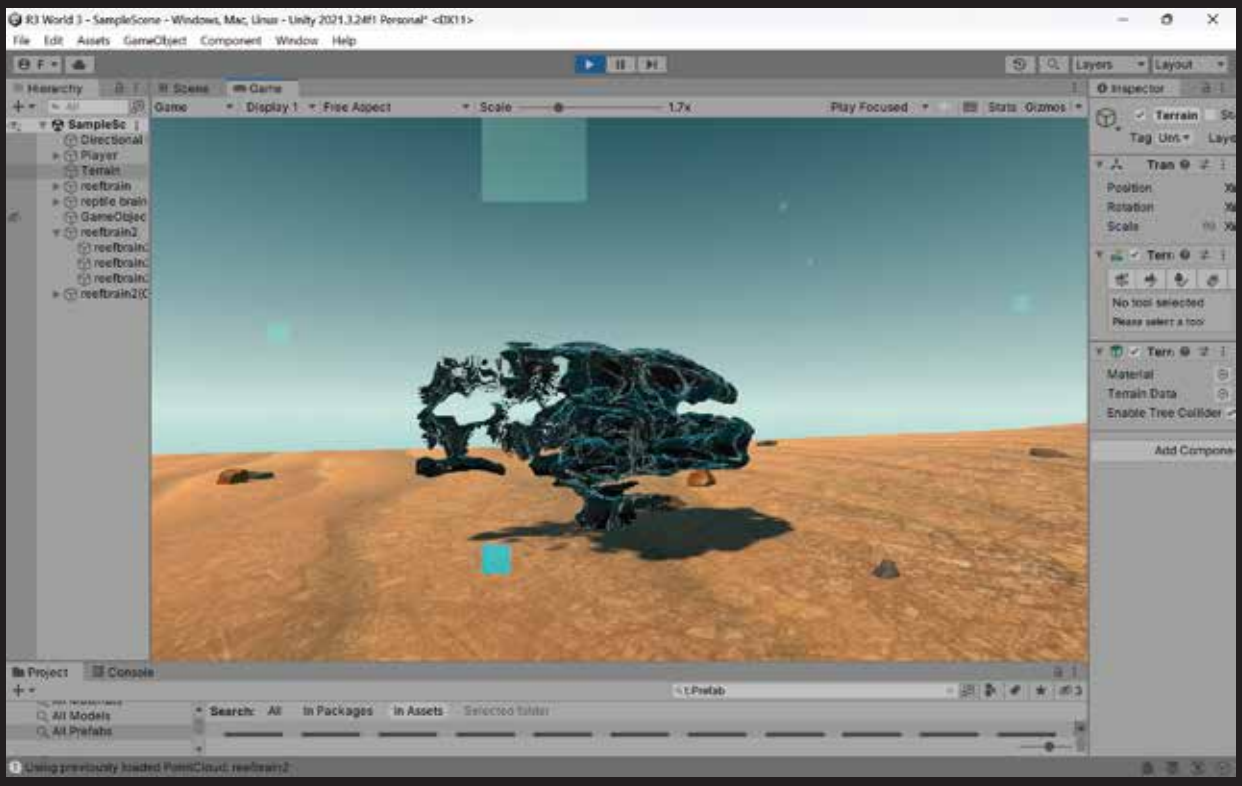

TRC x Common Sense Machines

Proposal Brief:

Common Sense Machines software offerings provide their organization with many different ways in which to explore the rupture of how machines can understand and build worlds. TRC proposes an experiment that utilizes CSM's AI and machine learning technologies applied to 3d worlds, while also integrating approaches from conceptual art in an attempt to conceive of building experiments that will help discover how non-human intelligences can understand and build worlds. The methodology that TRC proposes is ***Constructing Conceptual Landscapes***, an approach that seeks to experiment with how machines can create landscapes that produce new forms of thought.

