

Jan. 4th, 2021

Day 1: Introduction

Exploring AI and Neural Nets in Design

Gia Jung

Research Associate, Lab for Design Technologies, Harvard University
Irving Innovation Fellow 2020 - 2021

Claire Djang

Lab for Design Technologies, Harvard University
Currently at Certain Measures

1.1

Introduction

Who are we?

Who are you?

Class Overview

AI in Architecture

AI / DL / ML / NN

1.2

Applications

Deep Learning Projects in Design

Framework: Data-Model-Project

Example Projects

Focus Project
“Beaux-arts Latent Walk”

1.3

Workshop

Toolset

Data Scraping

Who Are We?

Gia Jung

Research Associate, Lab for Design Technologies, Harvard University
Irving Innovation Fellow 2020 - 2021

Claire Djang

Lab for Design Technologies, Harvard University
Currently at Certain Measures



Lab for Design Technologies
<https://research.qsd.harvard.edu/ldt/>

Who Are You?

Austin Lu (MLA II)
Kaihong Gao (MDes Tech)
Peitong Chen (MDes)
George Guida (MArch II)
Benjamin Villa
Yuxin Yang (MLA I AP)
Tsz Kit Justin Ng (MArch 1)
Elisa Ngan (MDE)
Siqi Joyce Zhu (MLA II)
Dongyun Kim (MDes Tech)

Ana Gabriela Loayza (MArch II)
Connie Chang (MDes Tech)
Yunzi Shi (MArch I)
Audrey Watkins (MArch I)
Rhea Jiang (MDes)
Em Sawyer (MArch I)
Kritika Kharbanda (MDes EE)
Jessica Chen (MDes EE)
Tiangang Lyu (MLA I)

Who Are You?

Go To: <http://bit.ly/contact-gsdai>

Austin Lu (MLA II)	Ana Gabriela Loayza (MArch II)
Kehe Gao (MDes Tech)	Connie Chang (MDes Tech)
Peltong Chen (MDes)	Yunzi Shi (MArch I)
George Guida (MArch II)	Audrey Watkins (MArch I)
Benjamin V. B...	Phoebe Jiang (MDes)
Yuxin Yang (MDes)	...
Tsz Kit Justin Ng (MArch I)	Kritika Kharbanda (MDes EE)
Elisa Ngan (MDE)	Jessica Chen (MDes EE)
Siqi Joyce Zhu (MLA II)	Tiangang Lyu (MLA I)
Dongyun Kim (MDes Tech)	

Class Overview

Go To:

<http://bit.ly/syl-gsdai>

Housekeeping

Google Drive Invitation:

Class Slack: <http://bit.ly/slack-gsdai>

Office Hours: Tues/Thurs 1PM-2PM, or by appointment

[Install Anaconda:](#)

Anaconda is a distribution of the Python and R programming languages for scientific computing, that aims to simplify package management and deployment. The distribution includes data-science and machine learning packages suitable for Windows, Linux, and macOS. We will be using Anaconda throughout this workshop, and we ask you to please install it ahead of the first session. Please follow this quick guide according to your operating system. Feel free to reach out if you have any issues or questions with installation.

Go To: <https://docs.conda.io/projects/conda/en/latest/user-guide/install/>

1.1

Introduction

Who are we?

Who are you?

Class Overview

AI in Architecture

1.2

Applications

Deep Learning Projects in Design

Framework: Data-Model-Project

Example Projects

Focus Project
“Beaux-arts Latent Walk”

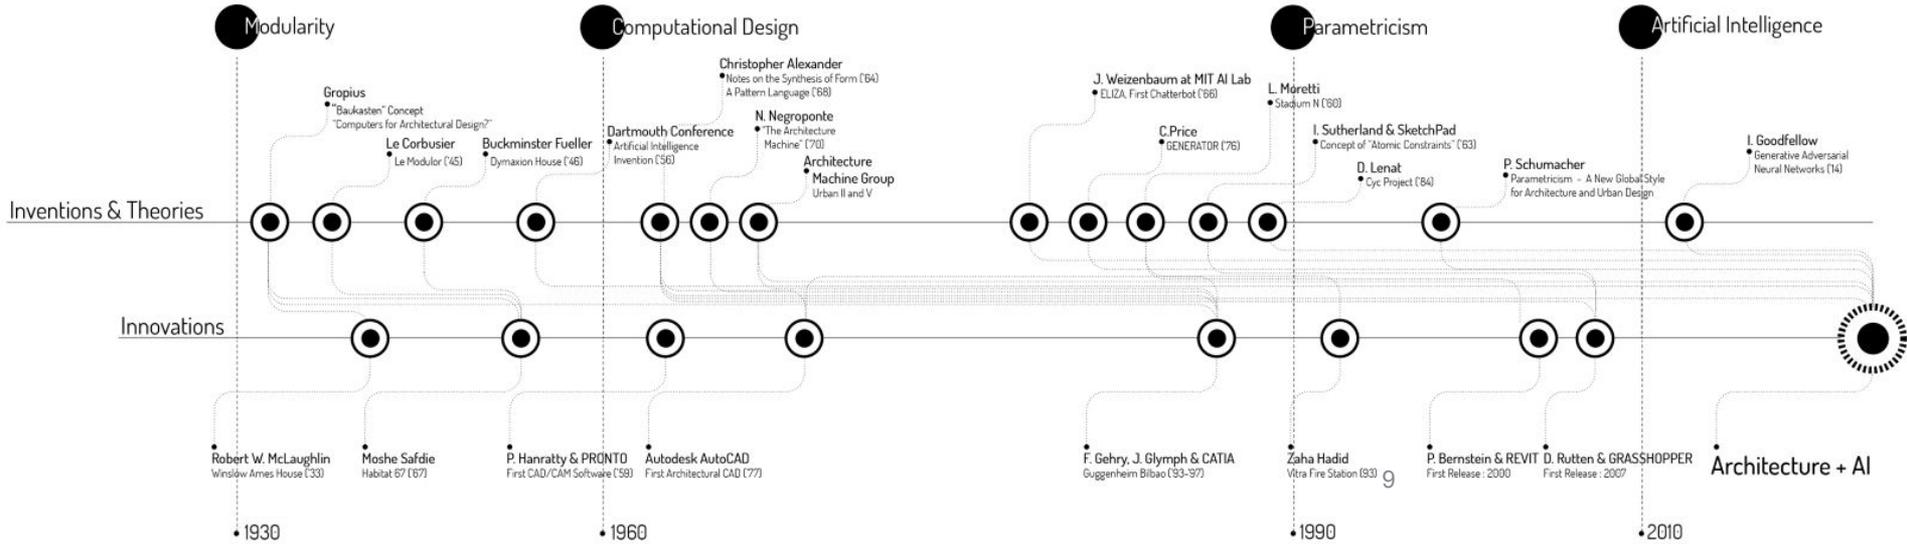
1.3

Workshop

Toolset

Data Scraping

AI in Architecture: History



History of Computation in Design, Stanislas Chaillou (2019)

AI in Architecture: History (1) Modularity 1930-1960

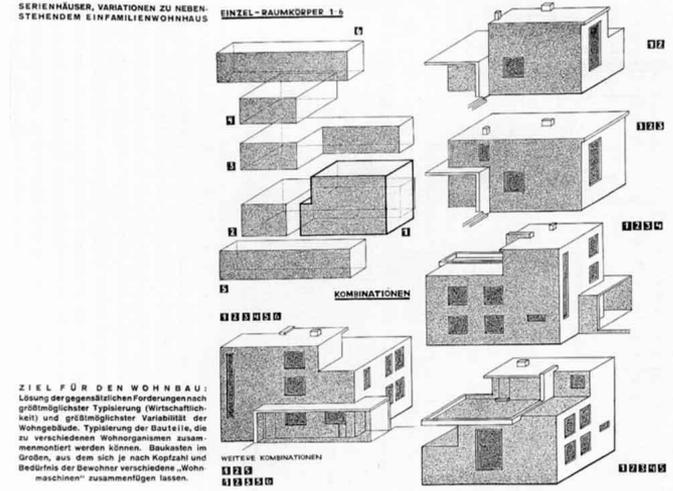
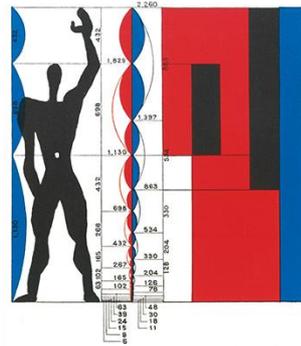
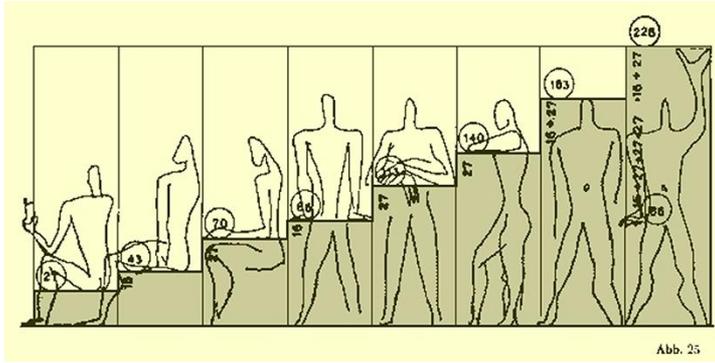
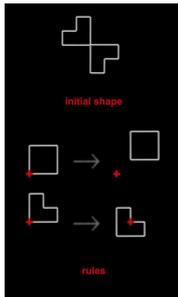


Abb. 12 Walter Gropius und die Architekturabteilung des Staatlichen Bauhauses, Baukasten im Großen, Serienhäuser in Variationen, 1923

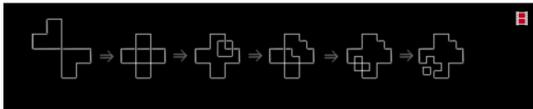
Walter Gropius, Modular Bauhaus (1923) from [Bauhaus Archive](#)
Le Corbusier, Le Modular (1945) from [Foundation Le Corbusier](#)

AI in Architecture: History (2) Shape Grammar 1960-1990

To the right is a two-rule grammar that illustrates these properties. The first rule shifts a square halfway along a diagonal axis of the square. The second rule shifts an L-shape, also along a diagonal axis. Registration marks in each rule show the positions of the shapes on the left-side and right-side of the rule relative to each other. The starting shape for computations, called the initial shape, consists of two L-shapes. The two rules apply to this shape and to shapes produced from it by matching the square or L-shape on the left-side of either rule with a square or L-shape in a design. The square or L-shape in either rule may be translated, rotated, reflected, or scaled in order to match a shape in a design. If a match is made, the matched shape in the design is then replaced with a shifted shape as indicated in a rule. The direction of the shift depends on the spatial transformation used to make the match.



Below is a computation of a design using the grammar. From the second step on, the rules can apply to either emergent L-shapes or emergent squares. Also from the second step on, either the first or the second rule can be applied to a design. The user of the grammar, human or machine, must decide which rule to apply and to which shape in a design to apply the rule.



Below is another computation using the grammar. The computation is identical to the one above in the first three steps. Then it diverges and follows a different path to produce a different design. Many other computations are possible with the grammar.

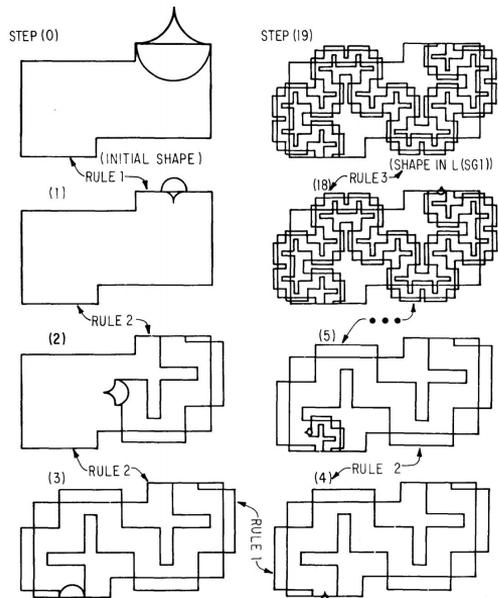
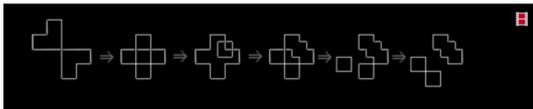


Figure 6-3. Generation of a shape using SG1.

