Creative A Generative Art & Al-Assisted Creativity

Eric Chu 12/10/2021 LINGO Group Presentation





Outline

- 1. My own path and projects
 - Developments in the machine learning world
- 2. Creative AI
 - Generative art
 - Al-assisted creativity
- 3. Connections to machine learning research
- 4. Connections to the real world...

My path: 2014, topological sculptures

algebraic topology





biomedical imaging & computational photography & computer vision





My path: 2014, topological sculptures

Eva Hild sculptures







DOI:10.1145/2001269.2001293

Building Rome in a Day

By Sameer Agarwal^a, Yasutaka Furukawa^a, Noah Snavely, Ian Simon^b, Brian Curless, Steven M. Seitz, and Richard Szeliski

Abstract

from large, unorganized collections of photographs such | tions two to three orders of magnitude larger. Fourth, the as those found by searching for a given city (e.g., Rome) on | algorithms must be *fast*—we seek to reconstruct an entire **Internet photo-sharing sites.** Our system is built on a set | city in a single day, making it possible to repeat the process of new, distributed computer vision algorithms for image many times to reconstruct all of the world's significant culmatching and 3D reconstruction, designed to maximize tural centers. parallelism at each stage of the pipeline and to scale gracefully with bot

available con strate that it i collections w less than a da

enormous-whereas prior methods operated on hundreds We present a system that can reconstruct 3D geometry or at most a few thousand photos, we seek to handle collec-

Creating accurate 3D models of cities is a problem of





"Can we do few-shot 3D reconstruction of these sculptures?"

"No. Too hard. Build me this tool."



My path: 2014, topological sculptures



A CAD tool for modeling 3D-printable topological sculptures

			🎏 Display for Window: Window	
tun. mid-radius 2.0	funnel-radius 6.0	ribbon topX 3.0	Options c s	-SLID
1.0 9.0	1.0 9.0	1.0 10.0	Update Time = 0.000000 Render Time = 0.000000	% slf_gRoot
tun. out-radius 3.0	funnel-alpha 9.0	ribbon topZ 8.0		C SLF_OFF
				C SLF BOUND
1.0 9.0	1.0 9.0	1.0 15.0		
tunnel dist/len 1.0	funnel-theta 0.0	ribbon bckY 10.0		• SLF_FULL
0.1 4.0	90.0 -90.0	1.0 15.0		
tunnel-alpha 90	funnel-phi	ribbon bckZ		74 ai 🗕 🗌
-180 180	180.0 -180.0	-9.0 10.0		
tunnel-theta	funnel x-pos	azimuth		
-90				
-120 120	-4.0 4.0	-90 90		Error Met
tunnel-phi	funnel y-pos	twist		ERROR_METRIC_E
	3.0			Subdivisions
-180 180	-4.0 6.0	-90 90		2
tunnel x-pos 0.0	funnel z-pos -8.0	sweepSlices 24		
			🏸 uiOffset 📃 🗖 🗙	Tolerance
-4.0 4.0	-8.0 0.0	8 80	LOD Shading Offset Draw Controls	
tunnel y-pos	drawSweep	drawSweep	SLF_OFF SLF_INHERIT SLF_OFFSET_ANGULAR SLF_OFF	0.0001
			Height	Facets Max
-4.0 6.0	0 1	0 1	0.00	0
tunnel z-pos	sweepSlices	drawPath		
4.0	24	1	-5.00 5.00	
-4.0 8.0			Strut Width 0.00	Weighted Edges
drawSween	funnel width	ribbon width		
0	2.0	1.0	-5.00 5.00	0
0 1	0.1 4.0	0.1 4.0		



Sep 2015

2

Leon A. Gatys,^{1,2,3*} Alexander S. Ecker,^{1,2,4,5} Matthias Bethge^{1,2,4} ¹Werner Reichardt Centre for Integrative Neuroscience

A Neural Algorithm of Artistic Style

and Institute of Theoretical Physics, University of Tübingen, Germany ²Bernstein Center for Computational Neuroscience, Tübingen, Germany ³Graduate School for Neural Information Processing, Tübingen, Germany ⁴Max Planck Institute for Biological Cybernetics, Tübingen, Germany

> [^]Medicine, Houston, TX, USA mail: leon.gatys@bethgelab.org

> tered the skill to create unique ex interplay between the conithmic basis of this process is ith similar capabilities. Howas object and face recognition rated by a class of biologically













