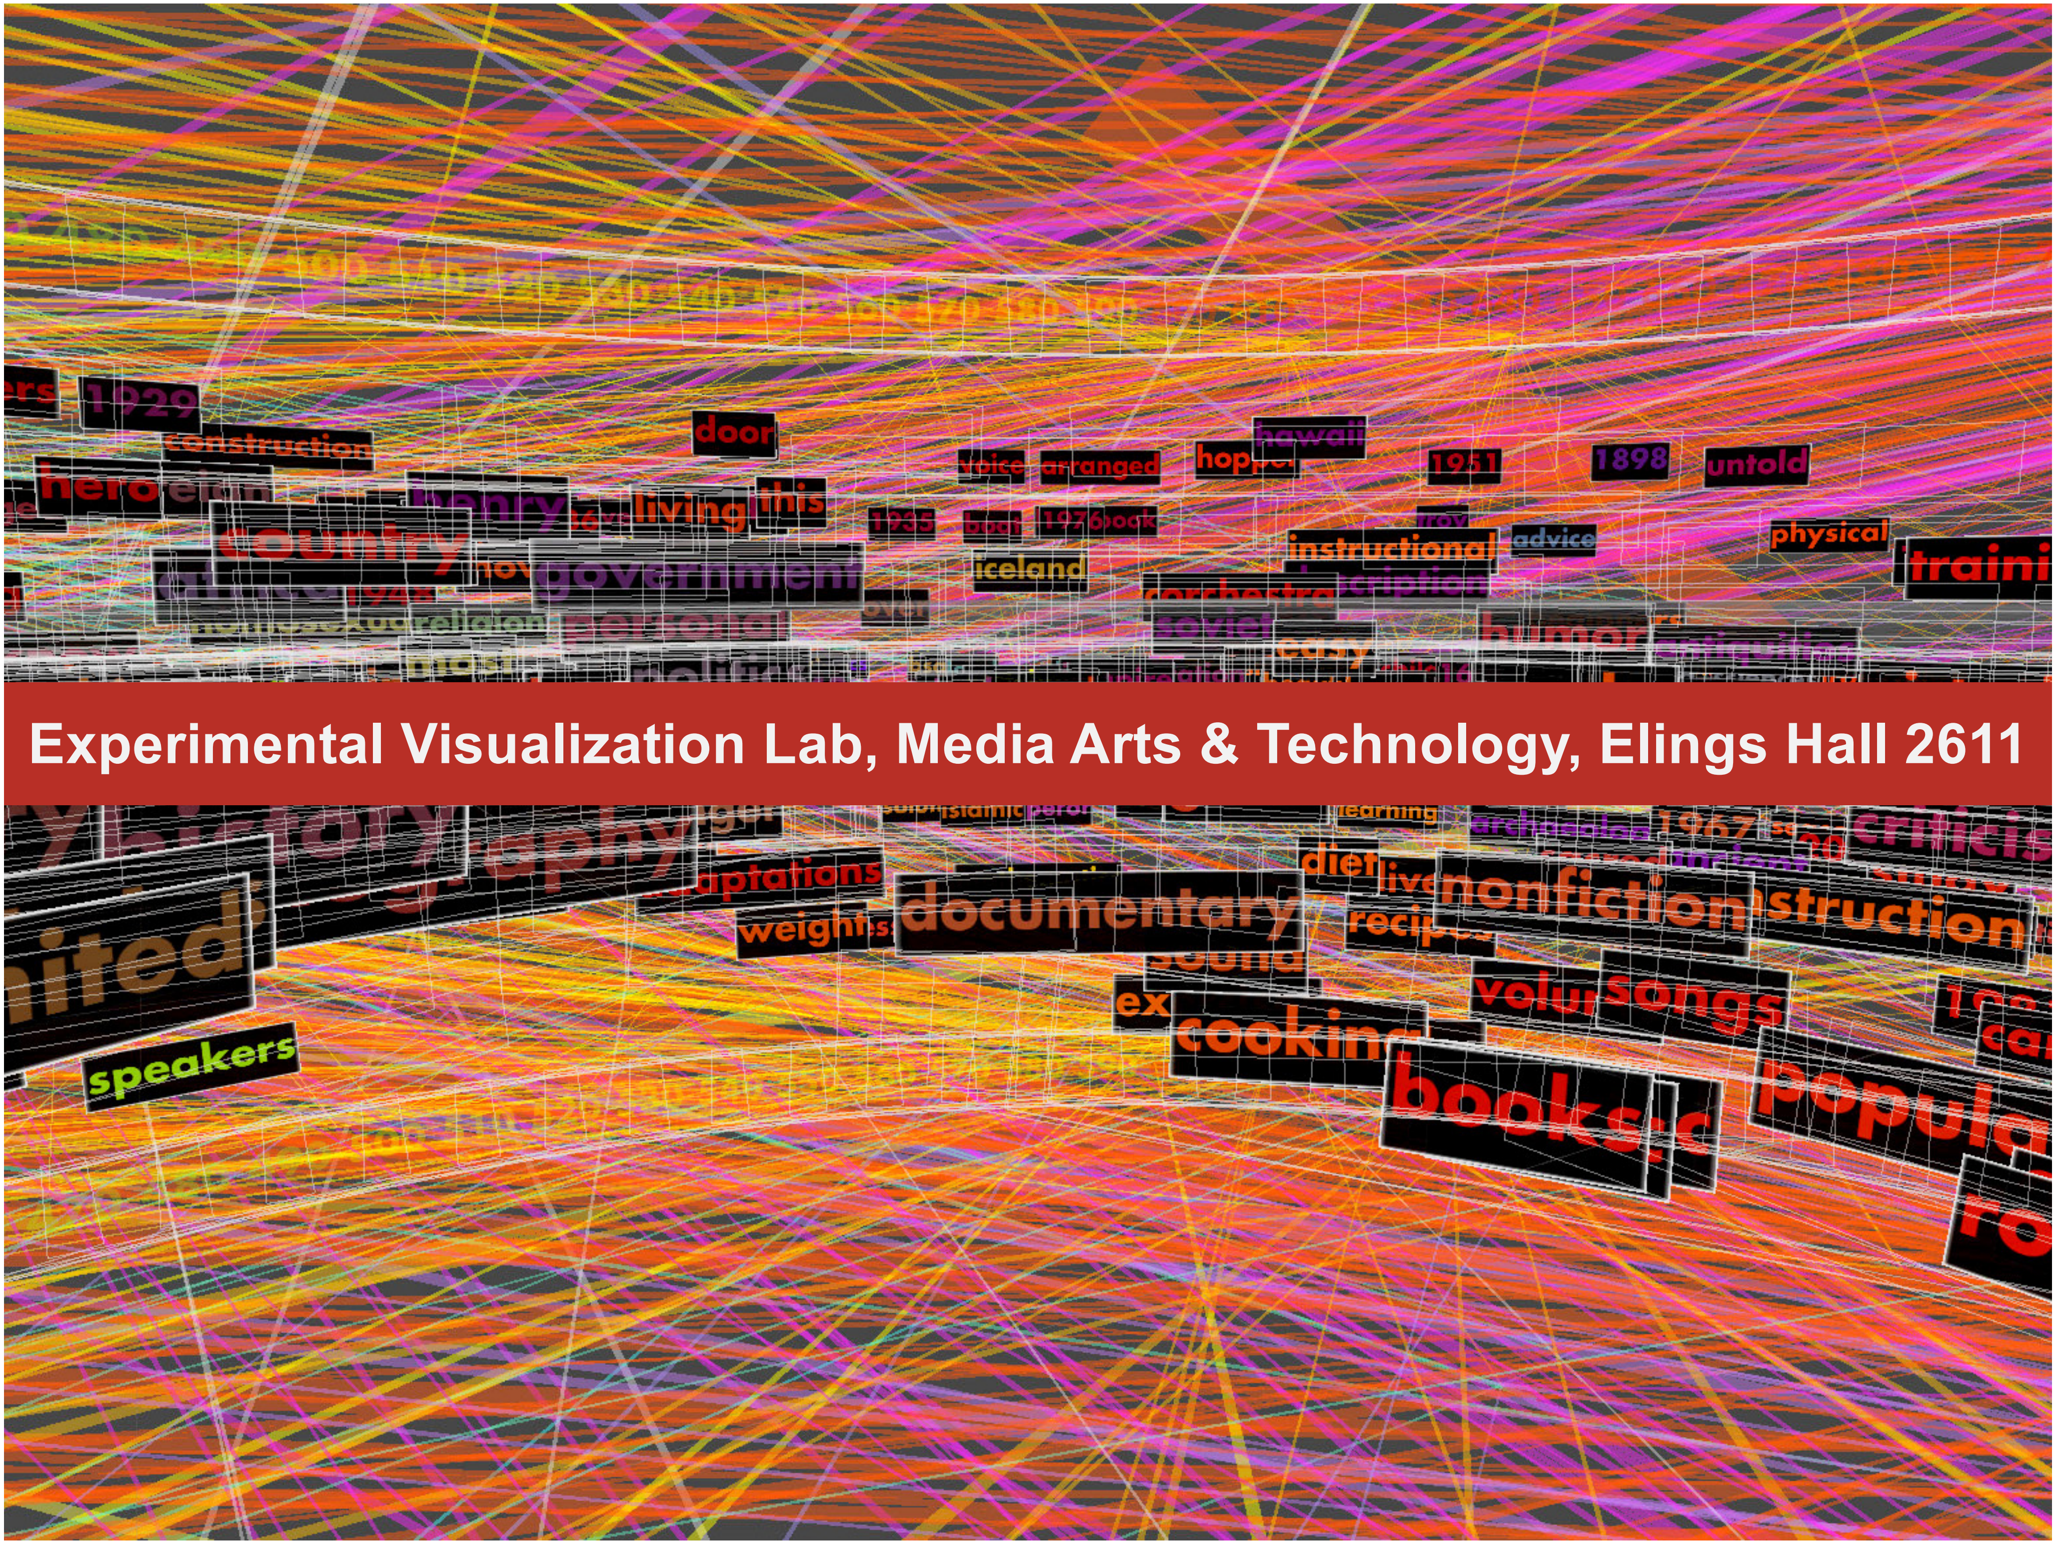

Experimental Visualization Lab, Media Arts & Technology, Elings Hall 2611

Experimental Visualization Lab

Founded in 2006 when MAT moved to CNSI, the lab focuses on creative explorations in the fields of:

- Data Visualization | Data Science
- Computational Imaging & Photography
- Mobile Multi-Camera Robotic Systems
- Digital Interactive Installations

