Rafael Igayara da Silva Ramos

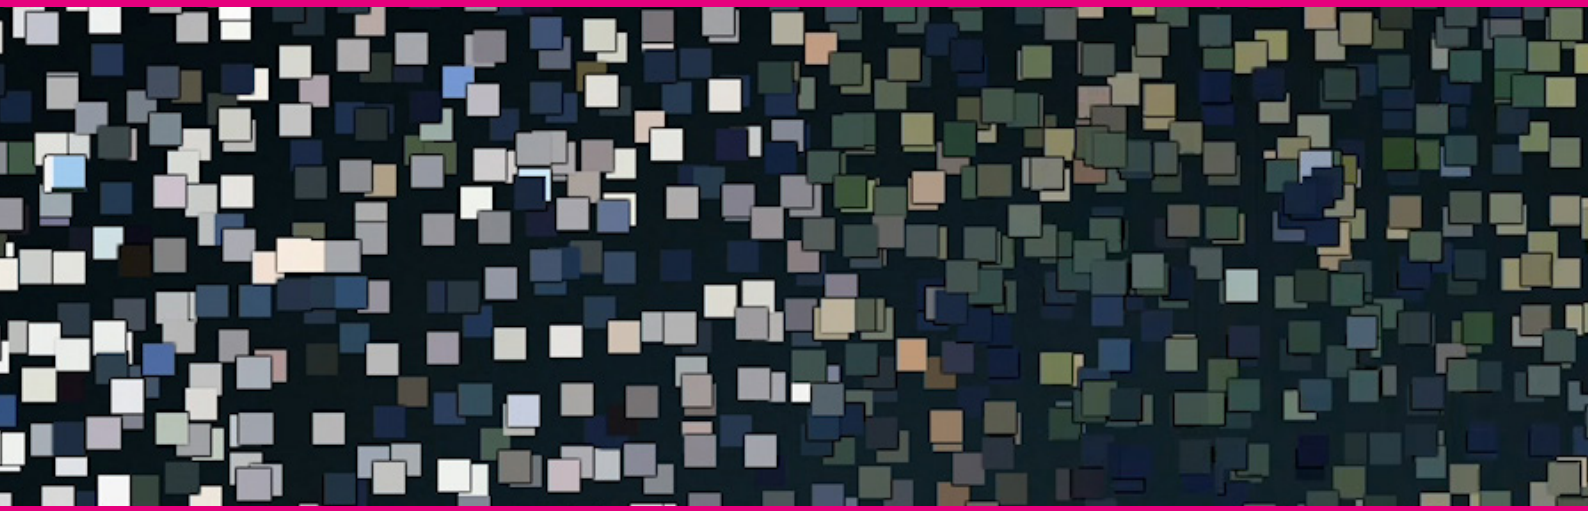

Point Cloud City

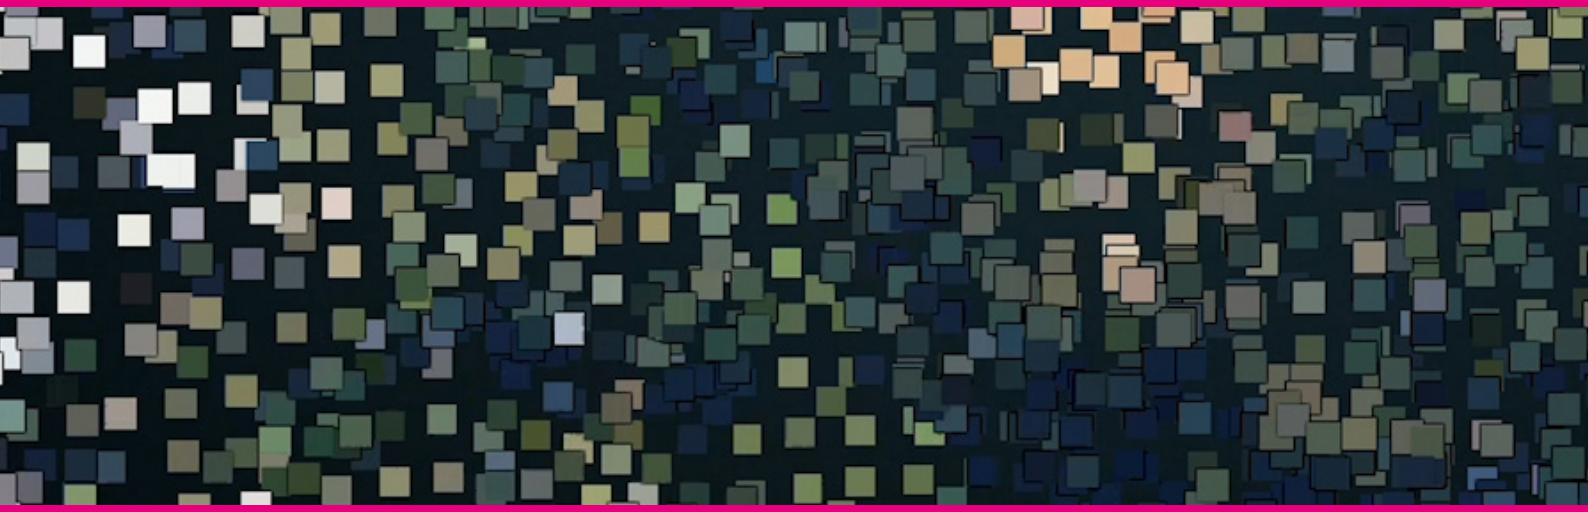

IAU5873 – Eternal Struggle – 2021

Contents

1. My Landscape Negotiation

IAU5873 – Eternal Struggle – 2021

As part of the Eternal Struggle course, we developed a series of workshops wherein we debated our particular understanding of the weekly seminars, and of the references and concepts that were discussed in the expositions.

I attempted to frame the assignments in terms of my **research landscape** and the struggles that it entails. As a MSc researcher, I am a student of Technology as it applies to Architecture and Urbanism. Specifically, I am studying how contemporary advances in artificial intelligence (AI) influence architectural design.

How do these areas of specialization respond to emerging Anthropocene issues and Ecological thinking? How might the Design and Architecture professions help bridge diverging strands of understanding of our contemporary condition emerging from disparate fields such as Environmental Humanities and STEM? Conversely, can we repurpose cutting-edge technologies from its current weaponization by capital into something that might help support a liveable planet and ecological democracy? Could a close look at worldwide indigineous, political, religious and ecological resistance movements serve as a lens with which to see through mainstream processes of ideological obfuscation?

In this first part of the pamphlet, I will present my workshop expositions. In attempting to answer some of these questions posed here and elsewhere, I have gained a clearer understanding of my field of study, the struggles and negotiations surrounding it, and specific strands of scholarship addressing these pressing issues.

Voyage Around My Environment

Introduce yourself regarding your Course Topic expectation, ecosophical curriculum and engagement, as well as final Course contribution prospective.

Our first assignment asked us to introduce ourselves in terms of our expectations and previous engagements with the course topics.

As a professional, I am a practicing architect, born and raised in the city of São Paulo, where I was also trained as an architect under the orthodox brutalist modernist tradition of the *paulista* school[1].