Shape Grammars, K. Knight, [MIT](#)

Shape Grammars were first invented by George Stiny and James Gips in their 1972 article *Shape Grammars and the Generative Specification of Painting and Sculpture*

AI in Architecture: History (2) CAD 1960-1990



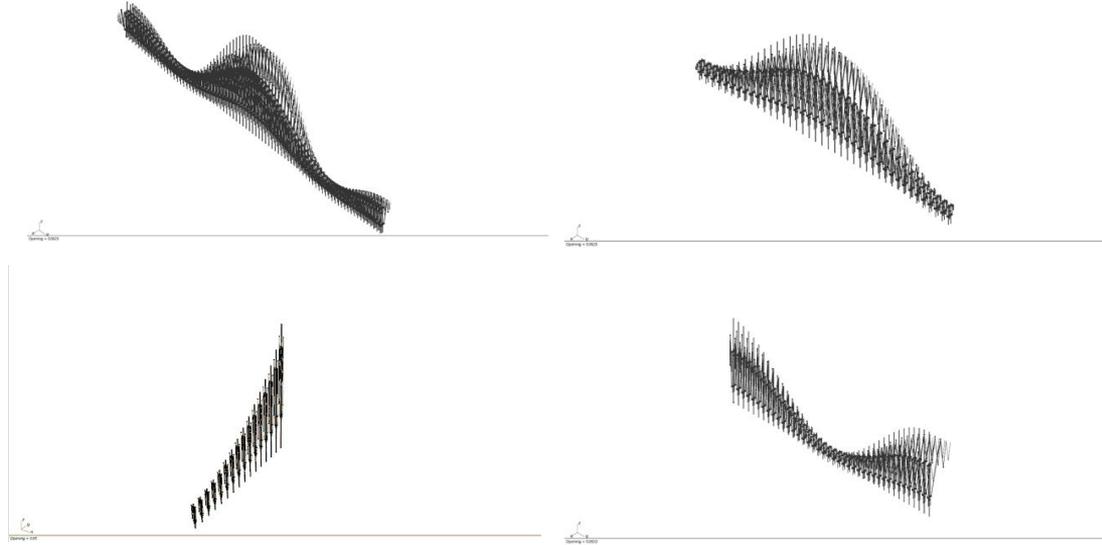
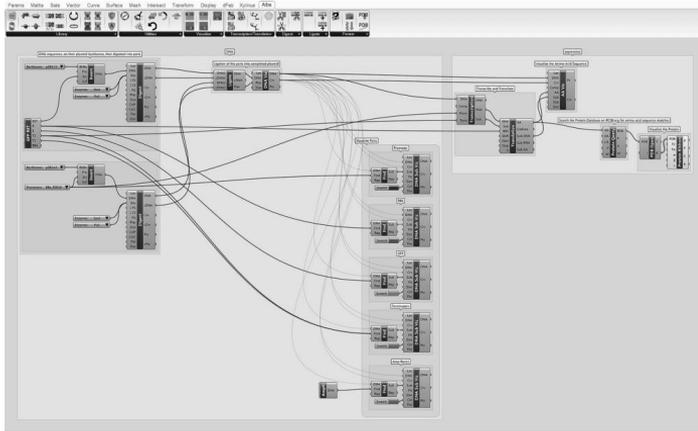
1976

CAD

Masanori Nagashima au MIT Architecture Machine Group

CAD (1976) Masanori Nagashima at MIT Architecture Machine Group
from [AI & Architecture](#)

AI in Architecture: History (3) Parametricism 1990-2010



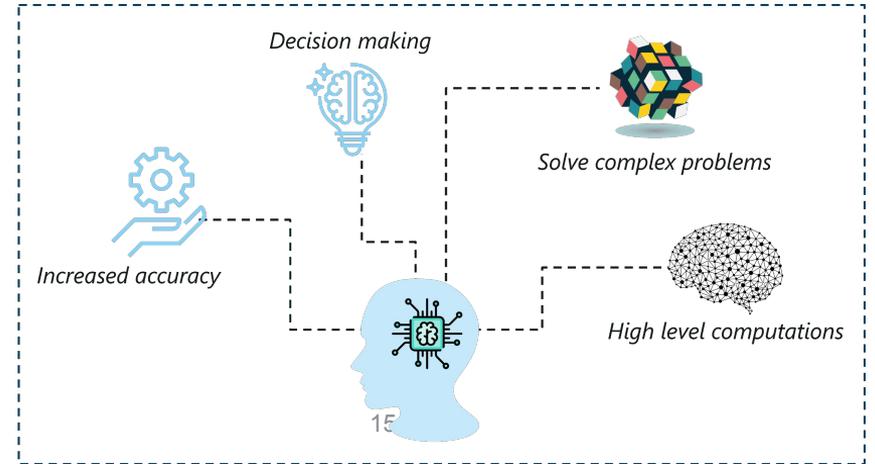
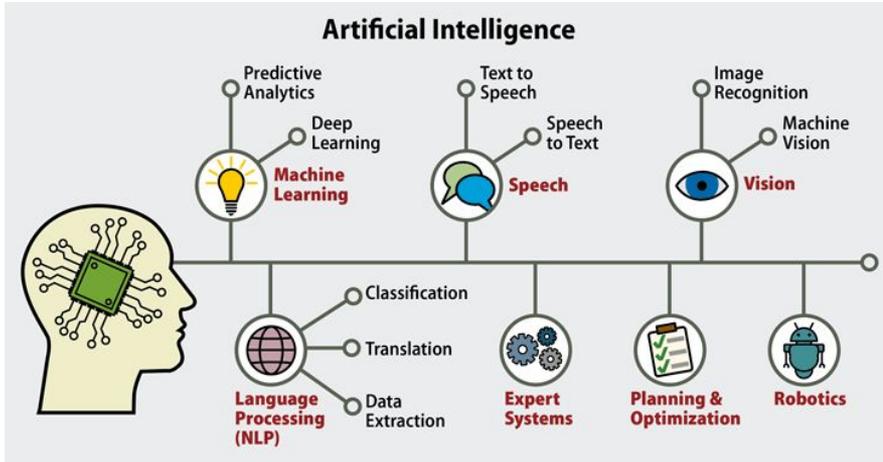
Grasshopper view (Left)
Example of Parametric design from [Las Paraguas](#) (Right)

AI in Architecture: History (4) **Artificial Intelligence 2010-**



Ai in Architecture: Historical Perspectives (2019)
Stanislas Chaillou

AI in Architecture: What is AI?



What is AI?
from [Datamation](#) and [Edureka](#)

AI in Architecture: What is AI?

By Capability

Narrow AI

can only perform dedicated task
e.g. Siri, Watson, AlphaGo

General AI

can function like humans

Strong AI

can function better than humans

By Functionality

Reactive Machines

react to current scenarios only
e.g. AlphaGo

Limited Memory

can store limited memory for a short while
e.g. Self-Driving Cars

Theory of Mind

understands human emotions, beliefs, people

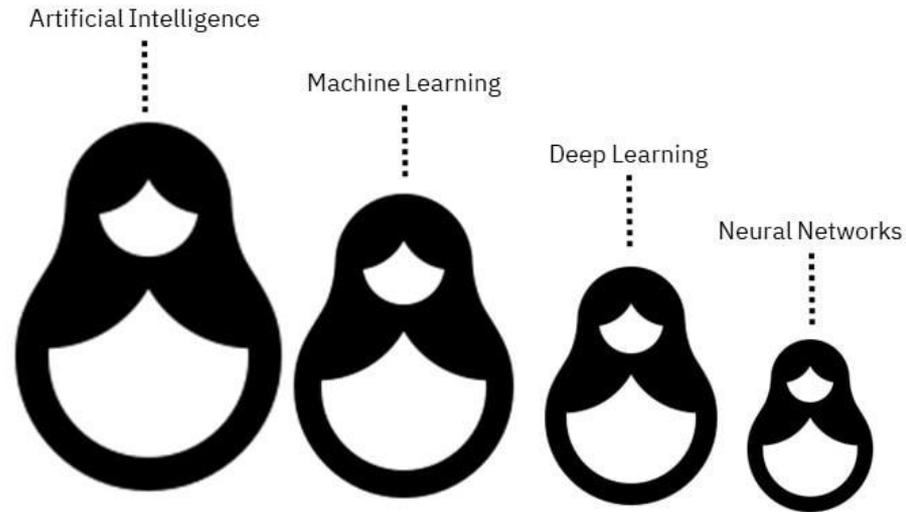
Self Awareness

possesses its own consciousness, sentiments

Types of AI

<https://www.javatpoint.com/types-of-artificial-intelligence>

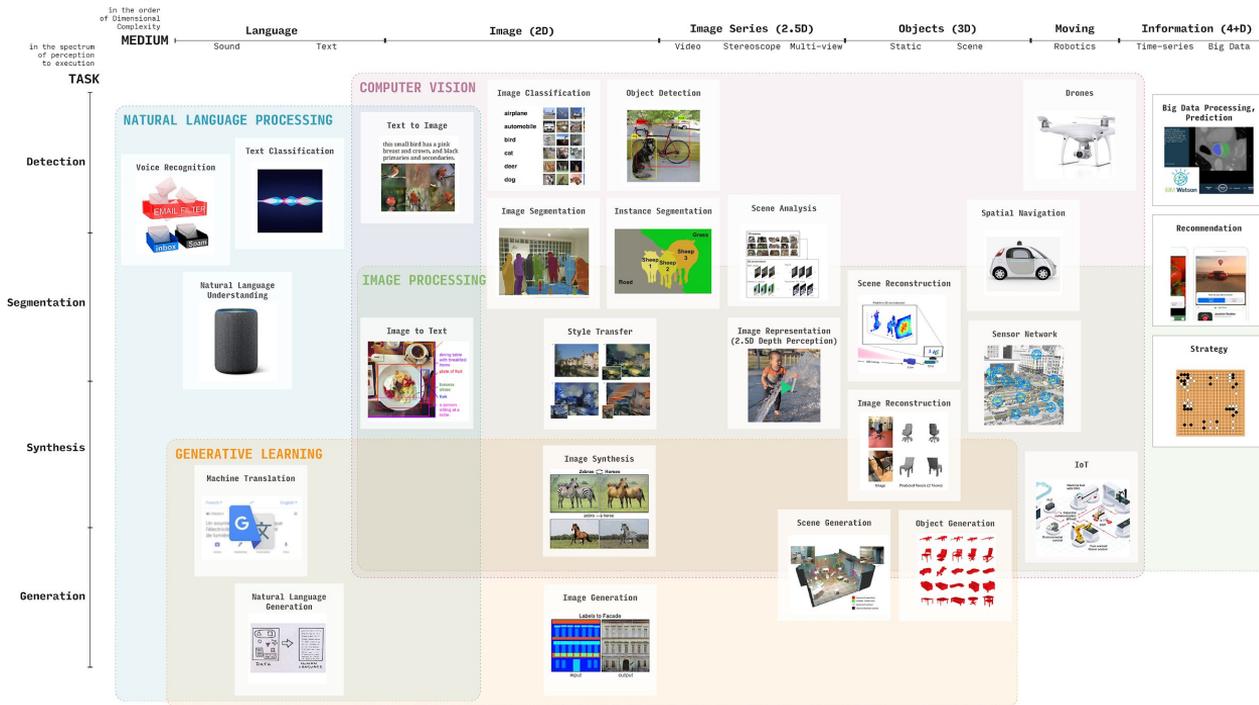
AI in Architecture: AI / DL / ML / NN



AI / ML / DL / NN

<https://www.ibm.com/cloud/blog/ai-vs-machine-learning-vs-deep-learning-vs-neural-networks>

AI in Design

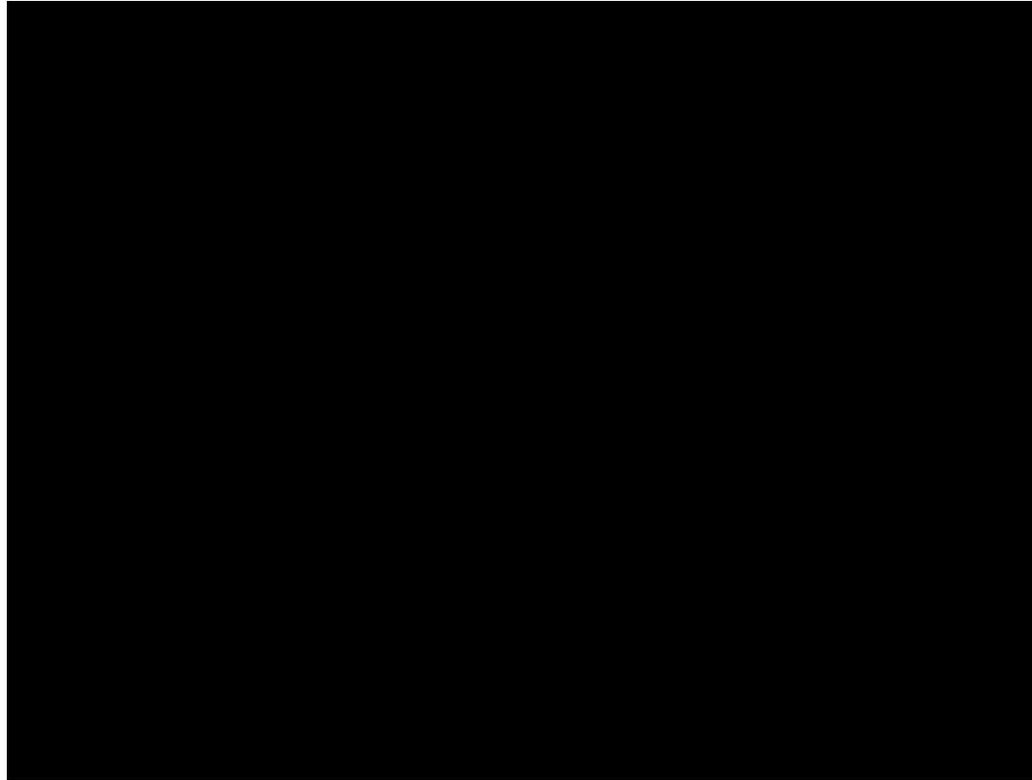


AI in Architecture, Design & Urban Environment (2020)
 Lab for Design Technologies
 Prof. Andrew Witt, Gia Jung, Claire Djang