1. Semantics







2. Emotional

3. Knowledge Grounded

"Roses" - Outkast The Roses of Heliogabalus, Sir Lawrence Alma-Tadema

WikiArt dataset

- 121,405 paintings
- 2,539 artists

Match based on

content_similarity(painting_title, lyric)
emotion_similarity(painting_title, lyric)
emotion_similarity(painting, lyric)

My Path: 2016, FlyArt

OH YOU MAD

Mario Tozzi Autoritratto

My Path: 2017, Artistic Influence GANs

- What if Banksy had met Jackson Pollock during his formative years?
- What if David Hockney had missed out on the Tate Gallery's famous **1960 Picasso exhibition?**

WikiArt dataset

- 121,405 paintings
- 2,539 artists
- Influence graph!

Can we develop a GAN to answer "what if" questions?

Step 1: compute painting embeddings

Step 2: compute artist embeddings

Gaussian mixture model over painting embeddings (painters often have multiple artistic periods)

Step 3: use artist embeddings to train Artistic Influence GAN

Step 4: at inference, specify influencers, generate art

My Path: 2017, Artistic Influence GANs

Single influencer

Morris Louis

William Turner

Two influencers

Paul Klee, William Turner Judy Chicago, Thomas Cole

Paul Klee

Three influencers

Caravaggio, de Kooning, Raphael, Picasso

Text prompt: "A colourful cubist painting of a parrot in a cage". VQGAN+CLIP made with NightCafe Creator.

cyberpunk forest

How well do CLIP-guided 2D generations capture 3D objects semantics?

to roll to curve to crease to lift to fold to inlay to store to impless to brind to fore to shorten to food to twist to smear to dapple to rotate to crimple to support to tear to support to tear to support to to split to support to cut	to scatter to enrange to repair to discard to enclose to encircle to hide to covar to discard to discard to discard to encircle to hide to covar to discard to discard to discard to discard to encircle to discard to discard to hide to covar
to simplify of entropy to differ of nature to differ of nature to disarrange of layering to open of layering to mix of fetting to splash to grasp to knot to tighten to spill to limite to droop to heap to flow to gather	to whave to systematize to join to systematize to match to refer to laminate to force to lond of mapping to hinge glocation to mark of context to grand of time to delute of carbonization to light to continue

"To Lift" - Richard Serra

Experiments

"to lift a sheet of vulcanized rubber" - VQGAN + CLIP

"to crumple metal" - VQGAN + CLIP

Can we generate CLIP-guided 2D art by first generating the 3D object?

Anamorphic art

DOUGLAS R. HOFSTADTER A metaphorical fugue on minds and machines in the spirit of Lewis Carroll

- 1. Generate 3D object and 2D view
- 2. Score using ImageNet model or CLIP
- 3. Repeat (optimization loop)

Optimization loop with genetic algorithm

Superformula: powerful, parameter-efficient generator

$$r(\varphi) = \left(\left| \frac{\cos\left(\frac{m\varphi}{4}\right)}{a} \right|^{n_2} + \left| \frac{\sin\left(\frac{m\varphi}{4}\right)}{b} \right|^{n_3} \right)^{-\frac{1}{n_1}}$$

broccoli

Petri dish

explosion

elephant's head

record player

elephant's head

dragonfly

hammerhead shark

"Thinking about designing not the object—but a process to generate objects."

<u>History</u>

- 1700's: Musikalisches Würfelspiel ("musical dice game")
- 1960's —: computer-aided generative art
- 1960's, 1990's : generative music
- 2001: "Processing" language Casey Reas, Ben Fry
 - p5.js Lauren McCarthy

uadrilatères

Sol LeWitt (check out Mass MOCA)

Six-part drawing. The wall is divided horizontally and vertically into six equal parts. 1st part: On red, blue horizontal parallel lines, and in the center, a circle within which are yellow vertical parallel lines; 2nd part: On yellow, red horizontal parallel lines, and in the center, a square within which are blue vertical parallel lines; 3rd part: On blue, yellow horizontal parallel lines, and in the center, a triangle within which are red vertical parallel lines; 4th part: On red, yellow horizontal parallel lines, and in the center, a triangle within which are red vertical parallel lines; 4th part: On red, yellow horizontal parallel lines, and in the center, a rectangle within which are blue vertical parallel lines; 5th part: On yellow, blue horizontal parallel lines, and in the center, a trapezoid within which are red vertical parallel lines; 6th part: On blue, red horizontal parallel lines, and in the center, a parallelogram within which are yellow vertical parallel lines. The horizontal lines do not enter the figures.