As such, I am very sensitive to and interested in contemporary framings of Modernity. Due to particularities in Brazil's history and politics, postmodernist discourse and aesthetics were never quite successful in setting a foothold in brazilian architecture; as much as that might have insulated us from some questionable aesthetic trends by the end of the past century, it also facilitated an untimely survival of naïve optimism in modernist rationality that survives to this day. Climate science, ecological and environmental studies thus offer an important path forward for brazilian architectural theory.

As a researcher, I have approached technology studies in Architecture from an interest in practical applications of computer programming and algorithms as design aids, a subfield known as **Computational Design[2].** This has led me to become attentive to state-of-the-art topics in Computer Science, including current advances in Artificial Intelligence. These can be contentious topics, as socially impactful as they're technically demanding.

Paulista school of brutalist modernist architecture

Computational Design

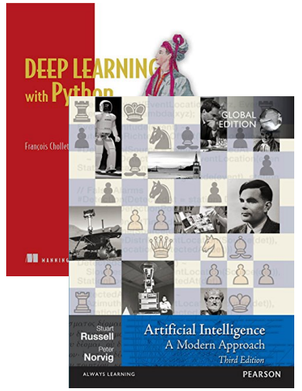

Artificial Intelligence

Mnemosyne Atlas

Through your personal investigation (daily life, living environment, regional ecologies & dynamics), and previous Works (academic, professional, volunteer); Provide a localised and intimate vision of the Anthropocene Pedigree and Epistemological Anarchism.

For this second task, I represented the way in which my ongoing efforts have already expanded my comprehension of the research topics as I originally formulated them.

Firstly, an important realization is the centrality of **Cybernetics** as a historical predecessor and theoretical foundation, not only of contemporary AI[3], but also with pioneering architects exploring Computational Design at a time where even Computer Science was still a nascent field of study[4]. Furthermore, as a general systems theory, Cybernetics from the very start places natural, social and machinic processes in equal footing[5], a welcome epistemological foundation for formulating thought and action suitable for the challenges of the Anthropocene.

From another angle, in previous research efforts I studied the work of feminist philosopher Judith Butler, whose notion of gender performativity takes another strike at common sense nature/culture dualisms[6]. In the work of Bruno Latour, I have found another opportunity for bridging this nature-culture-technology nexus, having brought forth environmental awareness by delineating the Anthropocene crisis while also having made influential works in Science and Technology Studies[7].

Lastly, for this task I also represented some important localized references from the University of São Paulo in the works of Vilém Flusser[8] and Giselle Beiguelman[9].

Cybernetics

Bruno Latour

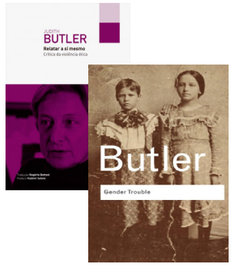

Judith Butler

Donna Haraway

Giselle Beiguelman

Mobilis in Mobili

Articulate your understanding of Intangible Landscape through non-Western /decentered European Scared and Ethics Nature/Culture relations.

For this assignment, I took the opportunity to query my research references and investigations for their engagement and geographical origins in non-Western or decentered European formulations. In doing so, I was able to find ideas and concepts of interest that explicitly relate the questions of technology, anthropocene crisis, and post-colonialist critique.

The first concept is that of **Cosmotechnics**, as formulated by chinese philosopher Yuk Hui: the notion that Technology is constrained by cosmologies which go beyond mere functionality or utility, such that *"there is no single technology but, rather, multiple cosmotechnics."*[10]

A second standout concept is **Inhumanism** as spearheaded by iranian philosopher Reza Negarestani. At the heart of the argument lies a critique both of Western essentialist humanism, but also to its shadow in anti-humanism, which stands in opposition but ultimately plays by the same rules; Negarestani argues instead that we ought to "*work our way through the problem of what it means to be human, and through this very exploration, reconstruct and reshape the human.*"[11]

Thirdly, as an investigation that is closer to my specific research focus, I'd like to draw attention to a group of scholars which define themselves as **Critical AI Studies**[12]. These are researchers which are studying and describing AI with a focus on its failures and fallacies, planetary costs, and the labour and resource extraction regimes associated with it.

Yuk Hui

Reza Negarestani

Liam Young (ed.)

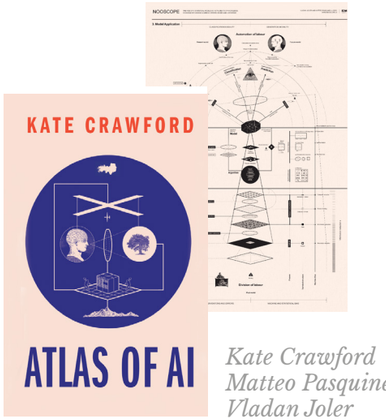

Kate Crawford
Matteo Pasquinelli
Vladan Joler

Dazibao

Articulate your understanding of Ecological Democracy through your country/region/community case studies.

How do these references help delineate an informed conception of Ecological Democracy? In an attempt to synthesize these various different, straying and often divergent pathways of thought, I have elaborated my understanding according to 4 topics, as listed to the right.

Items (1) and (2) form a description of our changing world — why formulate the question of the Anthropocene in relation to AI in the first place? It's not just that they share a place and a time, but that they emerge from shared economic, historic and environmental processes.

Item (3) offers a pessimistic account of the contemporary state of automation technology, as we have watched the rapid emergence of tech monopolies, tweeting autocrats, and decaying labor conditions worldwide.

Item (4) offers speculative account of what might lie beyond our bleak horizon, and perhaps a map for how to reach it.

1. Political participation of the non-human

Machines & Algorithms; Animals & Biomes; 'Gaia' & the Climate; populations previously dehumanized by colonialism and patriarchy — we are witnessing a proliferation of actors with varying degrees of agency.

2. Non-human territorialization

These 'new' actors have their own socio-spatial demands and regimes of territorialization. Climate refugee nomadism is a vibrant example, central to contemporary politics. Another example are the spatial prerequisites for useful representations in AI machine vision, which in turn facilitates various forms of automation.[13]

3. Automation as reification of capitalist order

We have reason to believe that current efforts to automate human activity might reify our current socioeconomic order — as contemporary AI is based on pattern recognition, we incur in numerous such risks[14], including '*failing to recognize or produce newness'* and '*replacing an episteme of causation'* to one of '*preemptive correlations'*.

4. Conceiving different cosmotechnics

Conversely, to end on an optimistic note, the Western teleology of technology as abstraction–domination is under intense questioning and scrutiny. We can and must bend the arc of technology towards the reproduction of systems and ecologies which are able to sustain life, justice and diversity on Earth.

2. Introduction

IAU5873 – Eternal Struggle – 2021

As we build up from this theoretical background, this introduction aims to frame Artificial Intelligence in terms of its definitions and objectives, historicity, contemporary status, and material demands.