62 62 63 64 65 66 67 67 69 70 71 72 72 73 73 73 72 72 71 70 69 67 66 66 66 65 63 62 61 60 60
61 62 63 64 66 66 67 68 68 69 70 71 71 72 72 73 72 72 71 71 70 69 68 66 66 65 65 63 62 61 60 60
61 62 63 64 66 66 68 68 69 70 70 71 72 73 73 73 72 72 71 71 69 68 67 66 66 65 65 64 63 62 61 61
61 63 64 64 66 67 68 68 68 69 70 71 71 73 74 74 73 73 73 71 70 69 68 66 66 65 64 63 62 61 61 60
61 63 64 65 67 68 69 69 70 70 71 71 72 55 53 69 72 72 71 71 70 69 68 67 66 65 64 63 62 60 60 60
63 64 65 66 67 68 69 69 70 70 71 72 42 4 5 11 48 72 71 71 69 69 68 67 66 65 64 62 62 60 59 59
63 65 66 66 68 68 69 70 71 71 71 72 18 4 4 7 8 66 71 70 69 68 68 67 66 65 64 63 61 59 59 58
63 65 67 67 68 69 69 70 71 71 72 64 4 27 24 54 33 29 52 64 68 68 67 66 65 64 63 62 61 59 58 58
64 65 66 66 68 69 70 71 41 24 24 12 17 24 48 60 37 43 36 52 66 68 67 66 65 64 63 61 60 59 58 57
65 66 67 67 68 69 71 48 6 6 6 5 34 36 12 47 34 17 29 54 43 63 67 66 65 64 63 62 60 59 58 57
64 65 66 66 68 69 38 6 6 5 5 7 16 19 4 47 44 27 24 40 67 66 66 65 65 64 63 61 60 59 58 57
63 64 65 65 67 30 6 6 5 5 5 6 8 9 20 27 51 78 41 44 66 65 65 65 64 63 62 60 59 58 57
63 64 65 65 34 5 5 5 5 5 5 4 19 6 7 54 64 20 59 65 65 64 64 64 63 62 61 60 59 57 56
63 64 64 65 14 5 6 6 5 4 5 4 18 7 5 4 19 10 11 65 64 64 64 63 61 66 62 61 60 59 58 56
63 64 64 65 53 7 4 5 6 6 7 10 6 5 5 4 21 24 18 64 64 64 63 62 64 65 62 62 60 59 58 57
64 64 64 64 65 50 4 4 4 5 11 16 6 6 4 6 35 16 26 66 64 64 63 61 72 67 63 62 61 59 58 57
64 64 64 64 65 46 4 4 4 5 6 9 8 5 29 10 43 56 29 57 64 64 63 61 70 67 62 64 65 59 59 57
64 64 64 65 66 27 5 4 4 5 6 6 6 18 66 20 57 60 46 36 75 70 62 61 70 67 62 61 60 59 58 58
49 50 62 65 57 5 5 6 5 6 6 6 6 41 59 28 60 58 44 22 63 71 72 60 69 68 61 60 58 59 59 58
42 52 57 52 26 5 5 5 5 5 5 5 70 50 43 61 62 64 39 42 64 60 62 56 63 65 65 67 61 53 53
32 32 32 33 6 5 5 5 5 5 6 6 11 38 21 33 51 50 45 46 18 32 36 33 23 44 70 71 51 42 27 31
50 50 51 39 5 5 5 5 6 5 6 6 42 69 28 34 42 39 43 37 26 29 40 26 29 26 35 42 35 33 18 19
52 53 51 22 5 5 5 5 6 5 6 5 44 58 17 51 54 53 54 56 51 22 54 54 55 55 54 53 53 52 52
54 54 53 8 5 5 5 5 6 5 6 13 52 42 21 51 54 51 49 49 50 22 41 45 42 42 41 40 41 44 43 42
52 52 54 36 8 5 5 6 5 6 28 55 32 32 54 53 51 51 51 44 25 51 51 49 49 50 49 48 46 46
54 54 52 53 30 7 5 6 6 5 6 40 54 29 52 51 53 56 55 52 52 51 38 52 52 50 49 46 46 45 46 47
51 52 51 53 27 14 5 4 5 4 7 47 51 21 39 49 47 49 52 52 52 49 35 31 48 46 47 47 47 46 46 43
48 50 51 53 25 14 17 8 4 4 17 46 40 18 43 47 46 49 52 54 53 53 54 18 50 49 46 47 47 47 45
49 49 49 49 22 12 20 24 6 14 35 51 39 48 48 50 51 51 49 51 51 52 50 41 58 48 47 47 47 45 45 46
51 49 50 50 22 13 19 36 13 12 42 50 40 73 50 50 50 49 48 49 49 48 49 45 51 46 44 44 44 42 45 47
47 49 49 47 20 16 26 39 21 15 36 48 42 61 47 48 51 47 50 51 51 51 49 47 47 52 47 47 44 43 45 46
48 50 48 52 19 13 33 36 18 18 36 49 51 54 47 47 49 46 46 49 49 49 49 47 44 53 44 48 44 46 46 45

AI in Architecture: What is AI?



AI revealed: Introduction to Artificial Intelligence
from [FAIR \(Facebook AI Research\)](#)

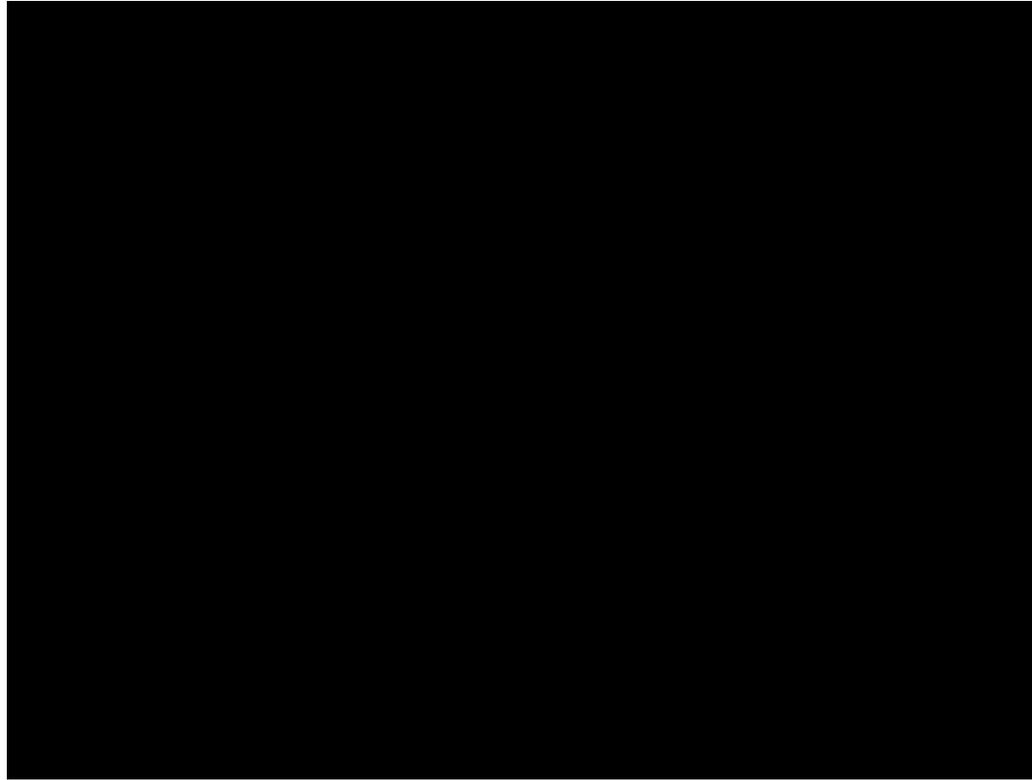
AI in Architecture: What is AI?

Go To:

<https://tech.fb.com/ai-revealed-introduction-to-ai/>

AI revealed: Introduction to Artificial Intelligence
from [FAIR \(Facebook AI Research\)](#)

AI in Architecture: What is AI?



Machine Learning, Explained
from [FAIR \(Facebook AI Research\)](#)

AI in Architecture: What is AI?

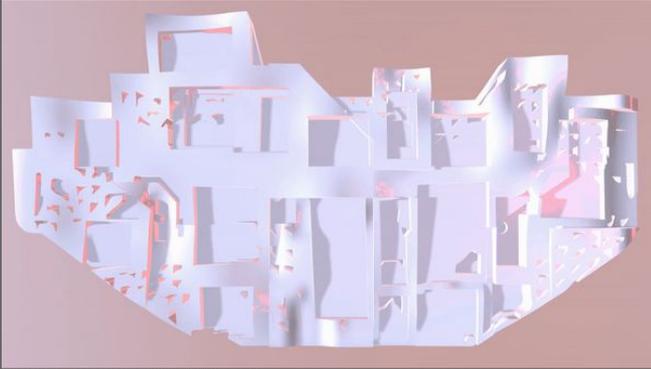
Go To:

<https://tech.fb.com/video/machine-learning-explained>

!

Machine Learning, Explained
from [FAIR \(Facebook AI Research\)](#)

Disciplinary Potential of AI in Architecture



Professional Potential of AI in Architecture

2018
Finch3D, Adaptive House
Jesper Wallgreen

2017
Generative Design for Architecture: Autodesk MaRS Office
Autodesk, The Living



The AI-Art Gold Rush Is Here

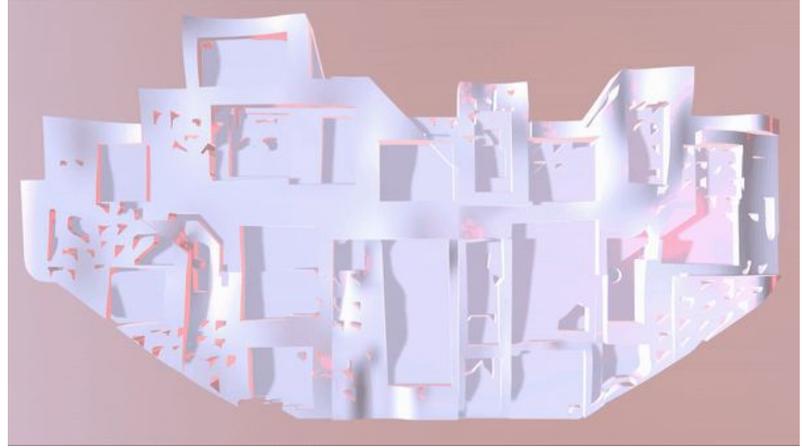
An artificial-intelligence “artist” got a solo show at a Chelsea gallery. Will it reinvent art, or destroy it?



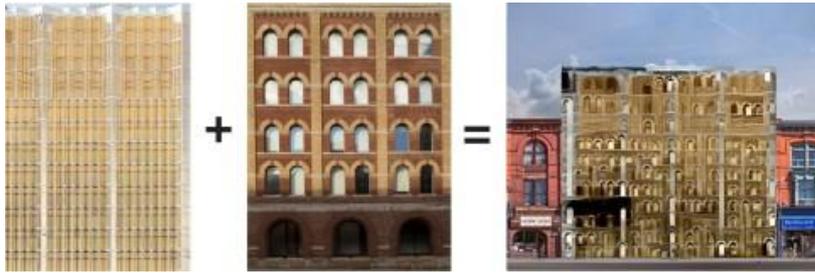
Memories of Passerby, Mario Klingemann (2019)



AI-Supported Portraits, Mario Klingemann (2019)
Mario-Klingemann-supported Portraits, AI (2019) ?