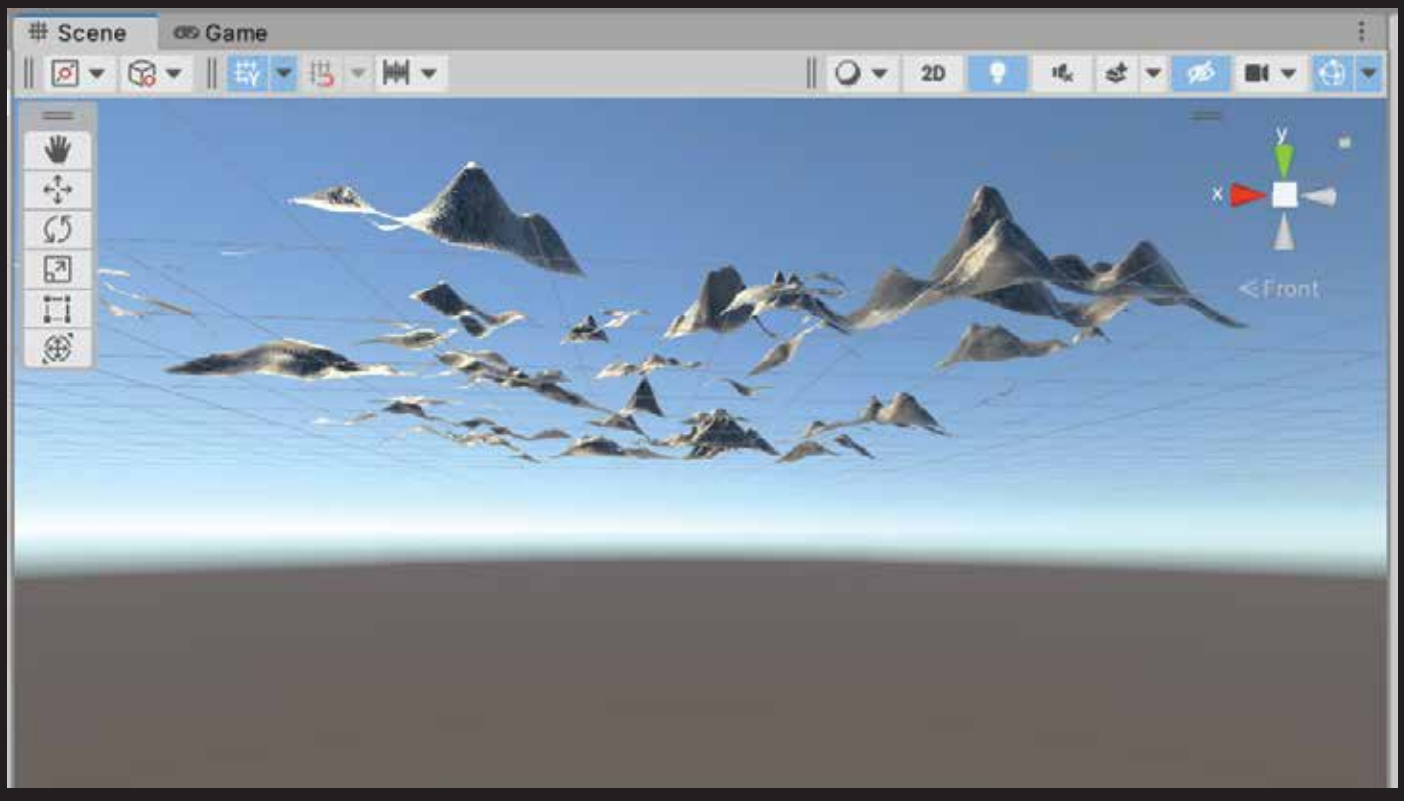

Conceptual Landscapes:

TRC proposes approaching CSM software as a machine to create landscapes that produce new forms of thought:

"Art does not imitate nature, it creates a world apart" – Paul Klee

"Poetry tries not to tame the forms which form language, not to procure the inscription which retains the event of the landscape. It tries to slip by before its withdrawal." – Jean-Francis Lyotard

Simulatations vs. Representations:

A simulation is a computer-based model that attempts to replicate the behavior of a real-world system or phenomenon. Simulations can be used to study complex systems that are difficult to observe or manipulate in the physical world, or to predict the behavior of a system under different conditions. For example, simulations can be used to model weather patterns, the spread of diseases, or the performance of a new product design.

On the other hand, a representation is a visual or conceptual depiction of an object, idea, or system. Representations can be static or dynamic, and can take many forms, such as diagrams, drawings, photographs, or 3D models. Representations can be used to convey information, to illustrate concepts, or to visualize data.

NeRFs vs. LiDAR:

NeRFs are a type of deep learning algorithm that can generate highly detailed 3D models of objects and scenes from 2D images or videos. By training a neural network on a large dataset of images, NeRFs can learn to predict the appearance of an object or scene from any viewpoint, even if that viewpoint has not been captured by a camera.

LiDAR is a remote sensing technology that uses lasers to measure distances and generate a 3D map of the environment. It works by emitting a laser beam and measuring the time it takes for the beam to bounce back off an object. This data is then used to create a 3D point cloud, which can be used for various applications such as autonomous driving, surveying, and mapping.

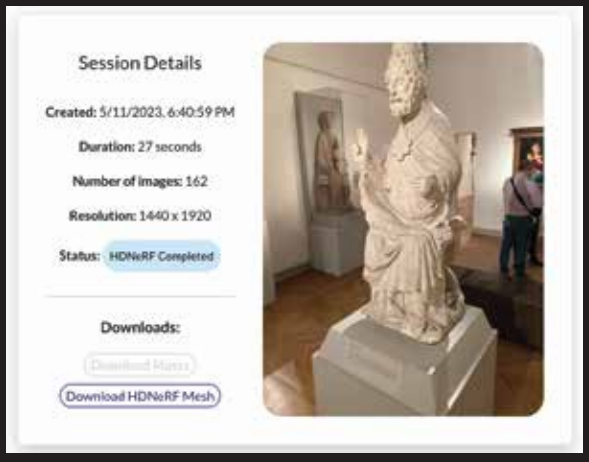

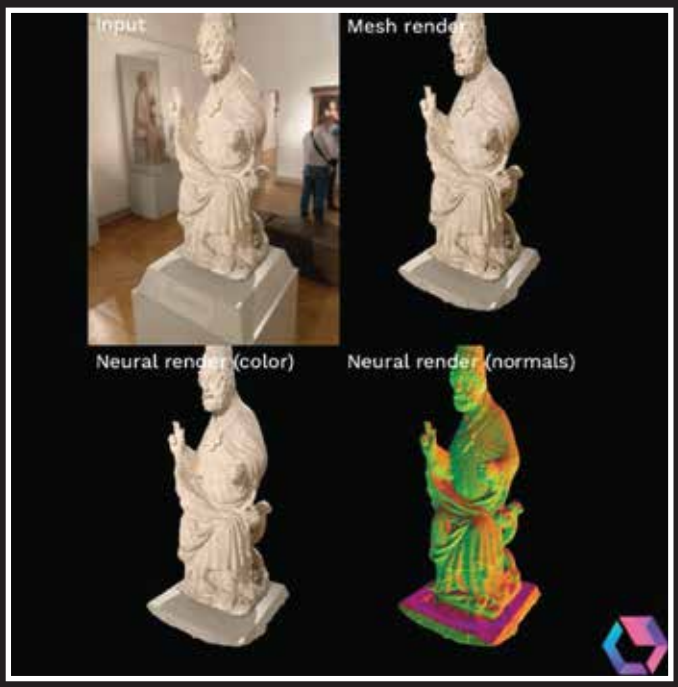

Artificial Imagination:

"... we still need to clarify the role of "artificial imagination" in machines since machines are no longer just a passive medium of support but also actively participate in the cognitive process of imagination. However, this is not to say that artifacts before the computational age had no capacity to constitute the imagination. Art would, of course, be the counterexample since the question of imagination is fundamental in artistic, that is, artificial creation and expression." – Yuk Hui

Can games be a catalyst to create new forms of thought? Through the use of CSM's technological offerings to build new digital worlds that are simulations of [new] realities, as opposed to representations of pre-existing cartesian spaces, we can utilize games to explore new landscapes of thought.

How can we build new types of games in these new 3-D simulated worlds that are constructed? What historical points of reference can we look at for inspiration for these games, such as Conceptual Art works that utilized writing to simulate writing that builds worlds such as Dan Graham's Schemas or Sol Lewitt's Algorithmic Outputs?

Conceptual Art & Games:

"The symbolic games of which we take part do not represent any universe of concrete experience, but on the contrary, this concrete experience represents games. **We live our concrete experience in function of games. Games are our ontological ground and all future ontology is necessarily game theory. Everything is fiction, nothing is real."**

"In face of every program, the question that emerges is **not: "how real is it?" but "how does it function?"** *The question of reality and falsity becomes a question in relation to the produced effect.* **What interests us in programs is not the input but the output."**

– Our Game, Vilem Flusser

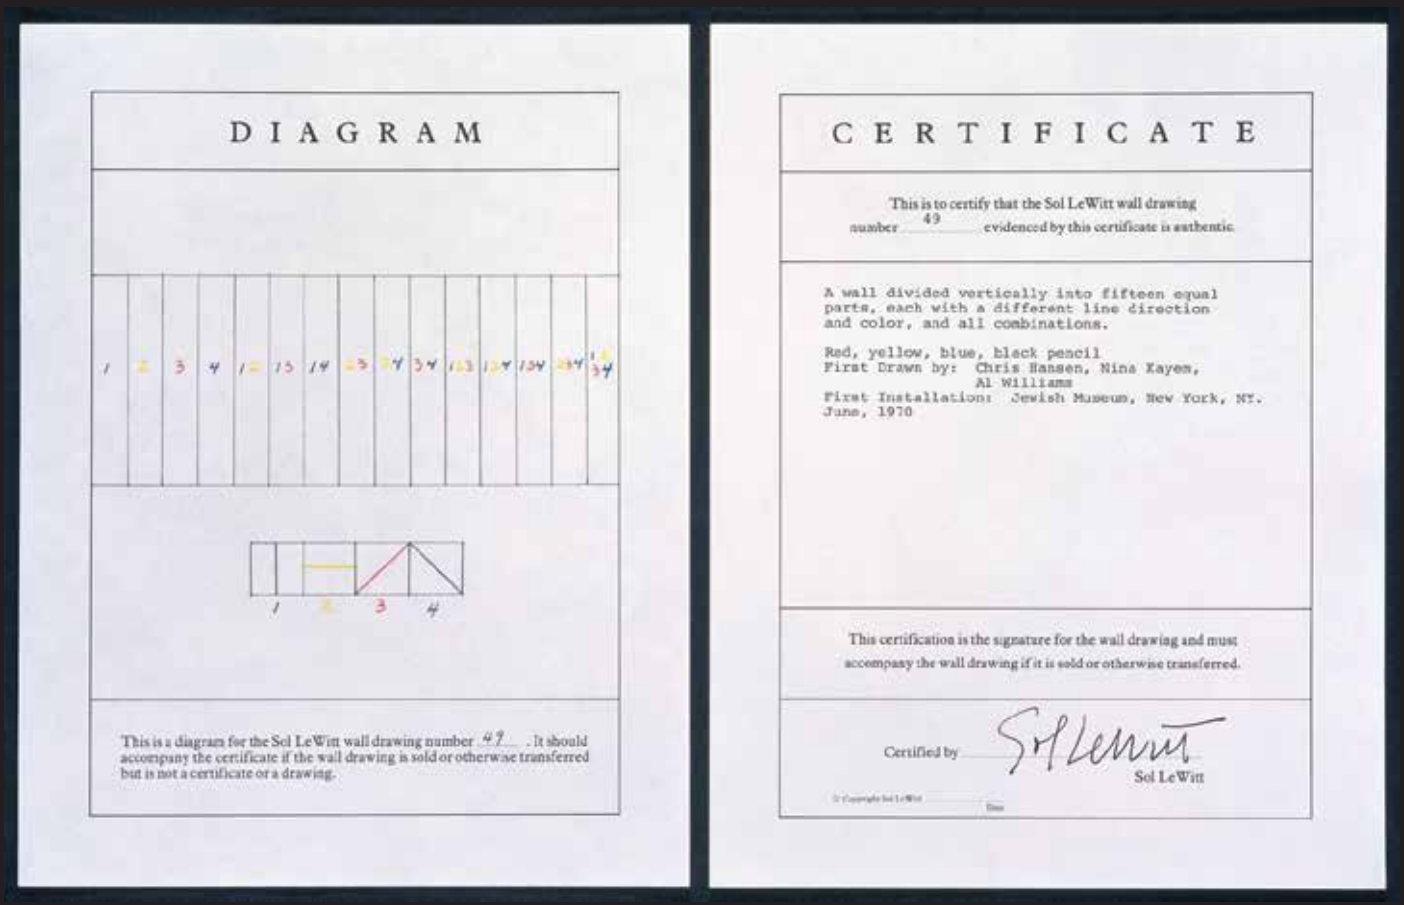

A Wall Divided Vertically into Fifteen Equal Parts, Each with a Different Line Direction and Colour, and All Combinations. (1970). Sol LeWitt

Conceptualism emphasizes ideas and concepts over traditional aesthetic concerns, which often involves the use of language, text, and architectural forms to challenge viewers' perceptions and expectations of art.