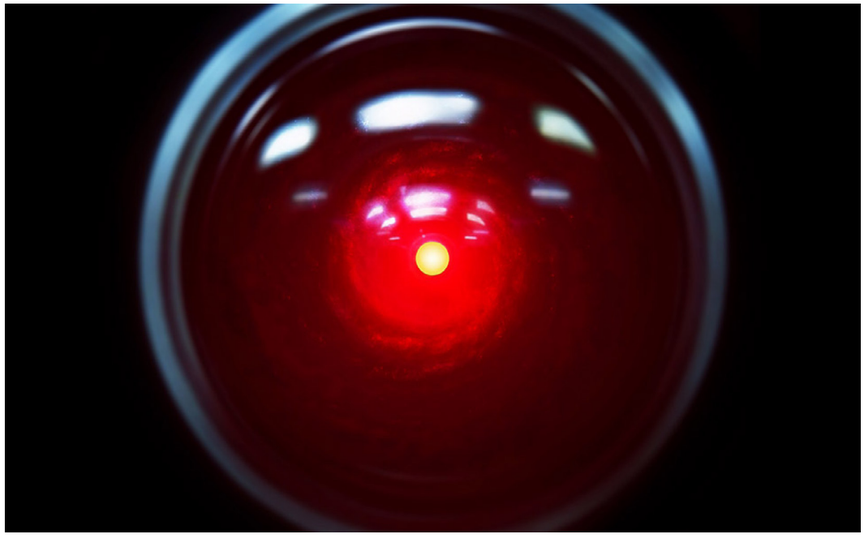

Artificial Intelligence

Our story begins with a definition of Artificial Intelligence. This is a contentious topic, surrounded by hype, overblown marketing, and existential anxiety.

On a basic level, AI can be defined as the *design of intelligent agents*[15]. What this might mean more specifically or measurably is something a bit harder to pinpoint.

Technically, when we speak of Artificial Intelligence nowadays, we are mostly speaking of Neural Networks. These are dense networks of elementary computational units, loosely inspired on the workings of animal brains. The fundamental strength of these kinds of cybernetic systems is their ability to learn when exposed to a set of training data – which is why the mainstream paradigm of AI today is called **Machine Learning**[16].

The term can be deceiving, however, since it seems to imply either the creation of a synthetic cognition or causal rationality. In fact, it does neither. AI achieves **pattern recognition** by modelling its dataset as a **probabilistic distribution**, or, as Pasquinelli and Joler put it, they're tasked with **automating the labor of perception**[17].

FIG. 1 — Organization of a biological brain. (Red areas indicate active cells, responding to the letter X.)

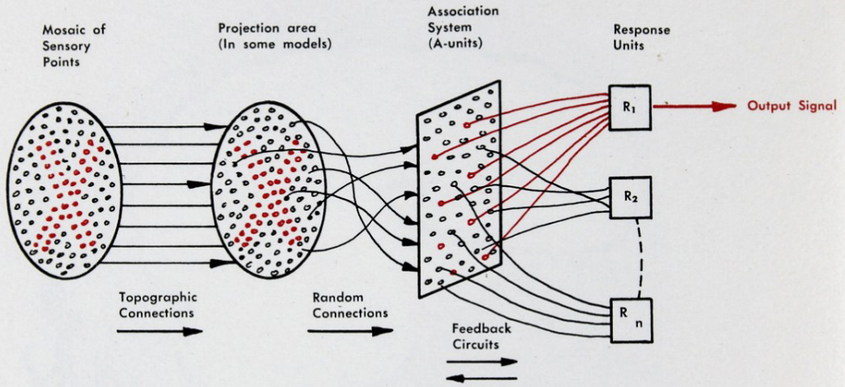

FIG. 2 — Organization of a perceptron.

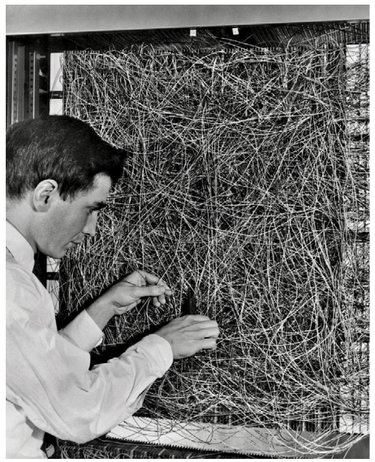

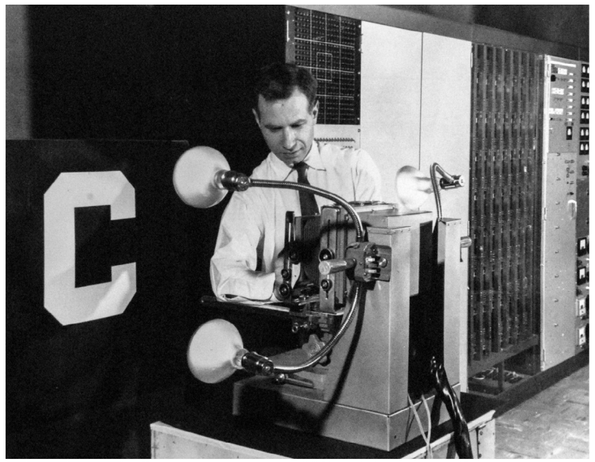

The Perceptron

The Mark 1 Perceptron, designed by the American Psychologist Frank Rosenblatt in 1958, is a famed early implementation of a Neural Network. It achieved remarkable success at the task of image recognition, and was built using specialized hardware.

As is implied by the name and diagrams, this was a system attempting to reproduce the dynamics of **Perception** in the biological brain. This focus is what sets connectivist AI apart from more traditional computer programs, which in turn leverage the logics of symbolic manipulation and if-statements to model rationality more directly. As illustrated by the messy wiring, this is a paradigm of emergence in complexity rather than elegant simplicity.

Rosenblatt was successful in his early studies and intuition, but this research path eventually went dormant[18], only to be revived in the past decade following the increased availability of computational power and accumulation of datasets from years prior, as the Internet became popularised worldwide[19].

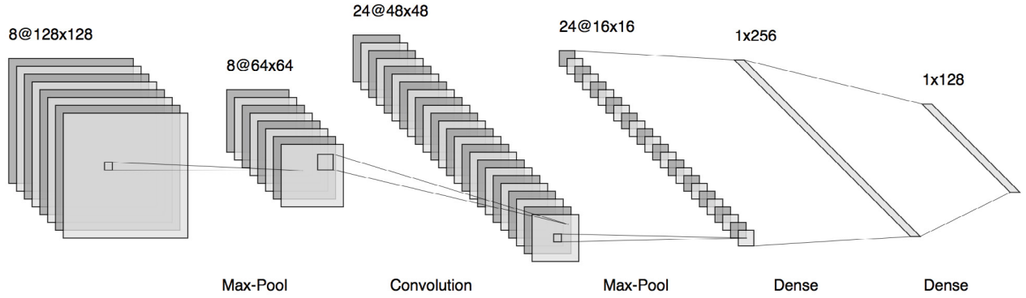

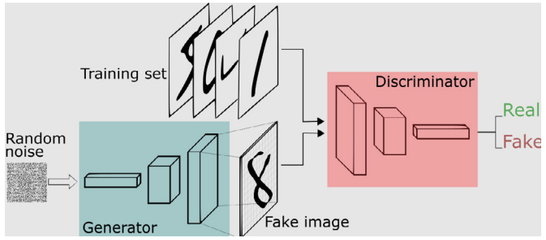

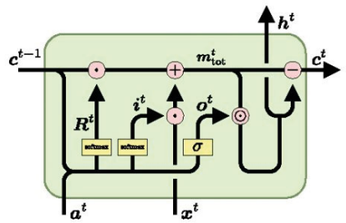

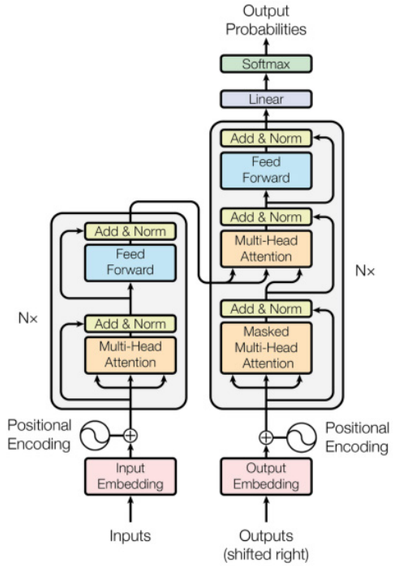

Figure 1: The Transformer - model architecture.