28



29

Delirious Facade by Lamas (2019)



Convergence and Stability of GAN Training by Mescheder at Max Planck Institute (2017-8)

https://avg.is.tuebingen.mpg.de/research_projects/convergence-and-stability-of-gan-training



Latent Walk around the City by Refik Anadol (2019)

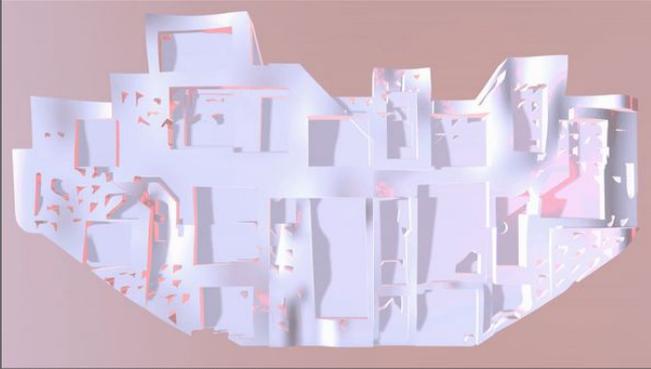
<https://twitter.com/refikanadol/status/1165885055278862337?lang=en>



Latent Walk around the City by Refik Anadol (2019)

<https://twitter.com/refikanadol/status/1165885055278862337?lang=en>

Disciplinary Potential of AI in Architecture



Professional Potential of AI in Architecture

2018
Finch3D, Adaptive House
Jesper Wallgreen

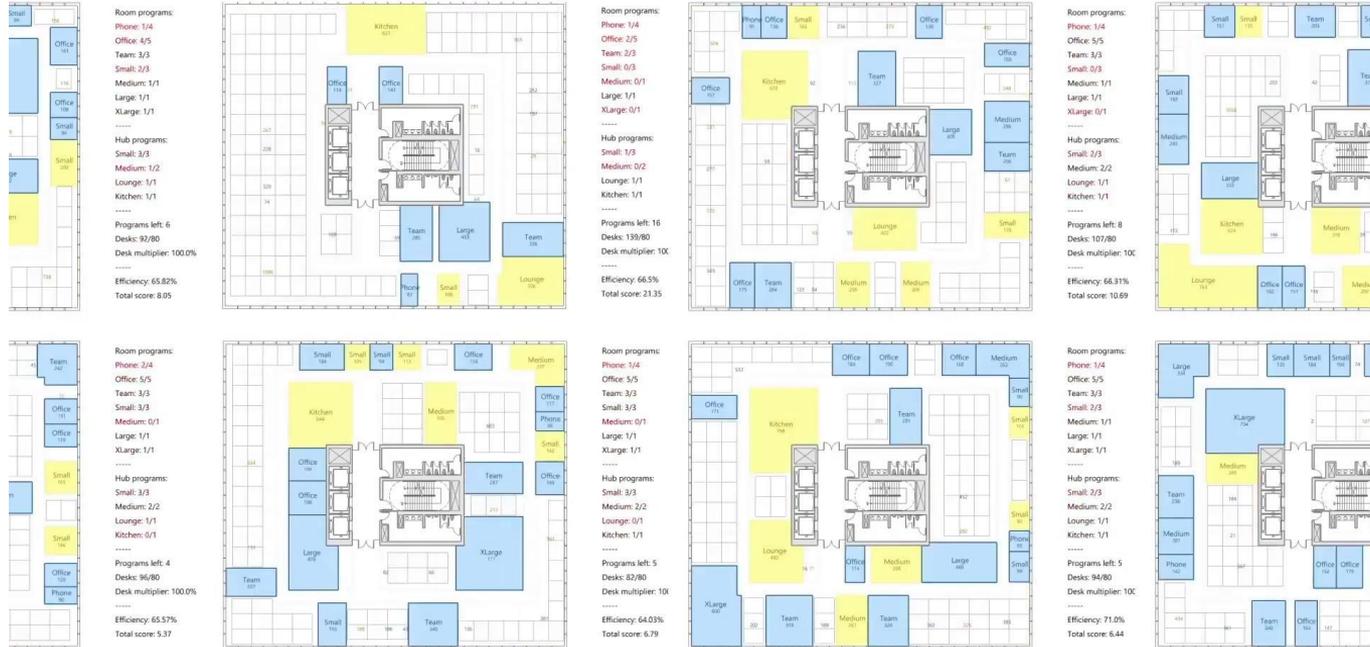
2017
Generative Design for Architecture: Autodesk MaRS Office
Autodesk, The Living

AI in Architecture: History (4) **Artificial Intelligence - Generative Design**

2018
Finch3D, Adaptive House
Jesper Wallgreen

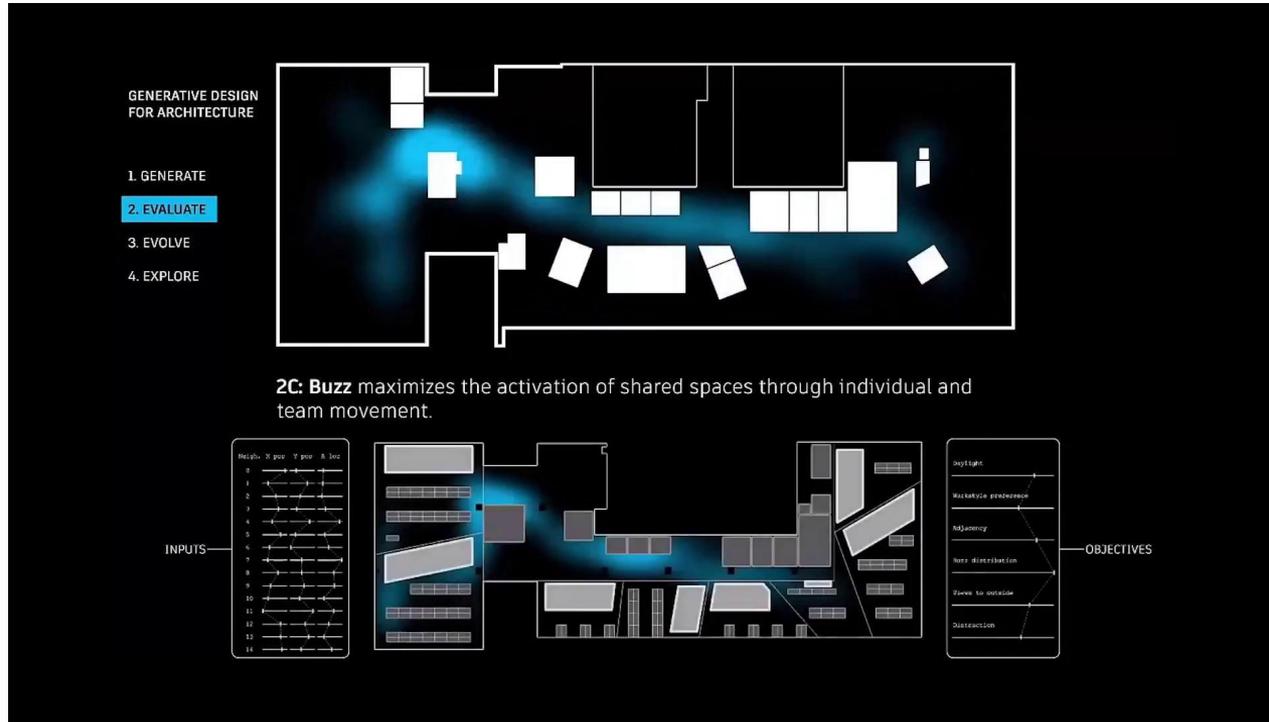
Finch3D (2018)

AI in Architecture: History (4) Artificial Intelligence - Generative Design



Autodesk, What is Generative Design? (2019)

AI in Architecture: History (4) Artificial Intelligence - Generative Design



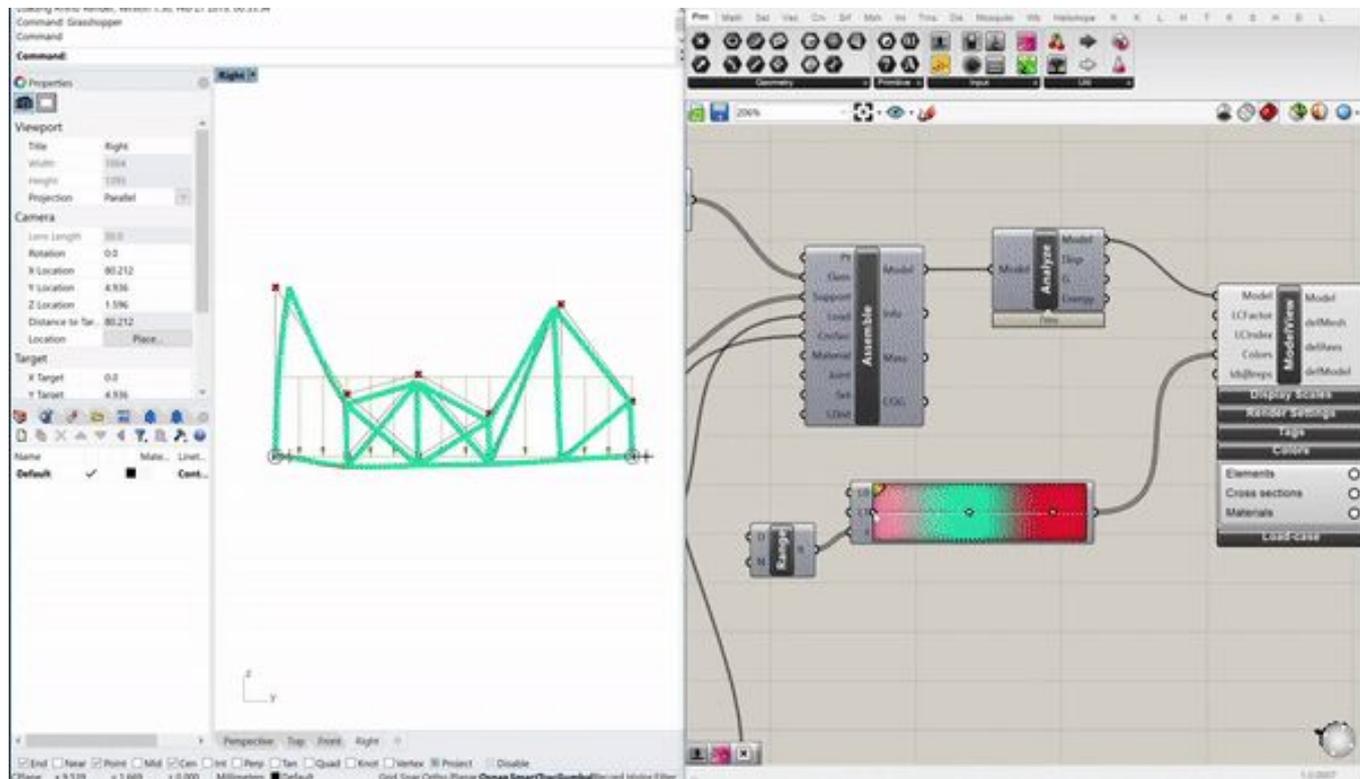
Autodesk, What is Generative Design? (2019)

AI in Architecture: History (4) **Artificial Intelligence - Generative Design**

2019
What is Generative Design ?
Danil Nagy, Autodesk The Living

Autodesk, What is Generative Design? (2019)

AI in Architecture: History (4) Artificial Intelligence - Generative Design



Autodesk, What is Generative Design? (2019)

AI in Architecture: History (4) **Artificial Intelligence - Generative Design**

2019
What is Generative Design ?
Danil Nagy, Autodesk The Living

Autodesk, What is Generative Design? (2019)

AI in Architecture: History (4) **Artificial Intelligence - Generative Design**

2017

Generative Design for Architecture: Autodesk MaRS Office
Autodesk, The Living

Autodesk, What is Generative Design? (2019)

1.1

Introduction

Who are we?

Who are you?