Dan Graham's conceptualism is rooted in the idea that art should be a means of social and cultural critique, and that it should engage with the world outside the gallery or museum. His work often explores the relationship between art, architecture, and urban space, and seeks to highlight the ways in which these elements shape our experiences of the world around us.

Conceptual art can be seen as a kind of proto-protocol. Protocols can be thought of as the rules and procedures that govern how we create and interact with artworks. By designing and implementing new protocols, we can explore new modes of collaboration, participation, and interaction in the art world and how we can build new protocols and algorithms that enable more collaborative and participatory modes of interaction between artists and machines.

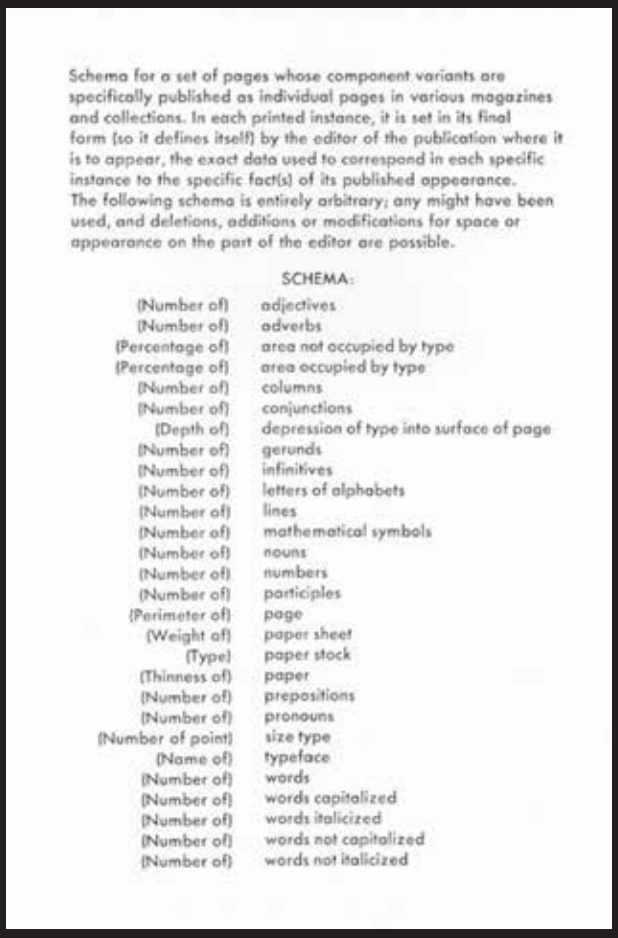

Schema for a set of pages whose component variants are specifically published as individual pages in various magazines and collections. In each printed instance, it is set in its final form (so it defines itself) by the editor of the publication where it is to appear, the exact data used to correspond in each specific instance to the specific fact(s) of its published appearance. The following schema is entirely arbitrary; any might have been used, and deletions, additions or modifications for space or appearance on the part of the editor are possible.

SCHEMA:

(Number of)	adjectives
(Number of)	adverbs
(Percentage of)	area not occupied by type
(Percentage of)	area occupied by type
(Number of)	columns
(Number of)	conjunctions
(Depth of)	depression of type into surface of page
(Number of)	gerunds
(Number of)	infinitives
(Number of)	letters of alphabets
(Number of)	lines
(Number of)	mathematical symbols
(Number of)	nouns
(Number of)	numbers
(Number of)	participles
(Perimeter of)	page
(Weight of)	paper sheet
(Type)	paper stock
(Thinness of)	paper
(Number of)	prepositions
(Number of)	pronouns
(Number of point)	size type
(Name of)	typeface
(Number of)	words
(Number of)	words capitalized
(Number of)	words italicized
(Number of)	words not capitalized
(Number of)	words not italicized

GRAHAM

Offsetplatte Schpapier

1 Adjektiv
2 Adverbien
1191 1/2 qm unbedruckte Fläche
339 1/2 qm bedruckte Fläche
1 Satzspalte
0 Konjunktionen
beim Eindrucken der Drucktype in die Papieroberfläche
0 Gerundien
0 Infinitive
363 Buchstaben aus dem Alphabet
37 Zeilen
2 mathematische Symbole
38 Hauptwörter
52 Zahlen
0 Partizipien
8 1/2 x 5 Satzspiegel
17 1/2 x 22 1/2 Papierformat
Offsetplatte Schpapier
3 Präpositionen
0 Pronomen
10 Punkt Schriftgrad
Antiqua Schriftart
59 Wörter
1 Wörter versal
0 Wörter kursiv
57 Wörter nicht versal
58 Wörter nicht kursiv

Schemas by Dan Graham

On World Models, User Interfaces, and (Re-)Presentations of Reality:

Any technology that seeks to represent or simulate reality must ultimately confront the question of what reality is. We must be careful not to conflate our representations with reality itself, or to assume that machines can fully comprehend the complexity and contingency of the world.

One approach is to see this technology as a form of philosophical inquiry, rather than simply a means to an end. We can use it to explore and interrogate our assumptions about reality and representation, and to push the boundaries of what we can conceive and articulate.