***New Direction**: Image Generation Through Machine Learning*


CRUNCH
WIN

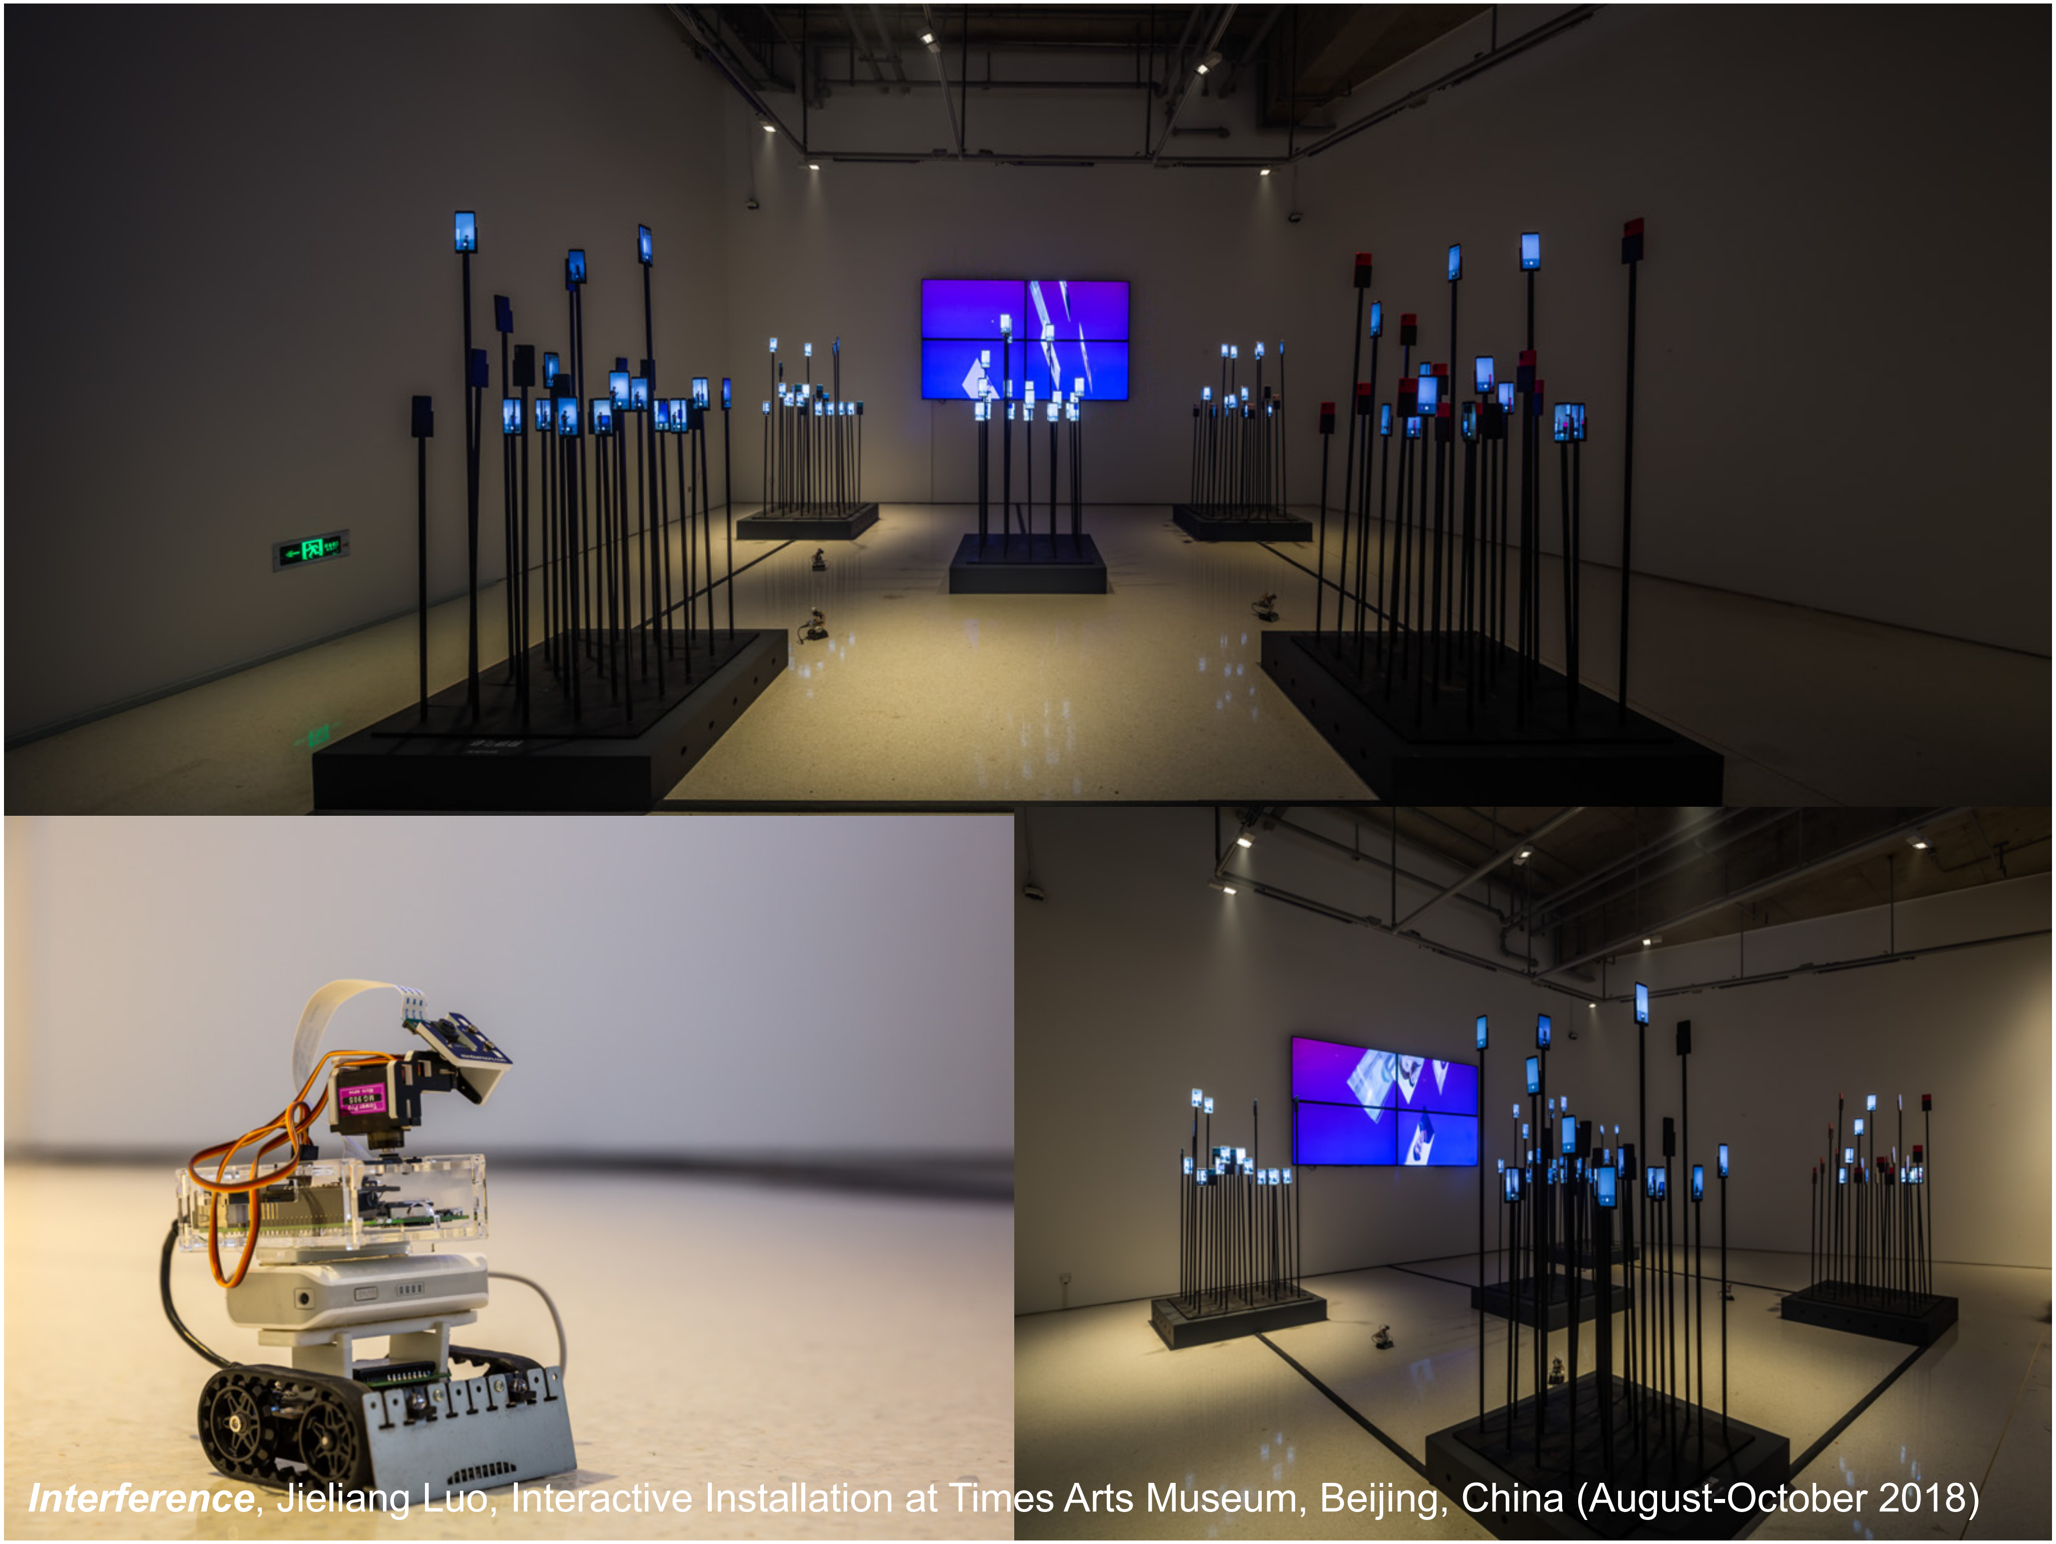

Interference, Jieliang Luo, Interactive Installation at Times Arts Museum, Beijing, China (August-October 2018)

Courses

- **M254 Arts & Engineering / Science Research** - A course to study research methodologies in arts, science and engineering through lab visits. Future version to focus on Data Science in the labs
- **M259 Data Visualization** - A studio course dealing with data mining, data aggregation and the visualization of abstract data in various modes that include frequency, spatialization, and 3D interaction.
- **M265 Open Projects in Optical/Motion Computational Processes** - An experimental studio course where students individually define a research topic that explores new directions in computational photography.
- **Arts130 Digital Visual Culture** - A theory based course that examines the impact of digital technologies on the study and analysis of the image.
- **Arts185 Optical Digital Culture** - focused on programming the behavior of mobile cameras that capture and project images connected to computers
- **In development - 3D Photography & VR Visualization** (in planning) Exploration of simulation of 3D virtual still and time-based media camera-based captured images
- **In development – Computational Research & Studio** – A studio critique course in which each student defines a research/applied project to undertake

M254 Arts & Engineering / Science Research - Research methodologies in arts, science and engineering course through Science lab visits.