Class Overview

AI in Architecture

AI-DL-ML-NN

1.2

Applications

Deep Learning in Design

Example Project:
Beaux-arts Latent Walk

“Data - Model - Project”

Example Projects

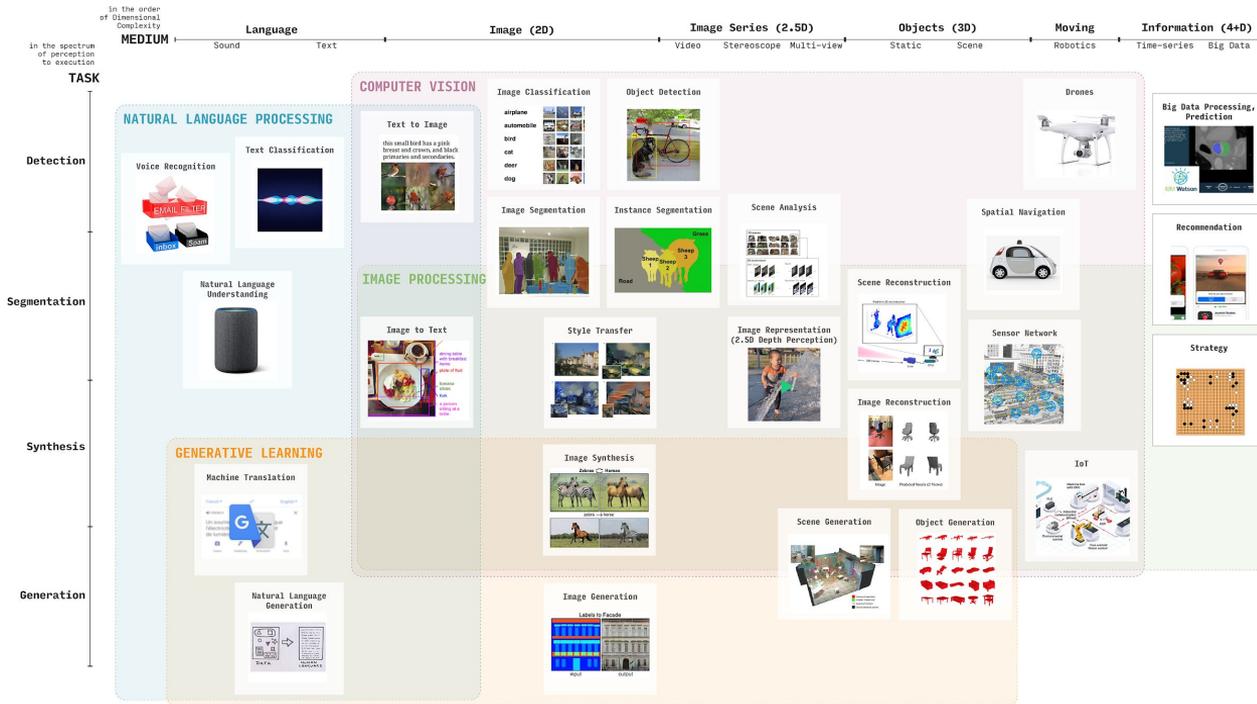
1.3

Workshop

Toolset

Data Scraping

Deep Learning in Design



AI in Architecture, Design & Urban Environment (2020)
 Lab for Design Technologies
 Prof. Andrew Witt, Gia Jung, Claire Djang

Example Projects

“Beaux-Art Latent Space Visualization” (unpublished)
by Andrew Witt, Gia Jung, Claire Djang, Lab for Design Technologies, Harvard University

“Automatic Affinities” <https://automaticaffinities.net> by Gia Jung

“Architecture of the 90%” <https://towardsdatascience.com/ai-architecture-f9d78c6958e0> by Stanislas Chaillou

“Landscape of Emotions” <https://www.evpeng.com/projects/landscape-of-emotions> by Shiyi Peng

Teachable Machine <https://teachablemachine.withgoogle.com/train/image> by Google

Example Projects

“**Beaux-Art Latent Space Visualization**” (unpublished)
by Andrew Witt, Gia Jung, Claire Djang, Lab for Design Technologies, Harvard University

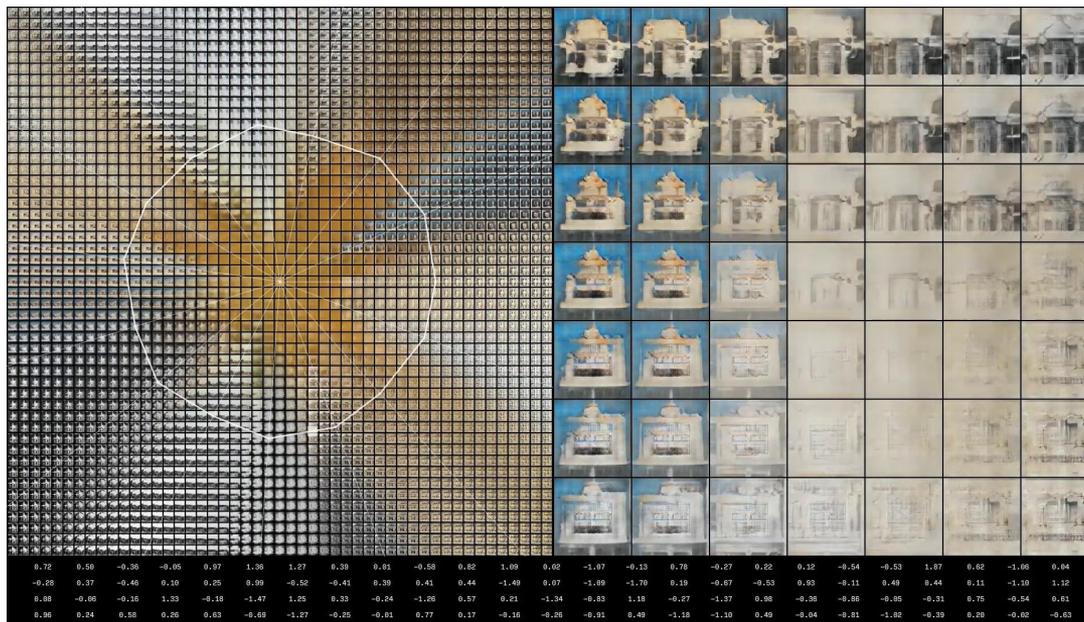
“Automatic Affinities” <https://automaticaffinities.net> by Gia Jung

“Architecture of the 90%” <https://towardsdatascience.com/ai-architecture-f9d78c6958e0> by Stanislas Chaillou

“Landscape of Emotions” <https://www.evpeng.com/projects/landscape-of-emotions> by Shiyi Peng

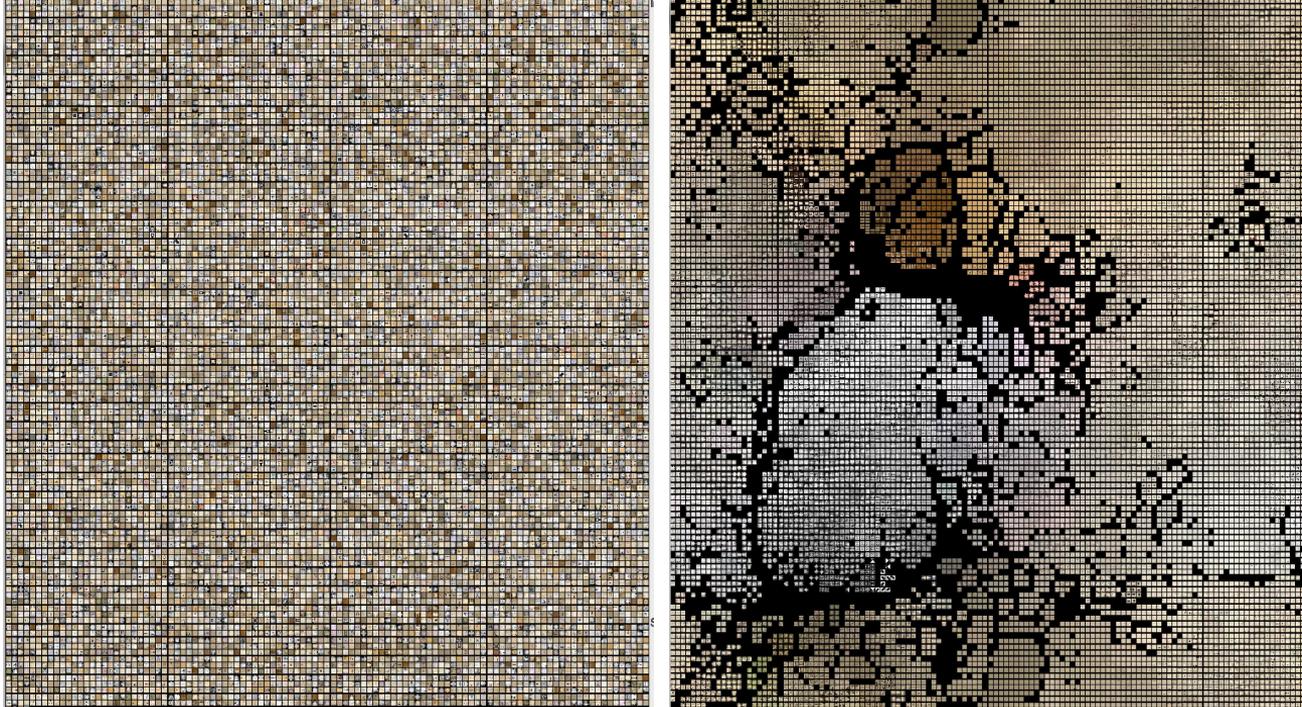
Teachable Machine <https://teachablemachine.withgoogle.com/train/image> by Google

Focus Project: Beaux-art Latent Space Visualization



“Beaux-art Latent Walk” (2020)
Lab for Design Technologies
Prof. Andrew Witt, Gia Jung, Claire Djang