Neural Networks

Whereas Rosenblatt's Perceptron was a physical machine, comprised of a few hundred neurons, and a single layer deep, contemporary deep neural networks have grown in complexity.

Today, they're no longer built as single-purpose machines, but implemented as computer programs. Rather than the shallow structure of Rosenblatt's scheme, contemporary Artificial Neural Networks are *deep*, as in they're arranged in multiple subsequent layers and intricate structures, facilitating specific tasks.

Furthermore, while the Perceptron already worked by leveraging the power of cybernetic feedback loops, current ANNs make use of increasingly refined techniques for deploying these loops, such as backpropagation algorithms. This has enabled constant increases in their precision, learning rates and generalization efforts[20].

As research shows that an increase in compute tends to equate in a comparable increase in measurable performance, we have witnessed an arms race in Neural Network size, with current state-of-the art Neural Networks such as GPT-3 having been trained with as much as **hundreds of billions** of parameters – each of these AI models costing millions of dollars to produce in energy consumption and compute power alone[21].

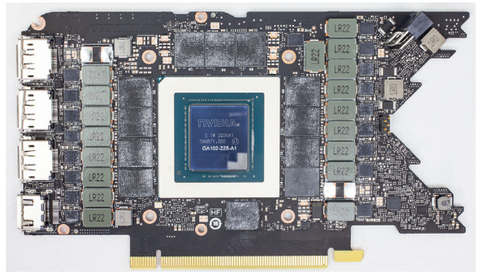

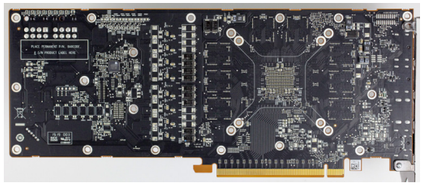

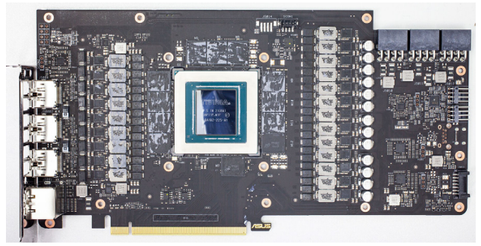

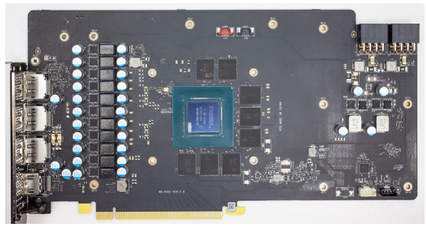

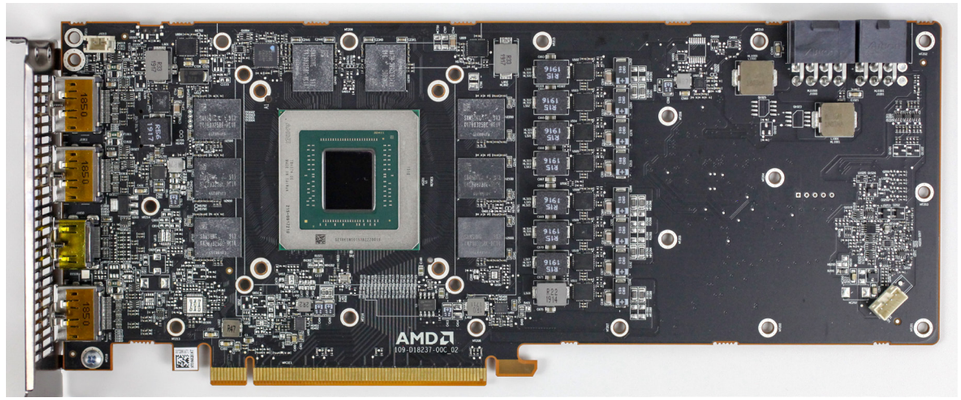

Graphics Processing Units

As we have shown, Machine Learning algorithms based on Neural Networks model biological Perception by establishing a dense, deep grid of nodes and connections. Because of its technical workings and its ultimate goal of processing vast collections of data into a single statistical model, Neural Networks require a lot of computational power; specifically, **parallel computing,** that is, a big number of small processors working in tandem.

The kinds of commercial computer processors optimized for parallel computing are called Graphics Processing Units, or **GPUs**. Over the past few years, GPU prices have soared as demand increased coming from fields such as Gaming, Machine Learning and Cryptocurrency Mining.

What are GPUs made of[22]? Apart from the versatile Silicon present in most computer chips, more specifically, GPUs require reliable transistors and capacitors to handle the very parallelism they're useful for. Some rare earth minerals used as these capacitors are Palladium and **Tantalum-Niobium**.

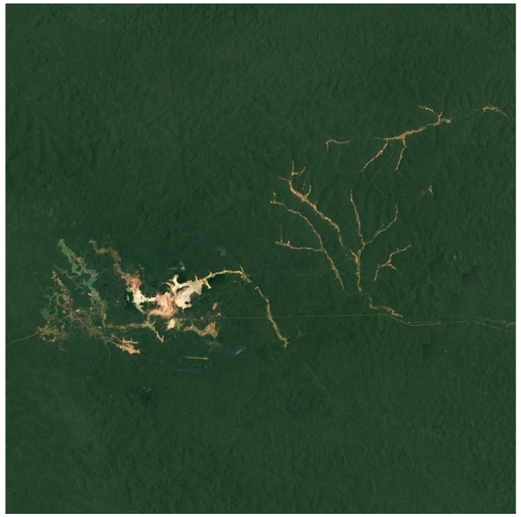

Pitinga Quarry

Tantalum-Niobium is a rare earth mineral, useful for a variety of electronic devices, with its current worldwide production mostly concentrated in Rwanda, Congo and Brazil. One of the largest Tantalum-Niobium mining spots in Brazil is the quarry of Pitinga, located in the middle of the Amazon rainforest, a few hours to the north of the city of Manaus.

Pitinga Quarry and the Balbina Dam to its south are a troublesome pair of developments of the early 1980s, set forth by the then governing brazilian military dictatorship. Both these projects take place in the former Indigenous Territory of the Waimiri-Atroari, which have been displaced and suffered great violence by the resulting invasion[23]. These projects have also caused disastrous environmental harm, contaminating waterways and harming local fauna, as well as its haphazard implementation having resulted in copious amounts of greenhouse gas emissions.[24]

Current trends in Brazilian politics echo this tragic legacy of the dictatorship era, as mining lobbyists pressure for lax environmental policies and for the deregulation of existing indigenous preservation territories.