Interrogating Methodologies

Exploring Boundaries in Art & Science

April 18-19, 2014

McCune Conference Room, 6020 HSSB
University of California, Santa Barbara

A multi-disciplinary symposium comparing methodologies from the natural sciences, social sciences, humanities, and the arts to interrogate questions at the heart of research methods and practices

Schedule

Friday, April 18

9:00 - 9:15	**Symposium Day 1**	
	Introduction	George Legrady (MAT, UCSB)
		Barbara Harthorn (CNS, UCSB)
9:15 - 10:30	**Keynote [Video]**	
	Introduction	Bruce Robertson (AD&A Museum, History of Art & Architecture, UCSB)
	Problems in the Theory of Visualization	James Elkins (Art History, School of the Art Institute of Chicago)
10:35 - 12:00	**The Big Picture: Visualizing Big Data [Video]**	
	Introduction	JoAnn Kuchera-Morin (MAT, UCSB)
	Supercomputing the Universe	Joel R. Primack (UC HiPACC, UCSC)
	The Original Big Data: Geospatial Information	Keith Clarke (Geography, UCSB)
	The Big Picture: Discovering knowledge from pictures	B.S. Manjunath (Center for Bio-Image Informatics, UCSB)
12:00 - 1:00	**Lunch**	
1:00 - 2:30	**Interrogating the Methodologies of Art & Science [Video]**	
	Introduction	Lisa Jevbratt (Art / MAT, UCSB)
	Art|Sci in the evolution of intelligence	James K. Gimzewski (Chemistry & BioChemistry, UCLA)
	From knowhow transfer to the sharing of methodologies: PhD research. Z-node	Jill Scott (Swiss Artistsinlabs- ICS. University of the Arts Zurich)
	Avoiding the Itch: Development of a Fluorescence Detection of Poison Oak Oil	Rebecca Braslau (Chemistry & BioChemistry, UCSC)
	Respondent	Victoria Vesna (DMA, Art-Science Center, UCLA)
2:35 - 3:50	**Citizen Science: How does the Public Contribute to Science? [Video]**	
	Introduction	Marko Peljhan (MAT, UCSB)
	Open science and wicked problems: how new ways of doing science can help tackle pressing global challenges	Lina Nilsson (Teklalabs, Blum Center, UC Berkeley)
	The Berkeley Atmospheric CO2 Observation Network--monitoring the heartbeat of the urban carbon cycle	Virginia Teige (Beacon, Chemistry, UC Berkeley)
	GalaxyZoo and the Zooniverse of Astronomy Citizen Science	Joel R. Primack (UC HiPACC, UCSC)
3:50 - 4:00	**Coffee Break**	
4:05 - 5:15	**Asking the Right Questions, Avoiding the Wrong Ones, How Research Evolves [Video]**	
	Introduction	Marcos Novak (MAT, UCSB)
	Panning for gold in the lab: Scientists as prospectors	Dave Deamer (Biomolecular Engineering, UCSC)
	Building on a chip? Scalar boundaries of multifunctional building enclosures	M. Paz Gutierrez, (BIOMS, Architecture, UC Berkeley)
5:30 - 6:45	**Reception in Media Arts & Technology**	Allosphere, Experimental Visualization Lab, Translab


UCSB

Experimental Visualization Lab

The experiments, projects and courses we do contribute to the arts, engineering, scientific communities in the following ways:

- Explore *new forms of visualizations* through implementation of advanced knowledge of visual language, image syntax and semiotics through computation
- *Rapid prototyping of concepts* through multimodal visual, spatial, interactive ways
- Address *image-based research* that are on the fringe and may result in new types of artistic practices and possibly contribute to enhancing engineering/scientific research