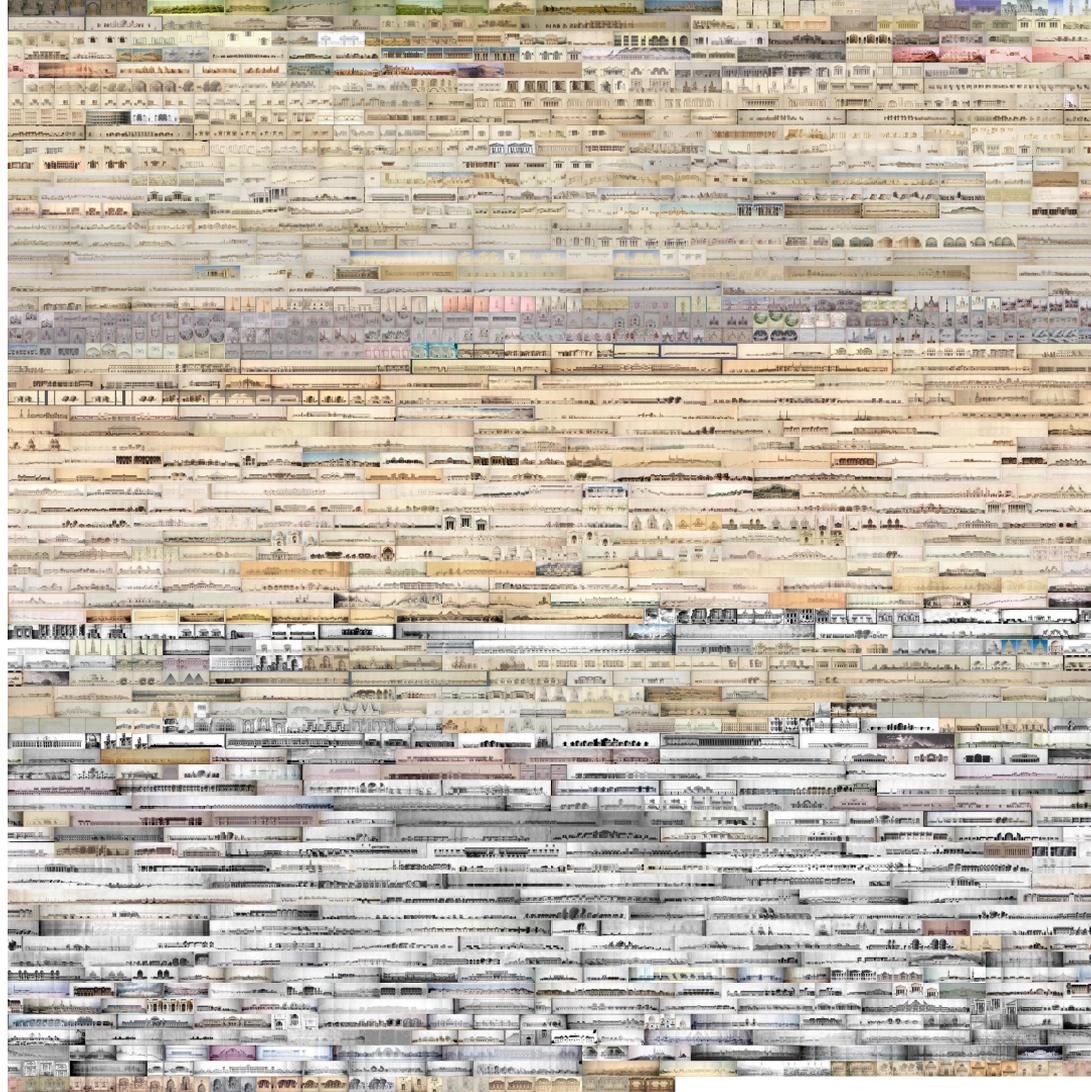
(Data) Data Mining



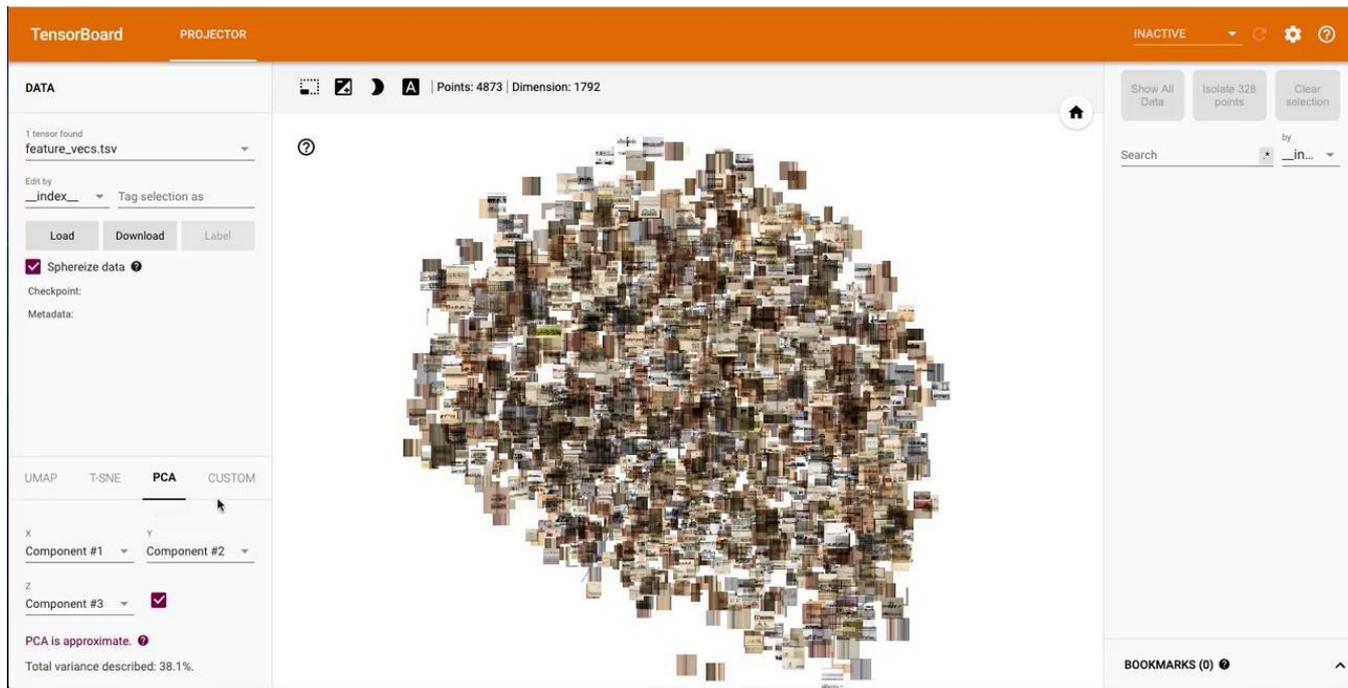
Beaux-Art AI (2020)

Lab for Design Technologies, Harvard University
Prof. Andrew Witt, Gia Jung, Claire Djang



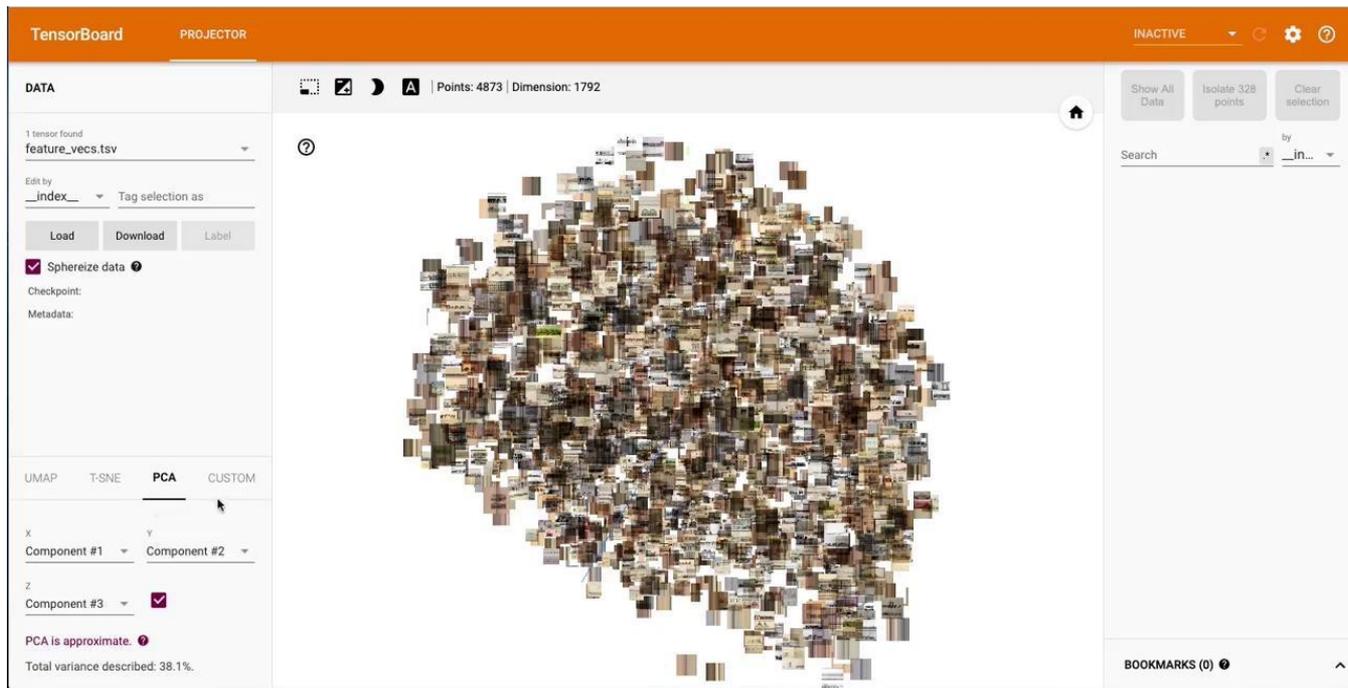


(Data) Mapping Data



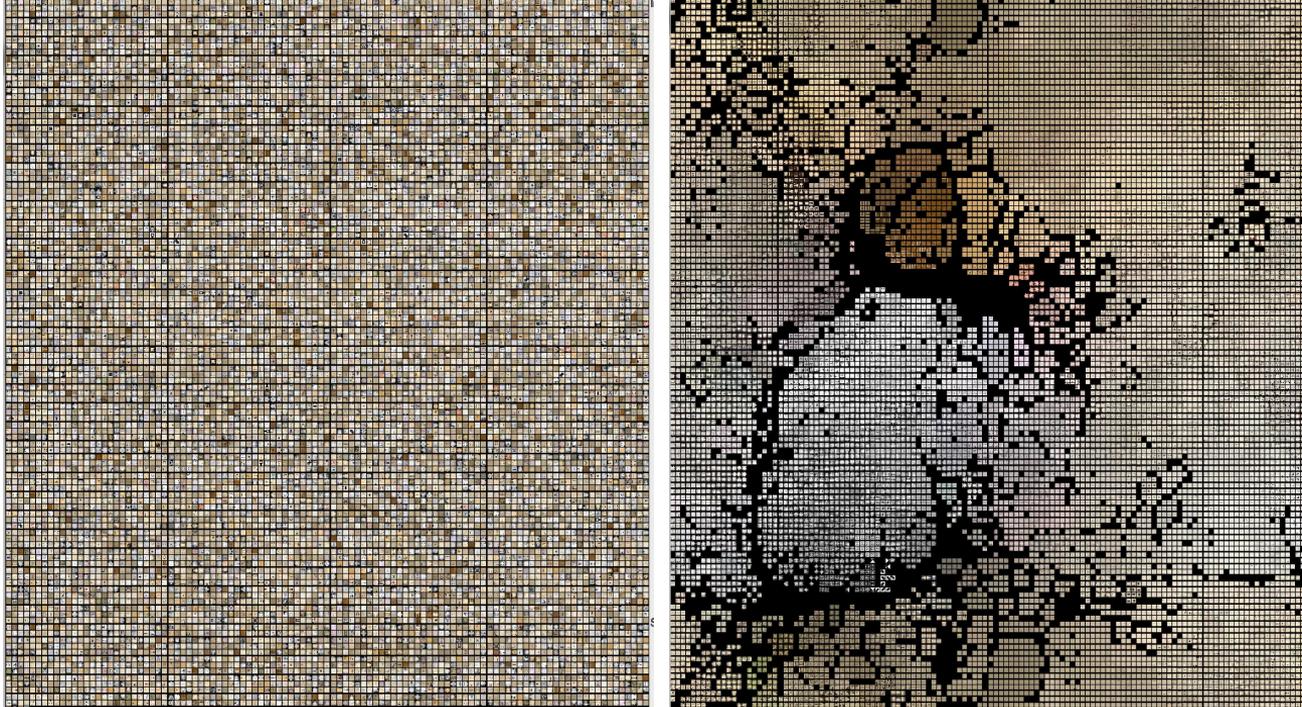
Beaux-Art AI (2020)
Lab for Design Technologies, Harvard University
Prof. Andrew Witt, Gia Jung, Claire Djang

(Data) Mapping Data



Beaux-Art AI (2020)
Lab for Design Technologies, Harvard University
Prof. Andrew Witt, Gia Jung, Claire Djang

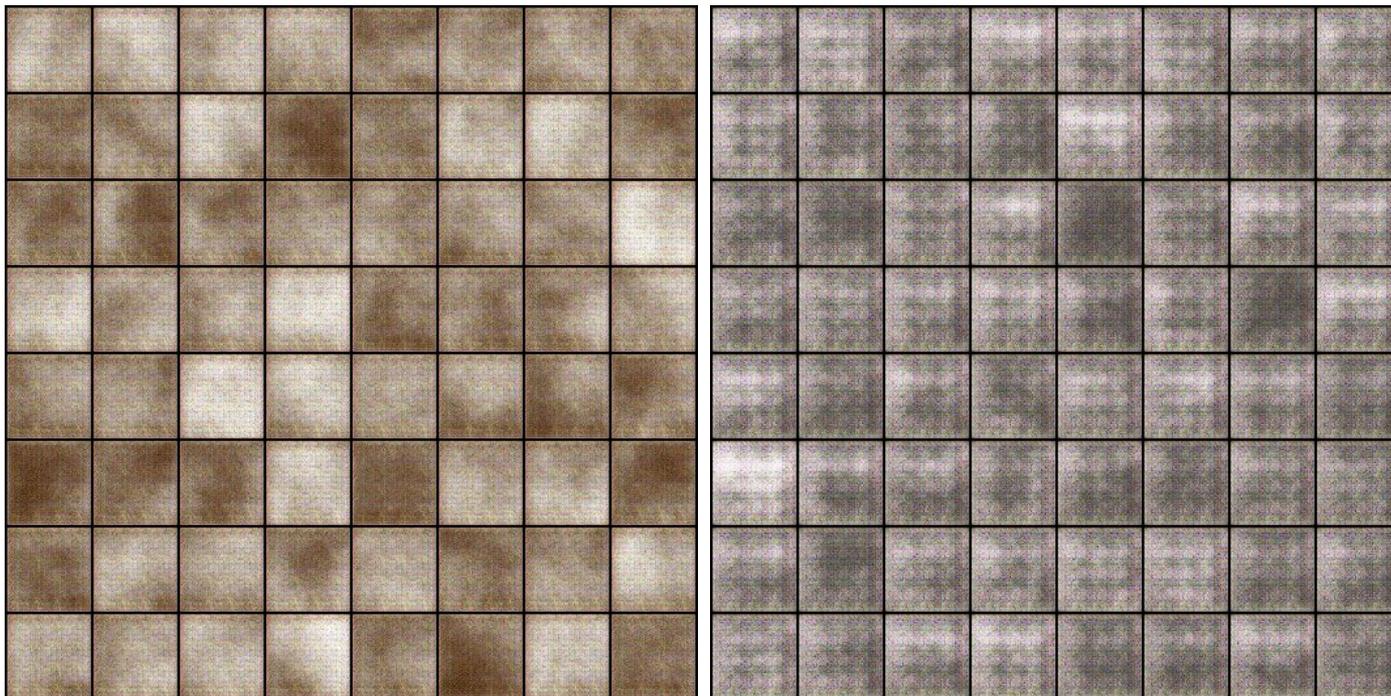
(Data) Data Mining



Beaux-Art AI (2020)

Lab for Design Technologies, Harvard University
Prof. Andrew Witt, Gia Jung, Claire Djang

(Model) Training Neural Net



Beaux-Art AI (2020)

Lab for Design Technologies, Harvard University
Prof. Andrew Witt, Gia Jung, Claire Djang

(Data) Mapping Data



Beaux-Art AI (2020)

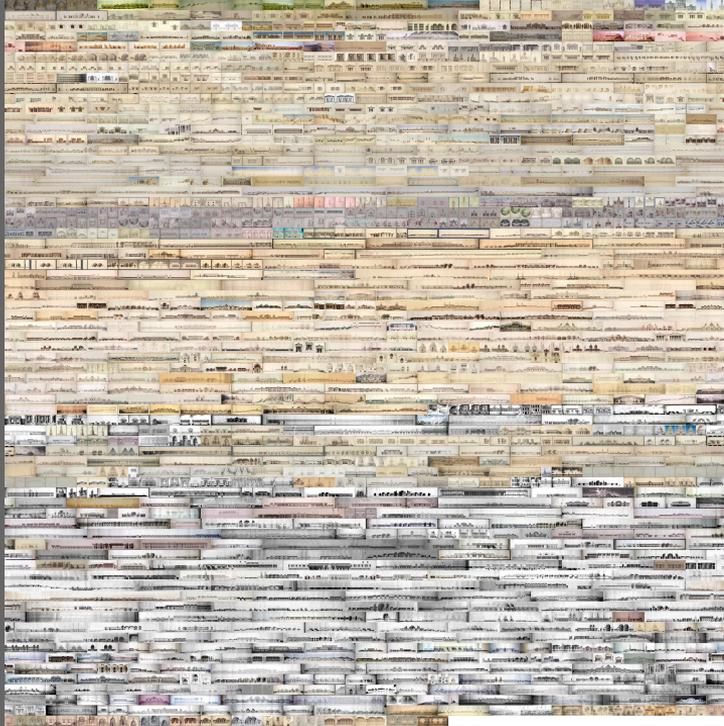
Lab for Design Technologies, Harvard University
Prof. Andrew Witt, Gia Jung, Claire Djang