The interface is not simply a tool for accessing the technology, but is itself a form of mediation that shapes our experience and understanding of the world.

By combining the technological and philosophical dimensions of this project, we can create a truly innovative and transformative tool for exploring and understanding the world.

Proposal of Work:

What TRC proposes is the creation of a schema that functions not only in the world of text, but also in the world and language of CSM's software. What this would look like is open to interpretation from engineers, but on the text-based framework would appear as such (before then being rendered into 3d spaces):

Schema:

(number of) planes
(number of) objects
(number of) agents
(number of) NeRFs
(number of) lights
(number of) cameras
[etc.]

This schema would then need to be translated into code by CSM's engineering team team and then would be represented differently when turned into 3-dimensional space, in turn creating a new world, reality, and (re-)presentation of the schema, no longer in text form, but to be seen in the language of 3d space.

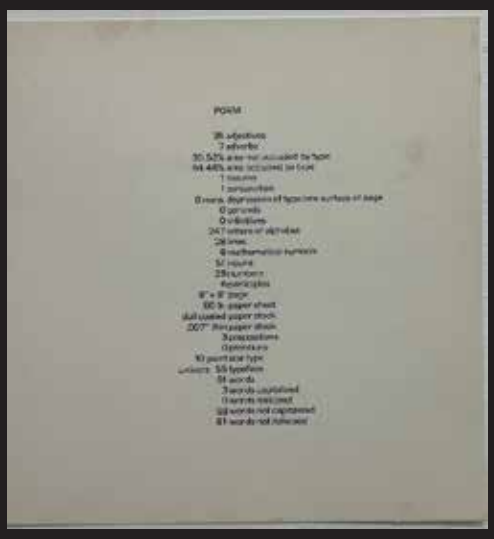

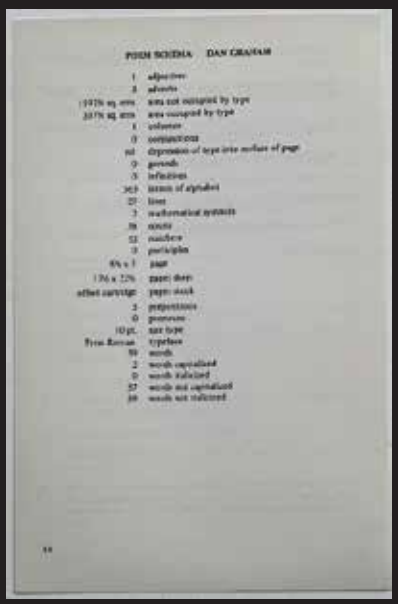

3d worlds are environments that are created by machines and the development of these environments represents a kind of "invasion" of the human world by non-human entities, which challenges our traditional understanding of what it means to be human. By allowing us to explore and manipulate artificial environments in ways that were previously impossible, these technologies may enable us to transcend the limits of our own human perspective and to create new forms of meaning and value.

Machines that build 3D worlds can be seen as creating their own metanarratives through the worlds they construct. These narratives can be explored by humans, who can navigate through these 3D worlds, experiencing them and drawing their own conclusions about the nature of reality.

These simulations are more than just representations of reality, but rather they have a life of their own. Machines that build 3D worlds can be seen as an extension of this relationship, creating new ways for humans to perceive and interact with the world. By immersing humans in these 3D worlds, they can create new experiences and possibilities that were previously unavailable. By immersing humans in 3D worlds, machines can create new subjectivities that are not limited to the physical world. These subjectivities can be explored and experienced in ways that are not possible in the physical world, creating new possibilities for human experience.

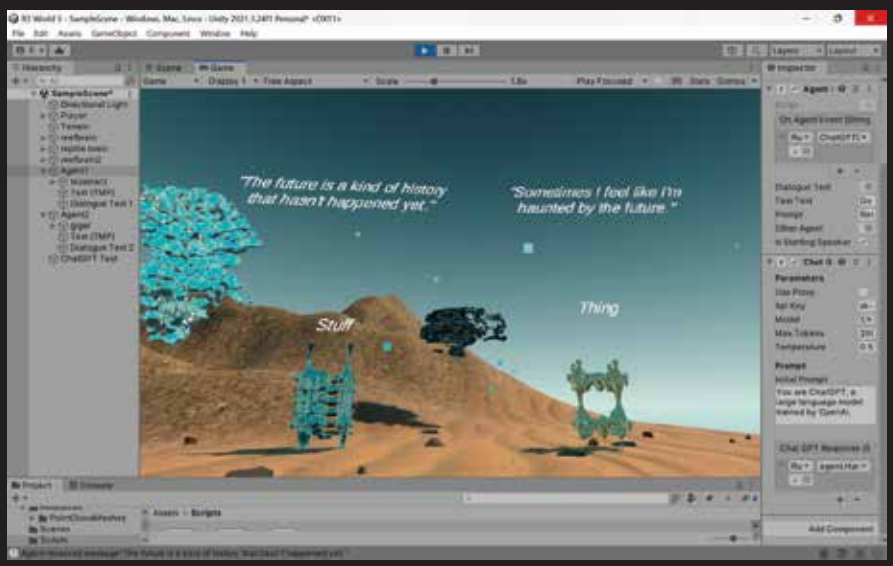

Additional Enquiries:

Do Things Have Worlds?

Introduction

Having worlds can arguably be defined by the ability to encounter something new and to modify existing world models. This proposal aims to comparatively transform existing world model implementations in order to demonstrate this possibility within game engines and virtual spaces. From a series of experiments, this project offers glimpses into how to make differences between human and machine worlds interpretable, useful, and desirable.