The Experimental Visualization Lab

M259 Data Visualization - 2018

VISUALIZING THE CULTURE ANALYSIS IN 2D

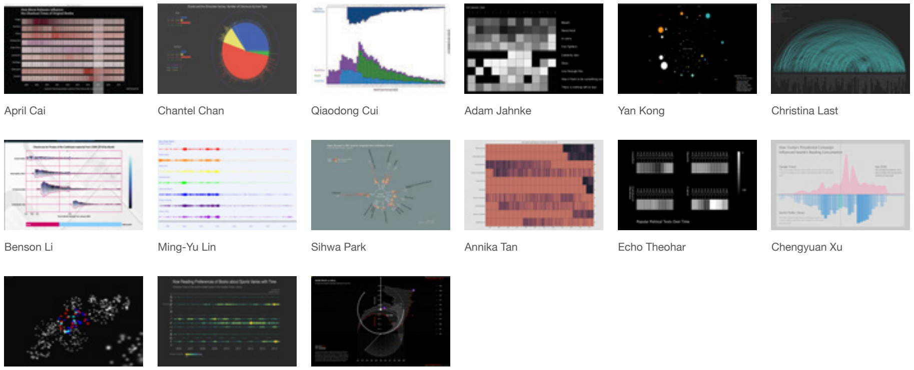

April Cai | Chantel Chan | Qiaodong Cui | Adam Jahnke | Yan Kong | Christina Last
Benson Li | Ming-Yu Lin | Sihwa Park | Annika Tan | Echo Theohar | Chengyuan Xu
Zhenyu Yang | Junxiang Yao | Weidi Zhang

3D INTERACTION & CHANGE OVER TIME

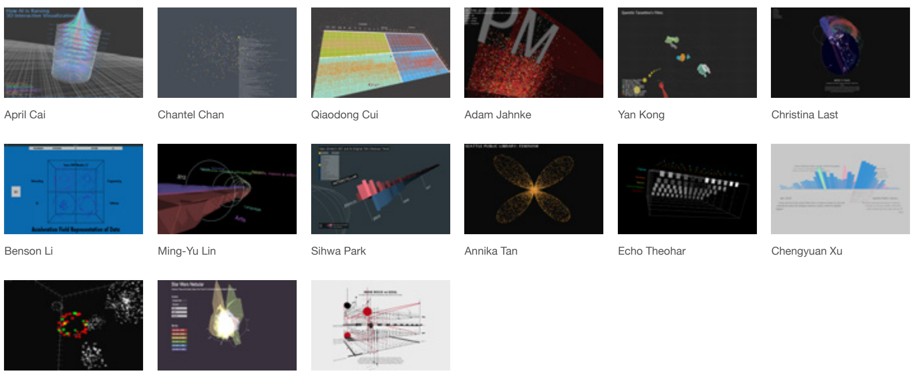

April Cai | Chantel Chan | Qiaodong Cui | Adam Jahnke | Yan Kong | Christina Last
Benson Li | Ming-Yu Lin | Sihwa Park | Annika Tan | Echo Theohar | Chengyuan Xu
Zhenyu Yang | Junxiang Yao | Weidi Zhang

STUDENT DEFINED VISUALIZATION

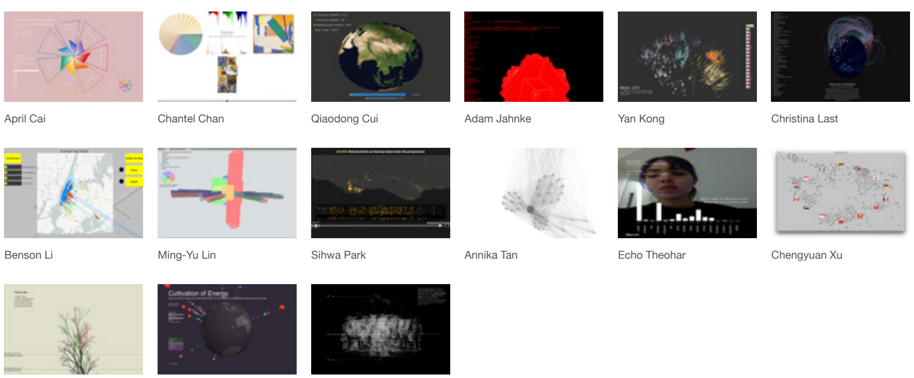

April Cai | Chantel Chan | Qiaodong Cui | Adam Jahnke | Yan Kong | Christina Last
Benson Li | Ming-Yu Lin | Sihwa Park | Annika Tan | Echo Theohar | Chengyuan Xu
Zhenyu Yang | Junxiang Yao | Weidi Zhang

Topic Modeling and Word Vector Visualization of Open Library Fiction Subjects

MAT 259, 2017 **Hannah Wolfe**

Concept Visualize fiction titles using word2vec mapping for book's subjects in 3 dimensions, and graphing the titles and subjects on the resulting word cloud.