3. Atmospheric Design

IAU5873 – Eternal Struggle – 2021

In our introduction, we have shown how investigating the technical foundations of AI leads us straight back to the Earth and to struggles surrounding natural resources and land ownership. AI is not a disembodied abstraction after all, but rather a complex system informed and enabled by the messiness and power dynamics of materiality.

As a technical system, however, Artificial Intelligence is often defined and assessed in terms of its resemblance to human behavior. For our Atmospheric Design, we have taken the inverse route – we assume the position of a curator/spectator in nonhuman processes, and probe the way these machinic systems render the territory. **What patterns does it extract? What patterns does it ignore? What patterns does it project? What patterns does it reinforce?**

Point Cloud City

As an approach to addressing these topics, we investigate a massive dataset of Point Cloud data spanning the entire city of São Paulo. This is a dataset that was created from 2017 onwards and has been released in 2019[25]. It has been produced by satellite LIDAR-scanning, the same kind of representation technology used in AI computer vision for autonomous vehicles and drones. **It represents the city as a raw, dense grid of colored points.** With an average density of 10 points/sqm, we arrive at a dataset comprised of roughly 15 trillion points. These points are then further processed to receive added metadata (e.g. differentiating between buildings and vegetation).

This kind of representation is in line with what Neural Networks demand – vast, raw, and atomic data, from where it can leverage our ever-increasing access to **computational power to recognize and extract patterns**. Text, Image, Sound or City – for AI, it's all the same – the methods are generalizeable and interchangeable, as long as your object can be represented as a simple collection of points, pixels or characters[26].

Compare this Point Cloud representation of the City with an usual architects' BIM model, predicated on hierarchical categorization and semantics of symbolic elements –that is a good data structure for the production of documentation, human legibility, project versioning and so on, but not so much for analysis by AI Neural Networks. We arrive at a trade-off: **Data that is legible to AI is not necessarily legible to humans.** This effect can be seen not only in the data itself, but on the fact that AI ingests huge amounts of data*sets,* each with so many data points that a single person cannot possibly even contemplate, let alone interpret all of it[27].

High Resolution

How does the machine see?

Our first exercise within Point Cloud City is a commentary on this tension between machine legibility and human legibility. If, as tech entrepreneurs will tell us, LIDAR-scanned point-clouds are the medium through which autonomous vehicles will see and navigate the world, then is worthwhile to delve into this representational scheme.

This struggle can also be expressed in terms of resolution. By any measure, for an object of this scale – The City – this is a high-resolution dataset. At 15 trillion points, the dataset is sufficiently dense that we can perceive urban shapes and forms.

As we progressively lower the point density – as a parsing algorithm might do, sacrificing precision for processing speed – then these shapes and patterns become increasingly indistinguishable to us. As we move away from a familiarity of a photographic image, it becomes clear that **this representation renders the City as a superposition of an abstracted Cartesian grid over the territory.**