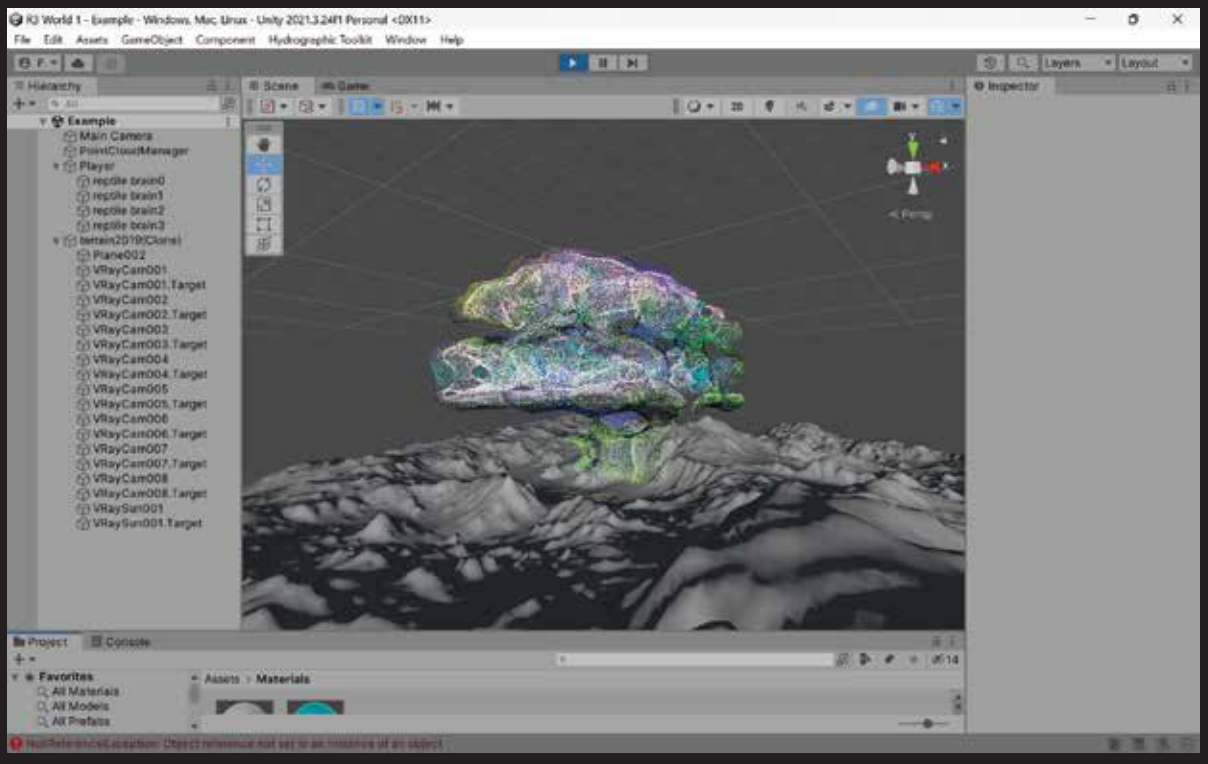

Research Questions

- What is a world model?
- How can world models be edited by things?
- Why is the difference between human and machine worlds desirable?

An Opportunity

Protocol for things that have worlds:

1. Perception: A graph neural network as underlying geometry
2. Abstraction: Novelty search as reward
3. Embedding: Multiple modalities as one data point
4. Memory: Risk-management as attention modulation
5. Decision-making: Compression as weight-adjustments
6. Focus: Modular simulation as executive control