Conway's Game of Life

•

Growing Neural Cellular Automata

· · · ·

· · ·

De.

Differentiable Model of Morphogenesis

'Unseen'

-	-	-	1	
			2	

.

•

Generative Art: visual generations

The Treachery of Images, René Magritte

The Treachery of ImageNet, Tom White (2017)

Fernanda Viégas and Martin Wattenberg, Live Wind Map, 2012

"Although we made the wind map as an artistic exploration, we've been surprised by the kinds of things people use it for: bird watchers have tracked migration patterns; bicyclists have planned their trips; and we've even seen conspiracy theorists use it to track mysterious chemicals in the air."

Al-assisted Creativity: ideation and assistive tools

Figure 1: The Wordcraft editor. Left: The assistant offers to continue the text at the user's cursor. Middle: After pressing "get continuations" the user is shown several continuations to choose from. Right: The assistant offers to rewrite or elaborate upon the selected text. Demo: http://bit.ly/wordcraft_video.

- Controllability (along human dimensions)
- Safety
- Creativity and intrinsic motivation / information theory
- Novelty and provenance of generations
- Emergent behaviors in multi-agent reinforcement learning
- Procedural generation

"I've made it to the forest." You think as a drop of water falls on your nose. You look up and see dark clouds filling the sky. "Uh oh."

★ You make a temporary forest shelter.

You decide to make a temporary shelter out of branches and leaves. The rain pours on your new temporary shelter, but you are dry. You are awakened by the crackling of the fire. You look up and see a man standing by it.

Formal Theory of Creativity, Fun, and Intrinsic Motivation (1990-2010)

Jürgen Schmidhuber

noun.

pranella

- a thin triangular piece of toast, usually fried, served as part of a meal "a sandwich of guacamole and pranella"
- 2. a word that does not exist; it was invented, defined and used by a machine learning algorithm.

Link / New word / Write your own

g (RL) to maximize or creation of *novel attern*, and *novelty*. regularity if it is tively short program predicting some of ise is unpredictable bserver, a pattern is *sing* if the observer able to *learn* it. The isely *measured* and arate RL controller nce the controller is rising data.

ment this idea. They tists or artists with a of the world and their models, they

AI Art at Christie's Sells for \$432,500

f 🔉 🖌 🛤 🏘 🗍

"Edmond de Belamy, from La Famille de Belamy," by the French art collective Obvious, was sold on Thursday at Christie's New York. Christie's

Robbie Barrat @videodrome

left: the "AI generated" portrait Christie's is auctioning off right now

right: outputs from a neural network I trained and put online *over a year ago*.

Does anyone else care about this? Am I crazy for thinking that they really just used my network and are selling the results?

10:31 PM · Oct 24, 2018 · Twitter Web Client

1,478 Retweets 433 Quote Tweets 4,066 Likes

Check out the NeurIPS Machine Learning for Creativity and Design Workshop

https://neuripscreativityworkshop.github.io/2021/

Machine Learning for **Creativity and Design**

What's the State of the 'Art' in Machine Learning? The Creativity workshop explores applying the latest ML technologies in art and design.

December 13, Online.

Elhoseiny, Mohamed and Jha, Divyansh. The Corona Monster. 2020. Al Art Gallery, Online.

Time	Event		
11:15	Welcome and Introduction Presented by Mattie Tesfaldet; Zoom		
11:30	Poster Session 1 All Posters; Discord		
12:30	Computers, Creativity, and Lovelace Speaker Presentation by Mark Riedl; Zoom		
13:00	AI for Augmenting Human Creativity Speaker Presentation by Devi Parikh; Zoom		
13:30	Interspecies Intelligence in Pharmako-AI Speaker Presentation by Kenric Allado-McDowell; Zoom		
14:00	Art Show Zoom		
14:30	Q&A Panel Discussion 1 Mark Riedl, Devi Parikh, and Kenric Allado-McDowell, moderated by Mattie Tesfaldet; Zoom + Rocketchat		
15:00	Social 1 Discord		
15:30	StyleCLIPDraw: Coupling Content and Style in Text-to-Drawing Synthesis Paper Oral by Peter Schaldenbrand et al.; Zoom		
15:40	Soundify: Matching Sound Effects to Video Paper Oral by David Chuan-En Lin et al.; Zoom		
15:50	Controllable and Interpretable Singing Voice Decomposition via Assem-VC Paper Oral by Kang-wook Kim et al.; Zoom		