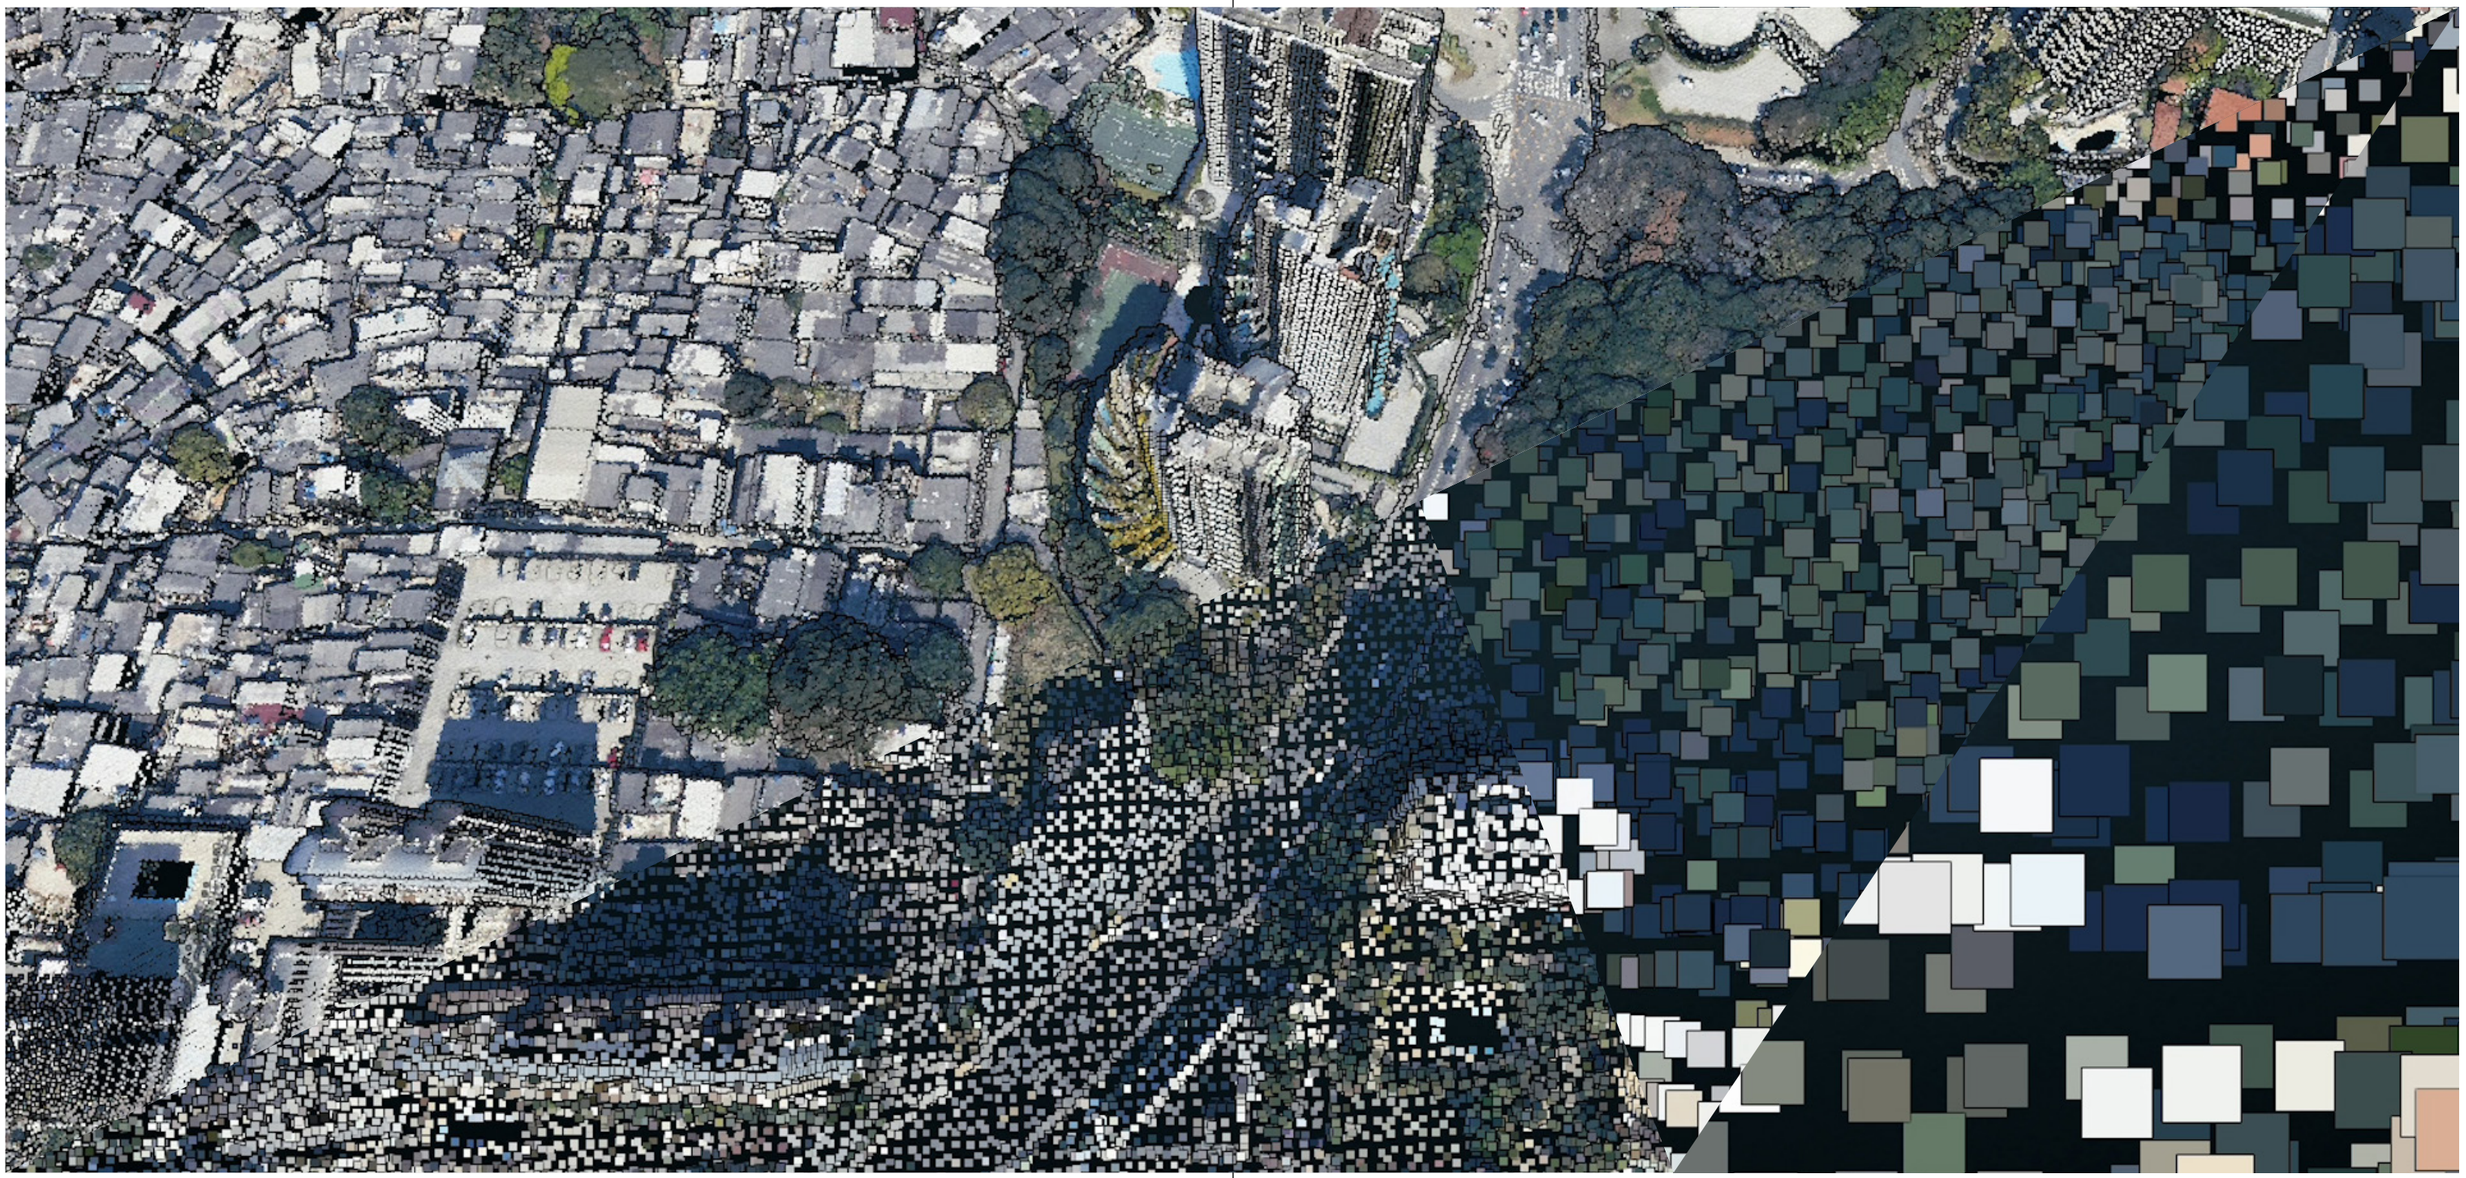

Presence

Absence

What does the machine see?

As high-resolution as this representation might be – 15 trillion points! – the city, alas, is not a single object. As we move away from the top-down view of the satellite and closer to the ground, the image of the City projected by the Point Cloud becomes ghostly and lacking.

Strolling through the densely built city center, the shape of buildings is suggested not by their façade or ground floor, but by its sky-facing rooftops. Shadows get rendered not as negative, as the presence of darkness, but as void, or the absence of data; there is no interiority to be seen either. **How can AI extract any patterns about that which cannot be captured into Data?**

This particular satellite projection, one might object, may well be enhanced and enriched by other Point Clouds captured by vehicles, drones, phones and cameras, forming together a Cloud of Clouds. But even then, which shadows and surfaces will remain unseen?

The LIDAR image of the City is at the same time dazzingly detailed but fallible and situated, its particular point of view made obvious by that which it fails to see. This is not the point of view of a Mongean projection, capturing the entirety of a 3D solid in a single 2D plane. In the Point Cloud City, nothing is solid; everything is represented as varying intensities in an Atmospheric field.