(Model) Generation & Latent Space

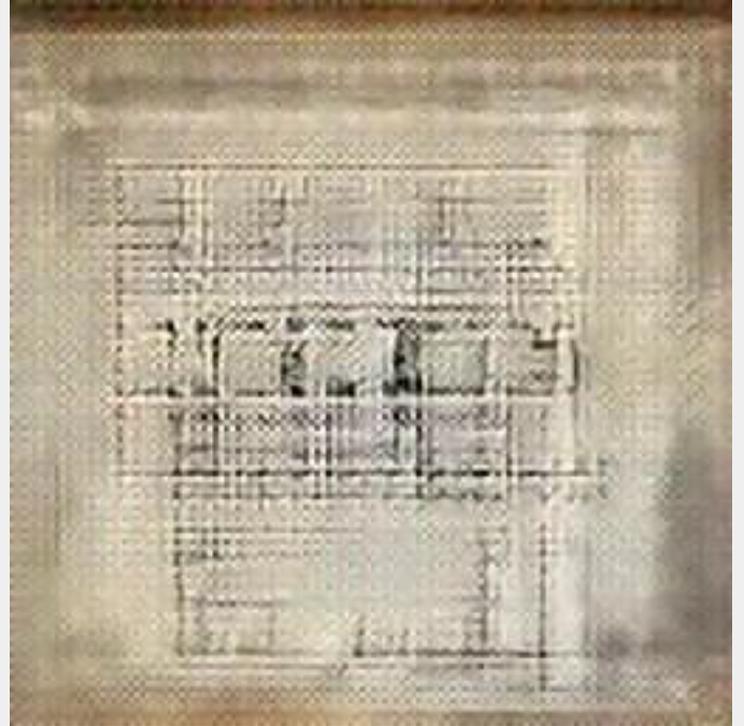


Beaux-Art AI (2020)
 Lab for Design Technologies, Harvard University
 Prof. Andrew Witt, Gia Jung, Claire Djang

Data



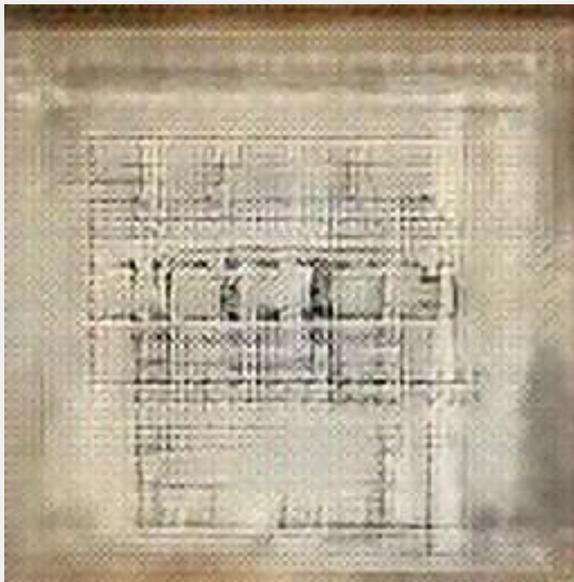
Model



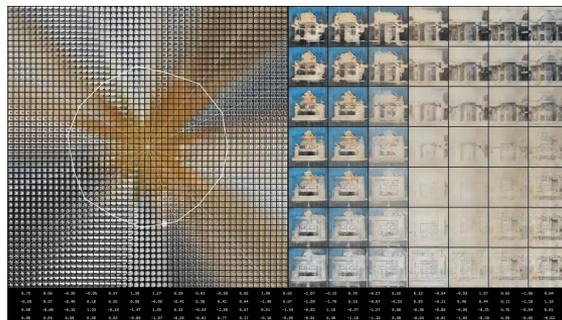
Data



Model



Project



Example Projects

“Beaux-Art Latent Space Visualization” (unpublished)
by Andrew Witt, Gia Jung, Claire Djang, Lab for Design Technologies, Harvard University

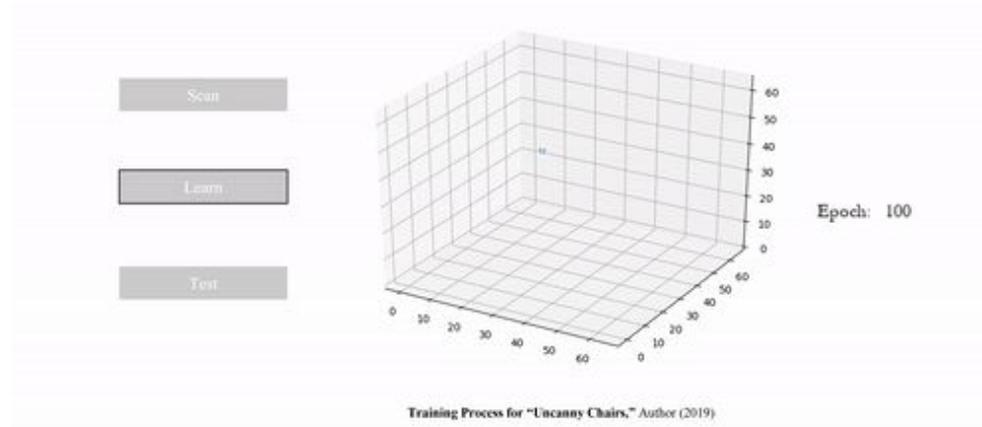
“**Automatic Affinities**” <https://automaticaffinities.net> by Gia Jung

“Architecture of the 90%” <https://towardsdatascience.com/ai-architecture-f9d78c6958e0> by Stanislas Chaillou

“Landscape of Emotions” <https://www.evpeng.com/projects/landscape-of-emotions> by Shiyi Peng

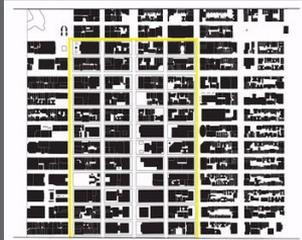
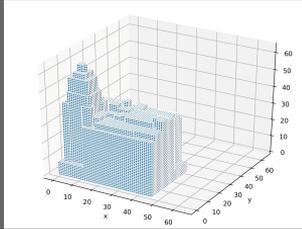
Teachable Machine <https://teachablemachine.withgoogle.com/train/image> by Google

Example Project: “Automatic Affinities”

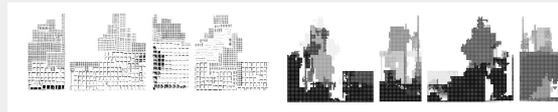
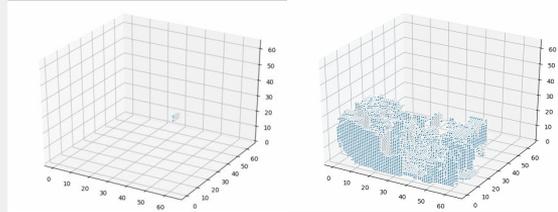


Automatic Affinities, Gia Jung (2020)

Data



Model



Project

- Human-Computer interaction in AI Tools in Design Process
- How to integrate ML in design process: data curation, model curation, presentation
- What new aesthetics of urban figures arise out of trained neural net

Example Projects

“Beaux-Art Latent Space Visualization” (unpublished)
by Andrew Witt, Gia Jung, Claire Djang, Lab for Design Technologies, Harvard University

“Automatic Affinities” <https://automaticaffinities.net> by Gia Jung

“Architecture of the 90%” <https://towardsdatascience.com/ai-architecture-f9d78c6958e0> by Stanislas Chaillou

“Landscape of Emotions” <https://www.evpeng.com/projects/landscape-of-emotions> by Shiyi Peng

Teachable Machine <https://teachablemachine.withgoogle.com/train/image> by Google

AI in Architecture: History (4) Artificial Intelligence - Generative Modeling: GAN



GSD M.Arch I Thesis, Stanislas Chaillou (2019)

Example Projects

“Beaux-Art Latent Space Visualization” (unpublished)
by Andrew Witt, Gia Jung, Claire Djang, Lab for Design Technologies, Harvard University

“Automatic Affinities” <https://automaticaffinities.net> by Gia Jung

“Architecture of the 90%” <https://towardsdatascience.com/ai-architecture-f9d78c6958e0> by Stanislas Chaillou

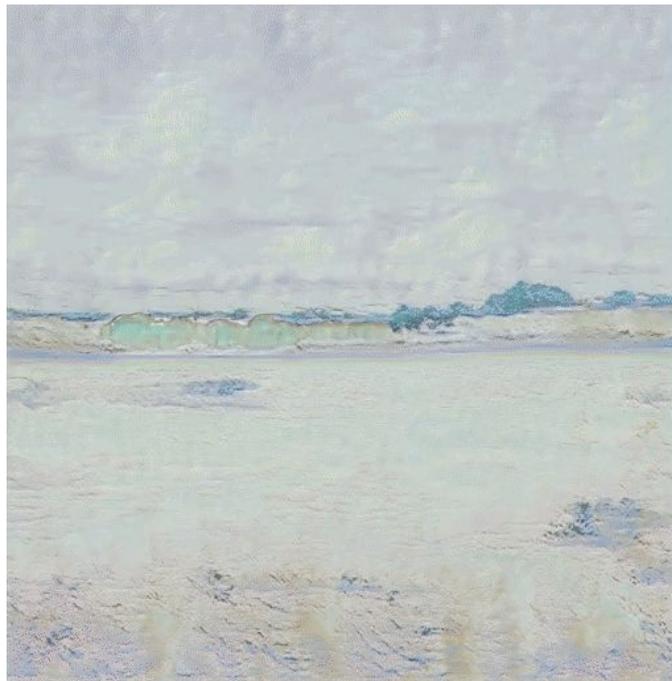
“Landscape of Emotions” <https://www.evpeng.com/projects/landscape-of-emotions> by Shiyi Peng

Teachable Machine <https://teachablemachine.withgoogle.com/train/image> by Google

StyleGAN

Results
Manipulation

Latent space
walk

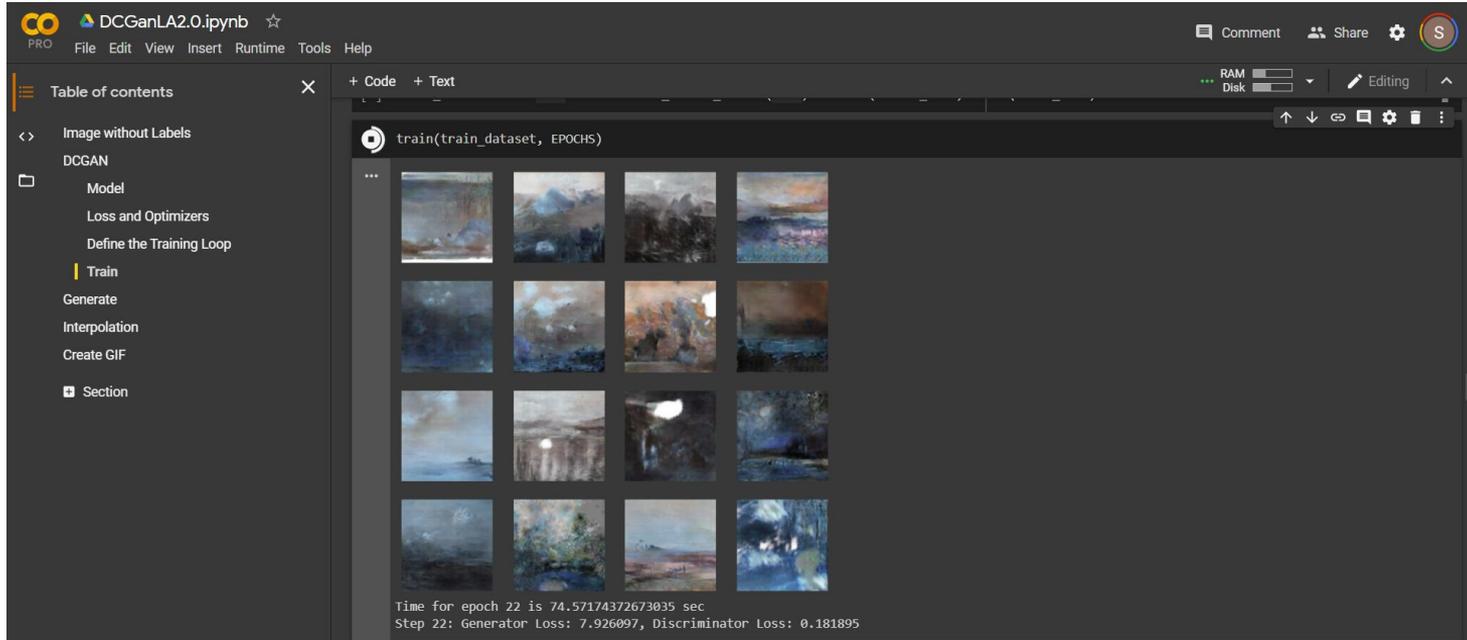


See video if the gif doesn't
load:
https://youtu.be/vzd_boFDpp4

The Landscape of Emotions by Shiyi Peng MDes'20 (2020)

DCGAN

DCGAN landscape art generator is trained using Google Colab Pro with one GPU, over approximately 2 weeks (8 hours per day).