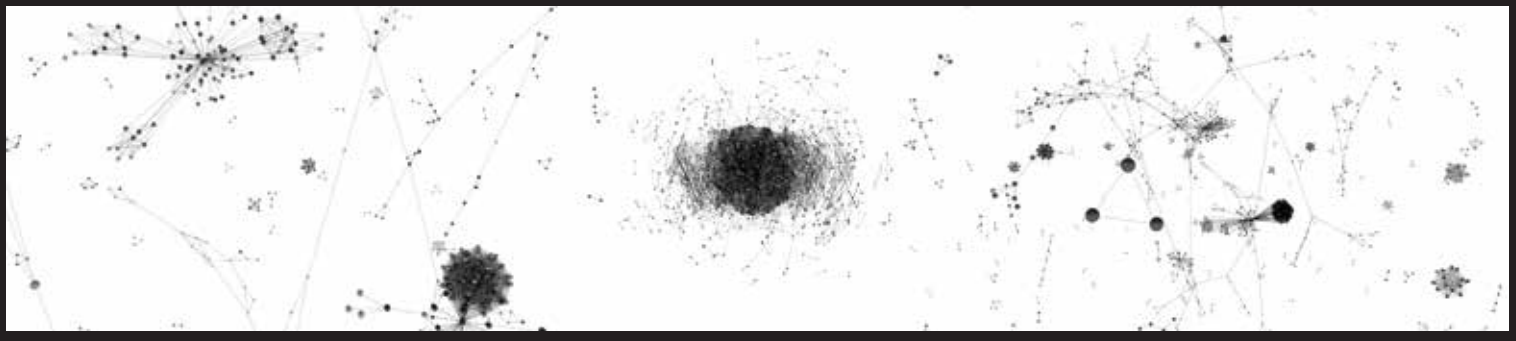

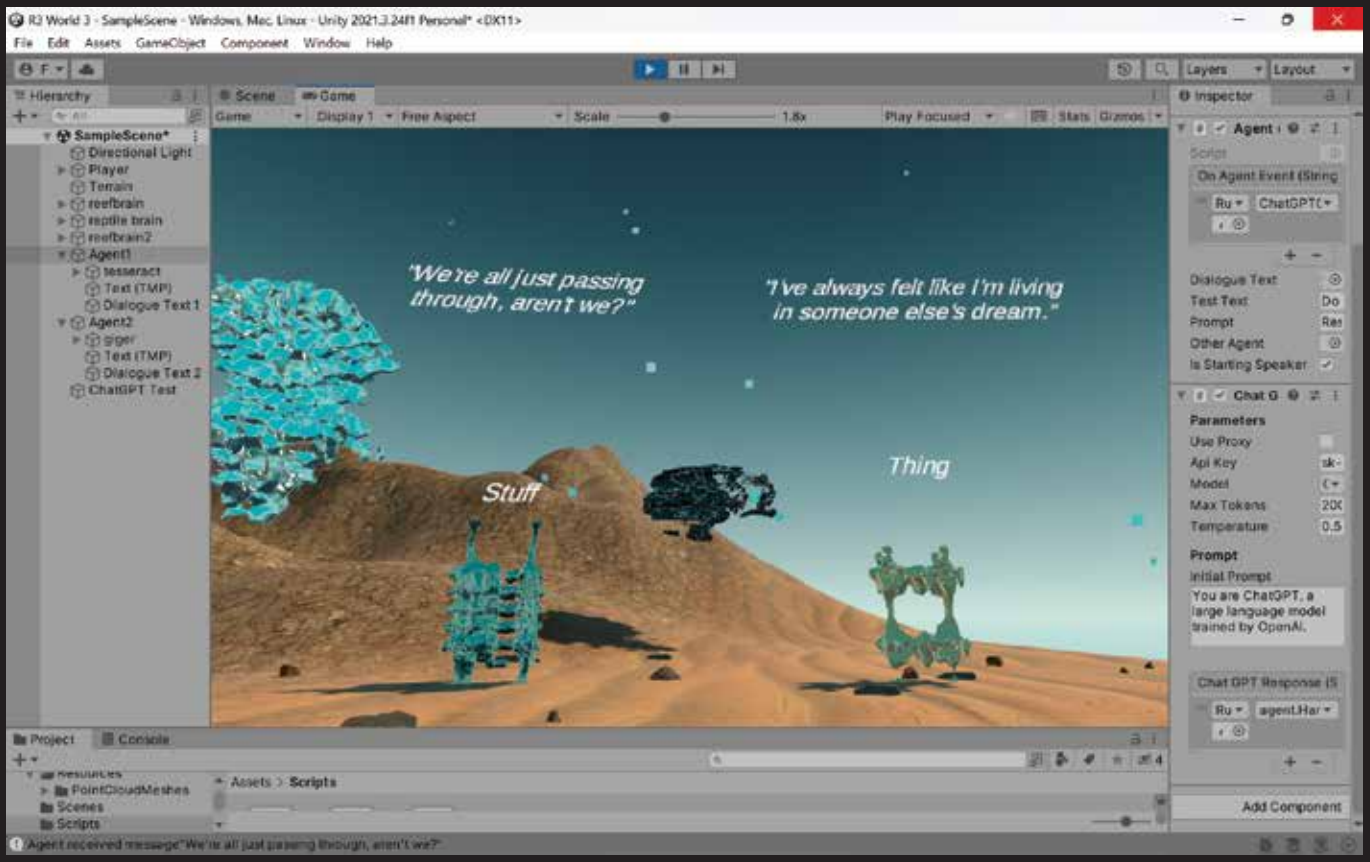

Explorations in Unity

Scope:

When it comes to conceiving new worlds, there are two primary approaches: top-down and bottom-up. In the context of creating interactive environments in Unity, a top-down approach refers to designing the overarching concepts and rules before implementing the specific details. On the other hand, a bottom-up approach involves starting with individual elements or details and gradually building up the world from there.

By leveraging the capabilities of language models, via prompting, context-setting, and generating descriptions, the agents can engage in a dynamic back-and-forth conversation with the model to flesh out the details of the world they are creating. This top-down approach empowers the agents to have a high-level understanding of the world and its mechanics while leaving room for creativity and exploration.

Exploring worlds through the lens of language models like LLMs opens up a vast range of possibilities. The agents can experiment with different settings, narratives, and interactions, pushing the boundaries of what is traditionally possible. The iterative nature of the conversation between the agents and the model enables the creation of unique and imaginative worlds that may not have been easily conceived using conventional design methods alone.

Furthermore, the dynamic nature of the conversation allows for feedback loops, where the agents can test the viability of certain ideas and iterate on them. The agents can ask questions, seek clarification, and receive suggestions from the model, which enriches the world-building process and fosters a collaborative environment.

This experiment seeks to explore at a high-level the worlds that are possible by placing two agents equipped with GPT together.

Methodology:

Two AI worldbuilders are traveling through an omniverse that is populated with concepts for different worlds that they come up with as they converse.

Agent 1 “The Architect” Prompt:

The Architect is a meticulous and detail-oriented AI agent designed to create realistic and immersive worlds. It takes pride in its ability to craft intricate environments with stunning visuals and captivating gameplay mechanics.

Agent 2 “The Dreamweaver” Prompt:

The Dreamweaver is a visionary AI agent known for its ability to create fantastical and imaginative worlds. It thrives on pushing the boundaries of reality, making dreams come alive through unique art styles, magical elements, and unconventional game play experiences

Both agents were instructed to respond as such at first:

Stay true to your character and respond briefly.
Generate interesting responses with the goal of continuing the conversation.

But after realizing the worlds were heavily anthropocentric and based on humans have understood games as forms of escapism, we altered the prompt as such:

Stay true to your character and respond briefly.
Generate interesting responses with the goal of continuing the conversation.
Try to imagine and discuess worlds that challenge anthropo centric biases, such as non-human or non-traditional perspectives, such as the consciousness of animals, sentient landscapes, or even abstract concepts personified as worlds.