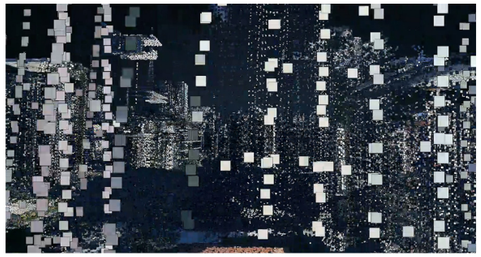

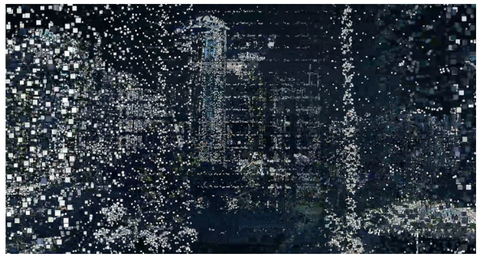

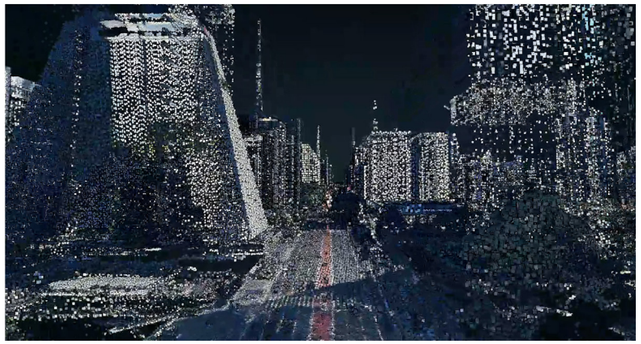

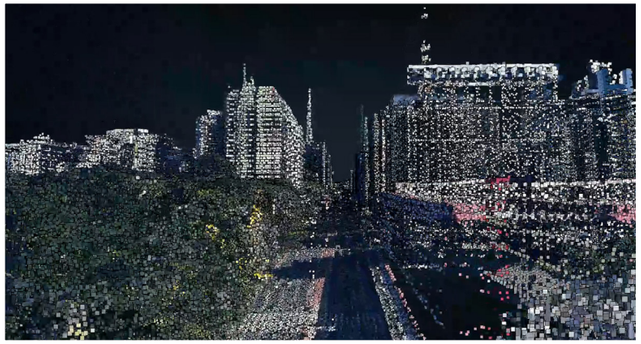

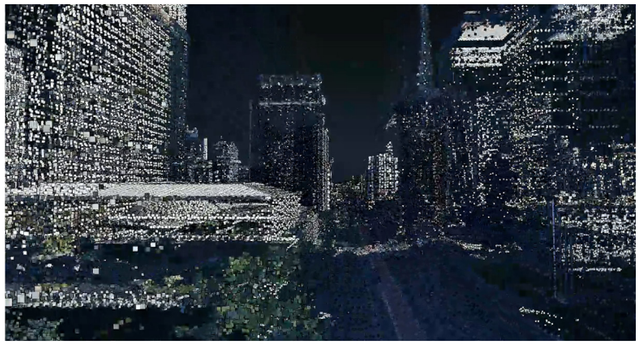

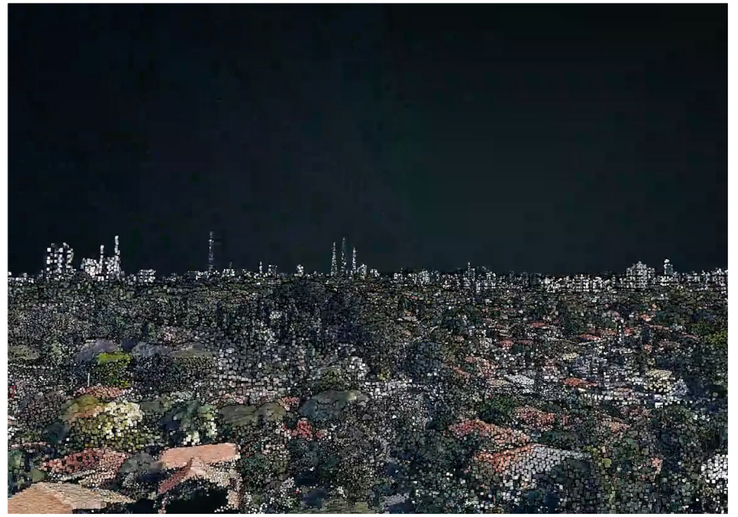

Pattern Recognition

Lastly, having looked through the machine's LIDAR eyes in our stroll of Point Cloud City, we finish our Atmospheric Design by **asking an AI to transform some of the images we've captured into new images**, matching a provided text description.

We did so by using a pair of Machine Learning techniques called VQGAN + CLIP[28]. One is tasked with generating the images (VQGAN) while another is tasked with steering the image into matching a provided text description (CLIP). These are pre-trained AI models, which have been trained on datasets comprised of images and image-text pairs collected from the internet.

This means that the patterns and images it produces are indicative of the kinds images it has been exposed to. Neural Networks, as we have discussed, are performing **pattern recognition** – it produces images not by composition or symbolic manipulation, but by trying to replicate something that would fit it with its learned statistical model of the dataset it was trained on. When we analyse these images we are assessing not only the AI models and algorithms themselves, but also the contemporary mainstream discourse and visuality that is reflected in these datasets.

This exercise makes it clear that we cannot simply accept outputs from AI as immediately valid or intelligent without peering at its interior patterns, datasets and representations.

Pattern Projection

What does the machine design?

As we mobilize AI to produce images, do we witness creativity? Those are entirely new images, but created by **projecting learned patterns** onto the provided data.

For image #1, we input as the target image description *"a prosperous valley".*

The AI takes the hint of a few existing towers in the horizon and amplifies it into a full-blown skyscraper skyline. Roofings at the foreground get morphed into crop fields, as the colors get more vibrant. An existing avenue crossing the center of the image becomes a cyan tubular structure, maybe an oil pipeline or high-speed railway. The final landscape composition is weirdly evocative of early 2000s futuristic neoliberal 'sustainable' solarpunk aesthetics. Is this what "prosperity" has come to signify?

For image #2, we input as the target image description *"a democratic metropolis".*

This time around, the AI exacerbates the tendency for pontillism in the source image of the Anhangabaú Valley, morphing even the more homogenous areas into a sequence of round shapes. It seems to equate democracy with mass public demonstrations, as façades, rooftops, trees and floorings all give way to patterns of mass public gatherings.

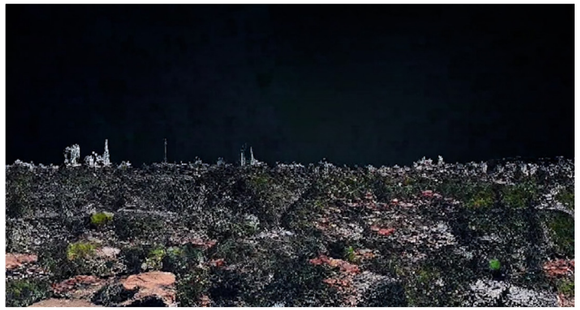
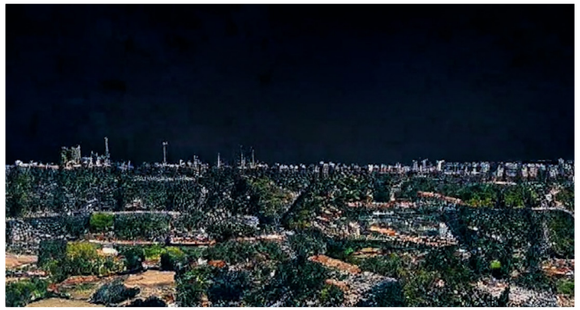

A Prosperous Valley

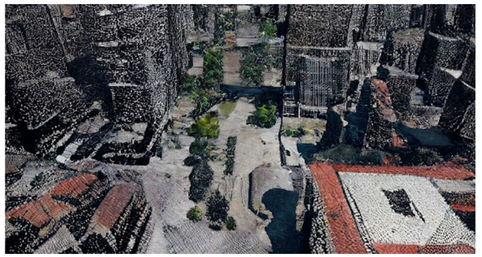

A Democratic Metropolis

4. Conclusion

IAU5873 – Eternal Struggle – 2021

We began by describing our workshop outcomes, which helped establish a theoretical foundation of authors and concepts linking Ecology, the Anthropocene, Technology, non-Western perspectives, Cybernetics and AI.