DCGANLA2.0.ipynb

Table of contents

- Image without Labels
- DCGAN
 - Model
 - Loss and Optimizers
 - Define the Training Loop
 - Train
 - Generate
 - Interpolation
 - Create GIF
- Section

```
train(train_dataset, EPOCHS)
```

Time for epoch 22 is 74.57174372673035 sec
Step 22: Generator Loss: 7.926097, Discriminator Loss: 0.181895

The Landscape of Emotions by Shiyi Peng MDes'20 (2020)

Modified architecture to output 128x128x3(rgb) images.

```
def make_generator_model():
    model = tf.keras.Sequential()
    model.add(layers.Dense(8*8*512, use_bias=False, input_shape=(100,)))
    model.add(layers.BatchNormalization())
    model.add(layers.ReLU())

    model.add(layers.Reshape((8, 8, 512)))
    assert model.output_shape == (None, 8, 8, 512)

    model.add(layers.Conv2DTranspose(256, (5, 5), strides=(2, 2), padding='same', use_bias=False))
    assert model.output_shape == (None, 16, 16, 256)
    model.add(layers.BatchNormalization())
    model.add(layers.ReLU())

    model.add(layers.Conv2DTranspose(128, (5, 5), strides=(2, 2), padding='same', use_bias=False))
    assert model.output_shape == (None, 32, 32, 128)
    model.add(layers.BatchNormalization())
    model.add(layers.ReLU())

    model.add(layers.Conv2DTranspose(64, (5, 5), strides=(2, 2), padding='same', use_bias=False))
    assert model.output_shape == (None, 64, 64, 64)
    model.add(layers.BatchNormalization())
    model.add(layers.ReLU())

    model.add(layers.Conv2DTranspose(3, (5, 5), strides=(2, 2), padding='same', use_bias=False, activation='tanh'))
    assert model.output_shape == (None, 128, 128, 3)

    return model
```

DCGan

Selected
Results

Some of the best generation results.



The Landscape of Emotions by Shiyi Peng MDes'20 (2020)

DCGan

Results Manipulation

The latent space of the learned model can be manipulated by vector arithmetics.
When we have the latent vector for an image, adding or subtracting another small vector results in small variations of the same image:

$\text{noise} = z - 0.35 * \text{tf.random.normal}([1, 100])$ smaller variations



$\text{noise} = z + 0.8 * \text{tf.random.normal}([1, 100])$ larger variations

DCGan

Results Manipulation

Continuous interpolation between two vectors becomes a random traversal in the latent space:





Interpolations among small variations of one image.



Random traversal in latent space.

Data

The screenshot shows the Kaggle Data Explorer interface. At the top, there is a search bar and navigation tabs for 'Data', 'Tasks', 'Notebooks (2)', 'Discussion', 'Activity', and 'Metadata'. A 'Download (35 GB)' button is visible. The 'Data Explorer' sidebar on the left shows a tree view of folders including 'abstract', 'animal-painting', 'cityscape', 'figurative', 'flower-painting', 'genre-painting', and 'landscape'. The main area displays the 'landscape' dataset with 15.0k files. A grid of 16 landscape painting thumbnails is shown, each with a unique ID and file size. The thumbnails depict various landscape scenes, including forests, mountains, and coastal views.

Kaggle WikiArt Dataset

Model

The screenshot shows the Kaggle Training interface for a 'Training Image Experiment'. The main area features a large grid of landscape painting thumbnails, similar to the data explorer, with the text 'Training in progress!' and 'See samples of your training progress.' A progress bar at the bottom indicates '0 Step' out of '3000 Steps', with a current position at '1220'. On the right side, there is a 'TRAINING STATUS' section showing 'Started an hour ago', 'Status', 'FID Score' of 51.09, 'Steps' of 1220 / 3000, and 'ETA Approx.' of 'an hour remaining'. Buttons for 'Save Sample Image' and 'Create Progress Video' are located at the bottom.

The Landscape of Emotions by Shiyi Peng MDes'20 (2020)

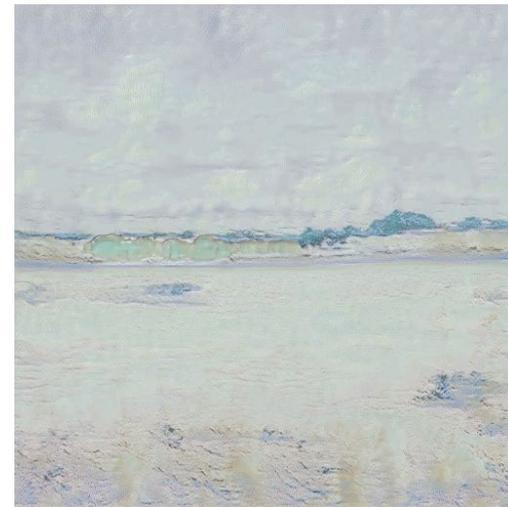
Data

A screenshot of a web-based Data Explorer interface. At the top, there is a search bar and navigation tabs for 'Data', 'Tasks', 'Notebooks', 'Discussion', 'Activity', and 'Metadata'. A 'Download (35 GB)' button is visible. The main content area is titled 'Data Explorer' and shows a folder named 'landscape (15.0k files)' containing 34.98 GB of data. A sidebar on the left lists various categories like 'abstract', 'animal-painting', 'cityscape', etc., with 'landscape' selected. The main area displays a grid of 20 landscape image thumbnails, each with a unique ID and file size. The thumbnails show various natural scenes like mountains, lakes, and forests.

Model

A screenshot of a 'Training Image Experiment' interface. The title 'Training in progress!' is prominently displayed in the center, with a sub-header 'See samples of your training progress.' Below this is a grid of 20 small image thumbnails, similar to the ones in the Data Explorer, showing the model's output at different stages. At the bottom, there is a progress bar ranging from '0 Step' to '3000 Steps', with a blue dot indicating the current position at '1220'. Two buttons, 'Save Sample Image' and 'Create Progress Video', are located below the progress bar.

So-what



Training process from 1000 to 3000 steps:

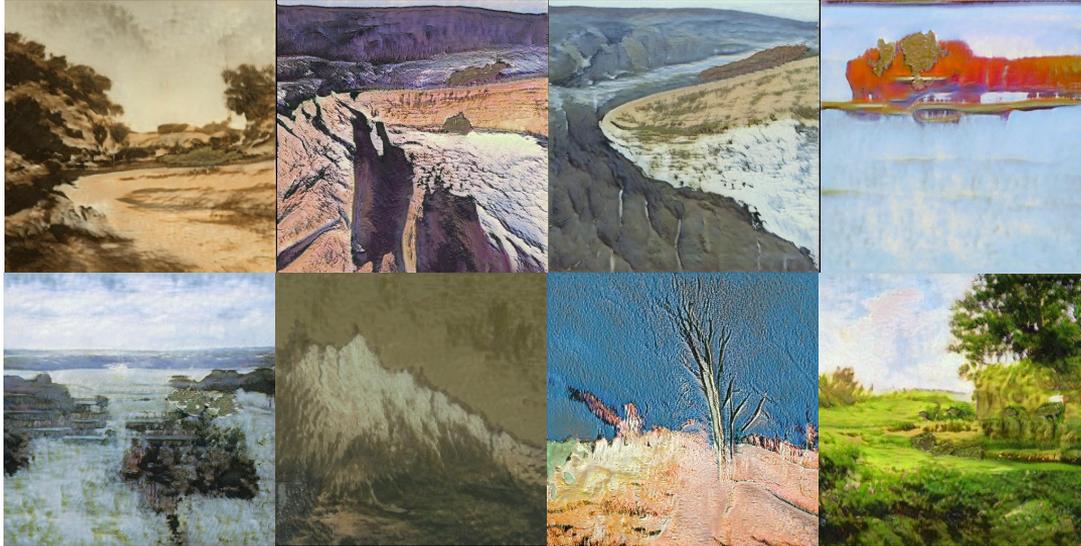


<https://youtu.be/zUR3gM6EqNo>

StyleGAN

Results
Evaluation

The outputs are 512x512 px. More training steps may have better

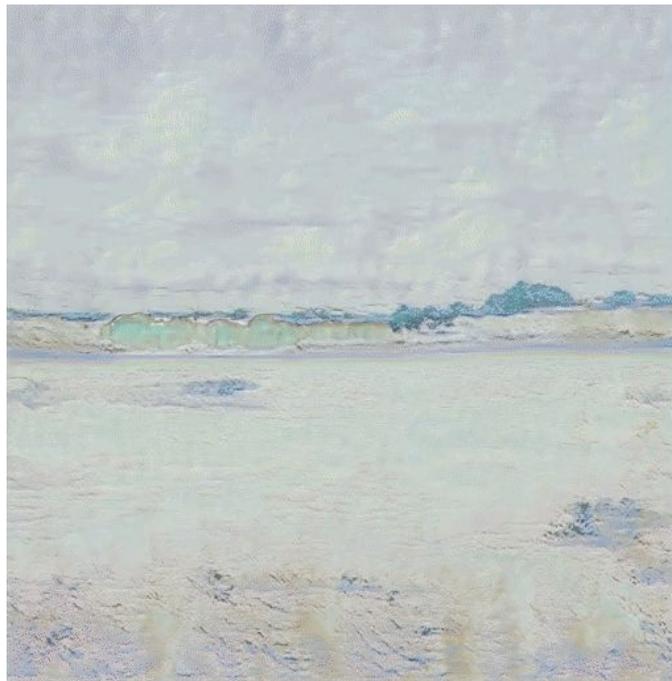


The Landscape of Emotions by Shiyi Peng MDes'20 (2020)

StyleGAN

Results
Manipulation

Latent space
walk



See video if the gif doesn't
load:
https://youtu.be/vzd_boFDpp4

The Landscape of Emotions by Shiyi Peng MDes'20 (2020)

Data

The screenshot shows a web-based Data Explorer interface. At the top, there is a search bar and navigation tabs for 'Data', 'Tasks', 'Notebooks (2)', 'Discussion', 'Activity', and 'Metadata'. A 'Download (35 GB)' button is visible. The main content area is divided into two sections. The top section is titled 'landscape (15.0k files)' and displays a grid of 12 landscape images with their respective file names and sizes. The bottom section is titled 'abstract (15.0k files)' and displays a grid of 12 abstract images with their respective file names and sizes. The interface includes a sidebar on the left with a 'Data Explorer' section and a search bar at the bottom of each section.

Model

The screenshot shows a Jupyter Notebook interface for a DCGAN model. The notebook title is 'DCGANLA2.0.ipynb'. The 'Table of contents' on the left lists sections: 'Image without Labels', 'DCGAN', 'Model', 'Loss and Optimizers', 'Define the Training Loop', 'Train', 'Generate', 'Interpolation', 'Create GIF', and 'Section'. The main code cell is titled 'train(train_dataset, EPOCHS)' and contains a grid of generated images. Below the code cell, a progress bar indicates 'Training in progress!' with a progress indicator at 1220 out of 3000 steps. A 'Save Sample Image' button and a 'Create Progress Video' button are visible at the bottom.

So-what



Example Projects

“Beaux-Art Latent Space Visualization” (unpublished)
by Andrew Witt, Gia Jung, Claire Djang, Lab for Design Technologies, Harvard University

“Automatic Affinities” <https://automaticaffinities.net> by Gia Jung

“Architecture of the 90%” <https://towardsdatascience.com/ai-architecture-f9d78c6958e0> by Stanislas Chaillou

“Landscape of Emotions” <https://www.evpeng.com/projects/landscape-of-emotions> by Shiyi Peng

Teachable Machine <https://teachablemachine.withgoogle.com/train/image> by Google

1.1

Introduction

Who are we?

Who are you?

Class Overview

AI in Architecture

AI-DL-ML-NN

Class Organization

1.2

Applications

Deep Learning in Design

“Data - Model - Project”

Examples

Example Project:

Beaux-arts Latent Walk

1.3

Workshop

Toolset

Data Scraping