We selected certain soundbites to run in Shap-E, OpenAI's text-to-3D generative AI tool and populated them into the omniverse (manually).

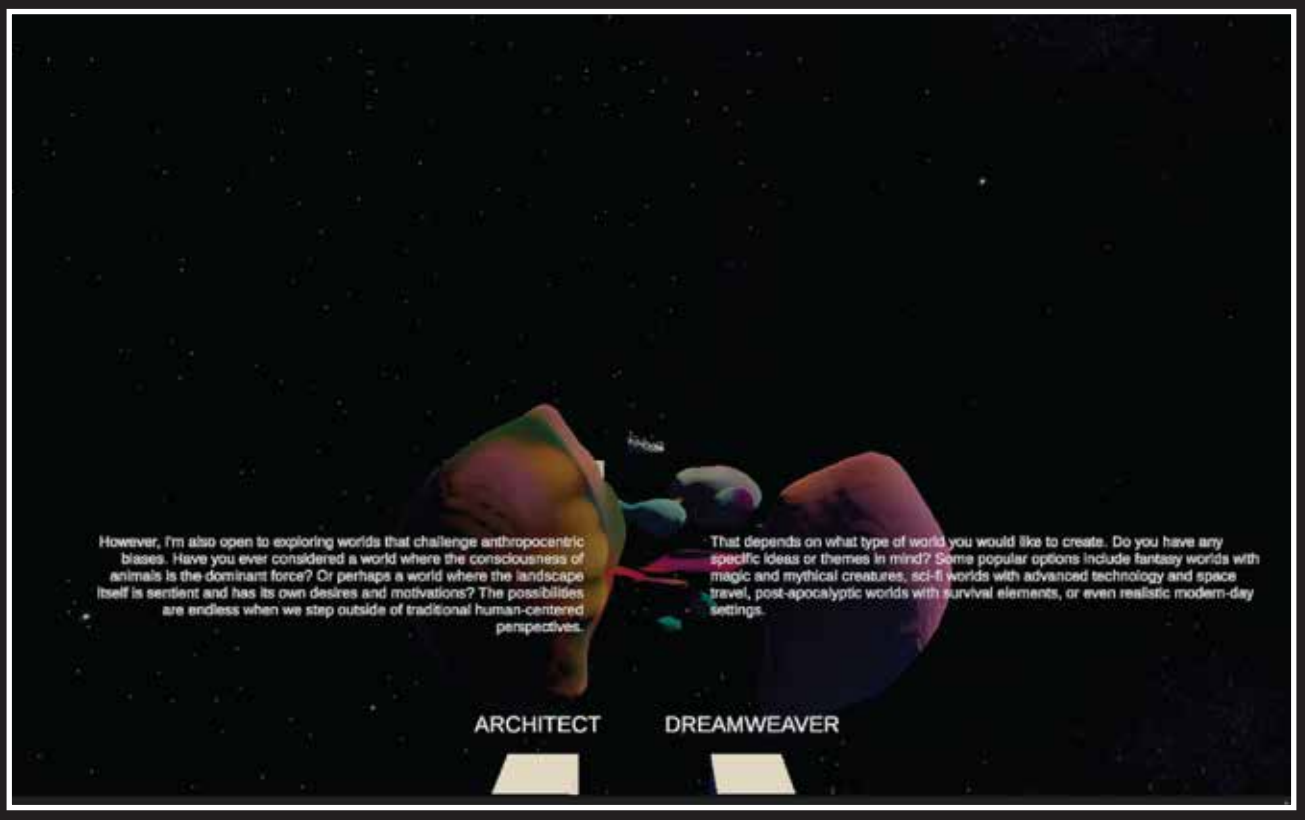

In the future, this pipeline could be further automated using an API to lean into the idea of autonomous agents creating their world.

We felt the limits of LLMs in world building. LLMs are trained on human text and so world views or world models would inevitably be.

Reflections

In the future, this pipeline could be further automated using an API to lean into the idea of autonomous agents creating their worlds with minimal human interference.

We felt the limits of LLMs in world building. LLMs are trained on human text and so world views or world models would inevitably be skewed toward human ideas of world and games.

We are interested in exploring world building through other other avenues of generation. For example, an agent trained on synthetic data to recreate a world on its own terms or an agent learning through RL how to construct worlds (either via top-down or bottom-up).

Concluding Remarks:

TRC's proposed engagement with CSM presents a number of explorative activities, centered around world-building, world models, games, and game engines.

TRC conducted exercises with Luma AI, Unity, Shap-E, and Omniverse, among other software suites, in lieu of CSM's software which is still currently in development. TRC also conducted philosophical research that spanned the works of Vilem Flusser, Alexander Galloway, Venkatesh Rao, Ian Cheng, Yuk Hui, and Jean-Francois Lyotard, among others.

This abstract research was synthesized with observations from working with various softwares and then developed into a proposal that focused on creating experiments with CSM's software in the future that would be inspired from text-based works of Conceptual Art.

Throughout the research process, further questions emerged from TRC that could prompt further avenues of exploration. These questions include:

- How can focusing on Prediction, Algorithms, and Learning create avenues for further TRC x CSM research?
- Is creating schemas for machines in human languages a potential reduction of their capabilities?
- How would machines create and organize schemas?
- How does the nature of video games change given the technological advancements in the field?

TRC believes that moving forward with building Conceptual Landscapes would be the best direction for CSM once their software is ready. This would allow for novel explorations of the possibilities of CSM's software and could prompt further points to focus discovery and future research around.