From there, in our Introduction we gave an account of Artificial Intelligence in terms of its definitions and objectives, historicity, contemporary status, and material demands. We have shown how investigating the technical and hardware foundations of AI leads us straight back to the Earth and to struggles surrounding natural resources and land ownership.

Finally, for our Atmospheric Design we have assumed the position of a curator/spectator in these nonhuman processes, and probed the way these machinic systems render the territory. We did so by first investigating a massive dataset of LIDAR Point Cloud data spanning the entire city of São Paulo, then by utilizing an AI model to transform some of the images we've captured into new images.

In doing so, we have glimpsed the particular ways in which these techniques render and model the world. As we bring its failures to attention, we are reminded that these are useful models and powerful economic forces, but not a totalizing force that is beyond political and ethical dispute.

Bibliography

BATESON, G. Steps to an ecology of mind: collected essays in anthropology, psychiatry, evolution, and epistemology. Northvale, N.J: Aronson, 1987.

BUTLER, J. Gender trouble: feminism and the subversion of identity. New York: Routledge, 2006.

CARDON, D.; COINTET, J.-P.; MAZIÈRES, A. Neurons spike back: The invention of inductive machines and the artificial intelligence controversy. Réseaux, v. n° 211, n. 5, p. 173, 2018.

CHOLLET, F. Deep learning with Python. Shelter Island, New York: Manning Publications Co, 2018.

CRAWFORD, K. Atlas of AI: power, politics, and the planetary costs of artificial intelligence. New Haven: Yale University Press, 2021.

CRAWFORD, K.; JOLER, V. Anatomy of an AI System. Virtual Creativity, v. 9, n. 1, p. 117–120, 1 dez. 2019.

HUI, Y. The question concerning technology in China: an essay in cosmotechnics. Second edition ed. Falmouth: Urbanomic, 2018

HUI, Y. MACHINE AND ECOLOGY. Angelaki, v. 25, n. 4, p. 54–66, 3 jul. 2020.

LATOUR, B. On technical mediation. Common Knowledge, v. 3, n. 2, p. 29–64, 1994.

MANOVICH, L. AI Aesthetics. Strelka Press ed. [s.l: s.n.].

MITCHELL, M. Artificial intelligence: a guide for thinking humans. New York: Farrar, Straus and Giroux, 2019

NEGARESTANI, R. Intelligence and spirit. Windsor Quarry, Falmouth, United Kingdom : New York, NY, United States: Urbanomic ; Sequence, 2018.

NEGARESTANI, R. The labor of the inhuman. In: MACKAY, R.; AVANESSIAN, A. (Eds.). . #Accelerate: the accelerationist reader. Third edition ed. Falmouth: Urbanomic, 2019.

PASQUINELLI, M.; JOLER, V. The Nooscope manifested: AI as instrument of knowledge extractivism. AI & SOCIETY, 21 nov. 2020.

PRATSCHKE, A. O barco, o mar e o timoneiro : processos de projeto e cibernética na cultura digital. Teoria e História do Urbanismo—São Carlos: Universidade de São Paulo, 31 jan. 2019

RUSSELL, S. J.; NORVIG, P. Artificial intelligence: a modern approach. Third edition, Global edition ed. New York: Pearson, 2016.

YOUNG, L. (ED.). Machine landscapes: architectures of the Post-Anthropocene. Oxford: John Wiley & Sons, 2019.

ZEIN, R. V. A arquitetura da escola paulista brutalista 1953-1973. Doctorate—[s.l.] UFRGS, 2005.

Endnotes

1 (ZEIN, 2005)

2 (MENGES; AHLQUIST, 2011)

3 For an introductory account of the history of modern AI, including the troubled origin of the term and its relation with Cybernetics, see (CARDON; COINTET; MAZIÈRES, 2018)

4 (PRATSCHKE, 2019)

5 On the topic of overcoming dualism, see (HARAWAY, 1985) or (HUI, 2020)

6 Butler's notion of gender perfomativity argues that gender is created and sustained by the very stylized acts that we retroactively ascribe to a stable, gendered identity. This is a notion drawn from Michel Foucault's concept of power as positive, that is, as actively producing subjectivities which are aligned to itself, rather than oppressing some inner-held truth. (BUTLER, 2006)

7 (LATOUR, 1994)

8 Vilém Flusser (1920-1991) was a Czech-Brazilian philosopher, who taught at the University of São Paulo in the 1960s.

9 Giselle Beiguelman (1962) is an artist, curator, researcher and professor at the Faculty of Architecture and Urbanism of the University of São Paulo. https://www.desvirtual.com/

10 (HUI, 2020)

11 (NEGARESTANI, 2019)

12 https://allmodels.ai/

13 (YOUNG, 2019)

14 (PASQUINELLI; JOLER, 2020)

15 (POOLE apud RUSSEL; NORVIG, 2016)

16 For an introductory overview of the technical workings of AI, see (RUSSELL; NORVIG, 2016)

17 (PASQUINELLI; JOLER, 2020)

18 https://news.cornell.edu/stories/2019/09/professors-perceptron-paved-way-ai-60-years-too-soon

19 While there are different accounts to explain the resurgence of Neural Networks and Machine Learning, the availability of parallel computing from the commercial development of GPUs and the ability to amass large-scale datasets from Internet data have certainly been contributing factors, among others, as detailed by (CHOLLET, 2018)

20 ibid.

21 https://bdtechtalks.com/2020/08/17/openai-gpt-3-commercial-ai/

22 We owe this idea of looking at AI in terms of its hardware demands to (CRAWFORD, 2021)

23 For an account of humans rights violations during the military dictatorship, including a specific chapter dedicated to the violations of the rights on indigenous peoples, see http://cnv.memoriasreveladas.gov.br/images/pdf/relatorio/volume_2_digital.pdf

24 https://amazoniareal.com.br/mina-do-pitinga-35-anos-de-controversias-e-nada-a-comemorar/

25 https://gestaourbana.prefeitura.sp.gov.br/noticias/geosampa-veja-a-cidade-de-sao-paulo-em-3d/

26 A very clear example of this effect can be seen in how researchers were able to use the "Transformer" architecture of neural networks for image processing. While this technique was initially devised for processing text data, understood as an ordered sequence of characters, it was able to process images just fine simply by parsing pixel data as if it were sequential data. (DOSOVITSKIY et al., 2021)

27 Recently, many projects have tackled the topic of bringing datasets to light; among these we may cite https://excavating.ai/ (CRAWFORD; PAGLEN, 2019)

28 VQGAN+CLIP ran on a Google Colab notebook authored by Katherine Crowson, available at https://github.com/crowsonkb with extended credits.