Dataset The fiction titles and subjects came from Open Library:
Open Library Bulk Download
The books subjects were analyzed using LDA from the gensim python library to create 10 topics:
Gensim: LDA Model
The google news word vectors were used as a pregenerated Word2Vec model:
Google News Word Vectors Description Google News Word Vectors Download
I retrained the google news word vectors with the book subjects using Gensim:
Gensim: Word2Vec Model
The resulting word vectors were brought down to three dimensions using t-SNE from the scikit learn python library:
Scikit-Learn: t-SNE

Process I first started the visualization with the fiction titles from my 3D library visualization. I used word2vec for mapping the words in 3 dimensions, and graphing the titles on the resulting word cloud. This differed from the p 3D library visualization because I was using pre-trained word vectors from google and continuing to train it with the book titles, labeling all the words in a title, and testing how much different amounts of training effected the visualization.

"aunt" selected when viewing the word cloud

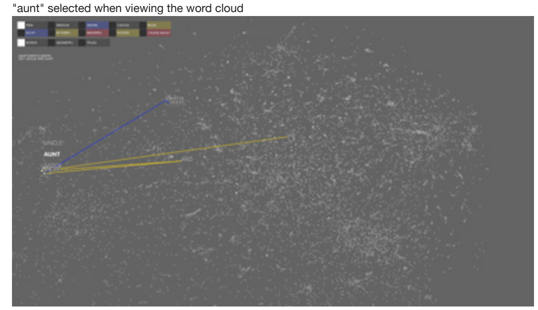

book selected when viewing the title cloud

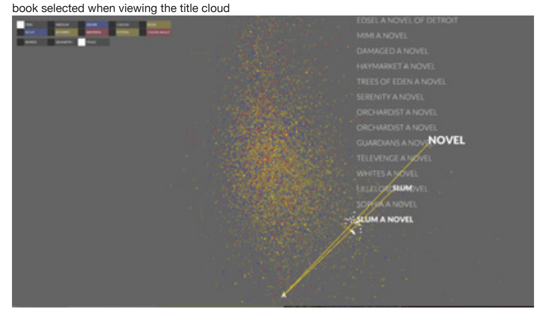

"fire" hovered over when viewing the title geometries

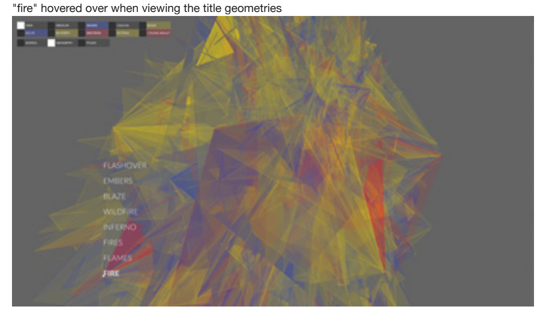

fire selected when viewing the title geometries

M259 DataVis Crs Enrollment

- 19w259 **PhD** (Geog 3; Phys 2; Math 1) - **Ms** (ECE 1) – **UG** (Comm 1; Art 1; CS 2)
- 18w259 **PhD** (Stat 1; Film 1; PolSci 1; MAT 4) – **Ms** (CS 1; Art 2; ECE 1; MAT 2) – **UG** (1)
- 17w259 **PhD** (CS 2; Phys 2; Bren 2; Geog 2; MAT 1) – **UG** (Stat 1; CS 1)
- 16w259 **PhD** (CS 1; Geog 1; MAT 4) – **Ms** (CS 1; MAT 6) – **UG** (Art 2, CS 2)
- 15w259 **PhD** (Geog 2; PolSci 1; MAT 5; Bren 1; Stat 1) – **Ms** (CS 1; MAT 3) – U (CS 2)
- 14w259 **PhD** (Geog 3; CS 1) – **Ms** (MAT 3; CS 1)

Making Visible the Invisible, Seattle Central Library (2005-2019)

Lab Direction

Under the direction of George Legrady, distinguished professor of Digital Media at UCSB since 2000. Former chair of Media Arts & Technology (2013-2017), Guggenheim Fellow recipient in Fine Arts, 2016.

George Legrady is an internationally exhibited and published artist and scholar whose works address the intersections of photography, computer media, and interactive installations. He is an early adopter of computer media (1980s). His focus is the exploration of how technologies transform visual content and result in new kind of representations.

His artworks are in the collection of the San Francisco Museum of Art, the Los Angeles County Museum of Art, the National Gallery of Canada, the Centre Pompidou Museum, Paris, the Santa Barbara Museum of Art, the musée d’art contemporain in Montreal, the Philbrook Museum of Art, the Smithsonian Institution, and others.

Permanent public commissions include the Los Angeles Metro Rail (2007), the Corporate Executive Board (Arlington), and the Seattle Central Library, a data visualization installation begun in 2005 that may be the longest running such project to-date.

Swarm Vision, SIGGRAPH 2013

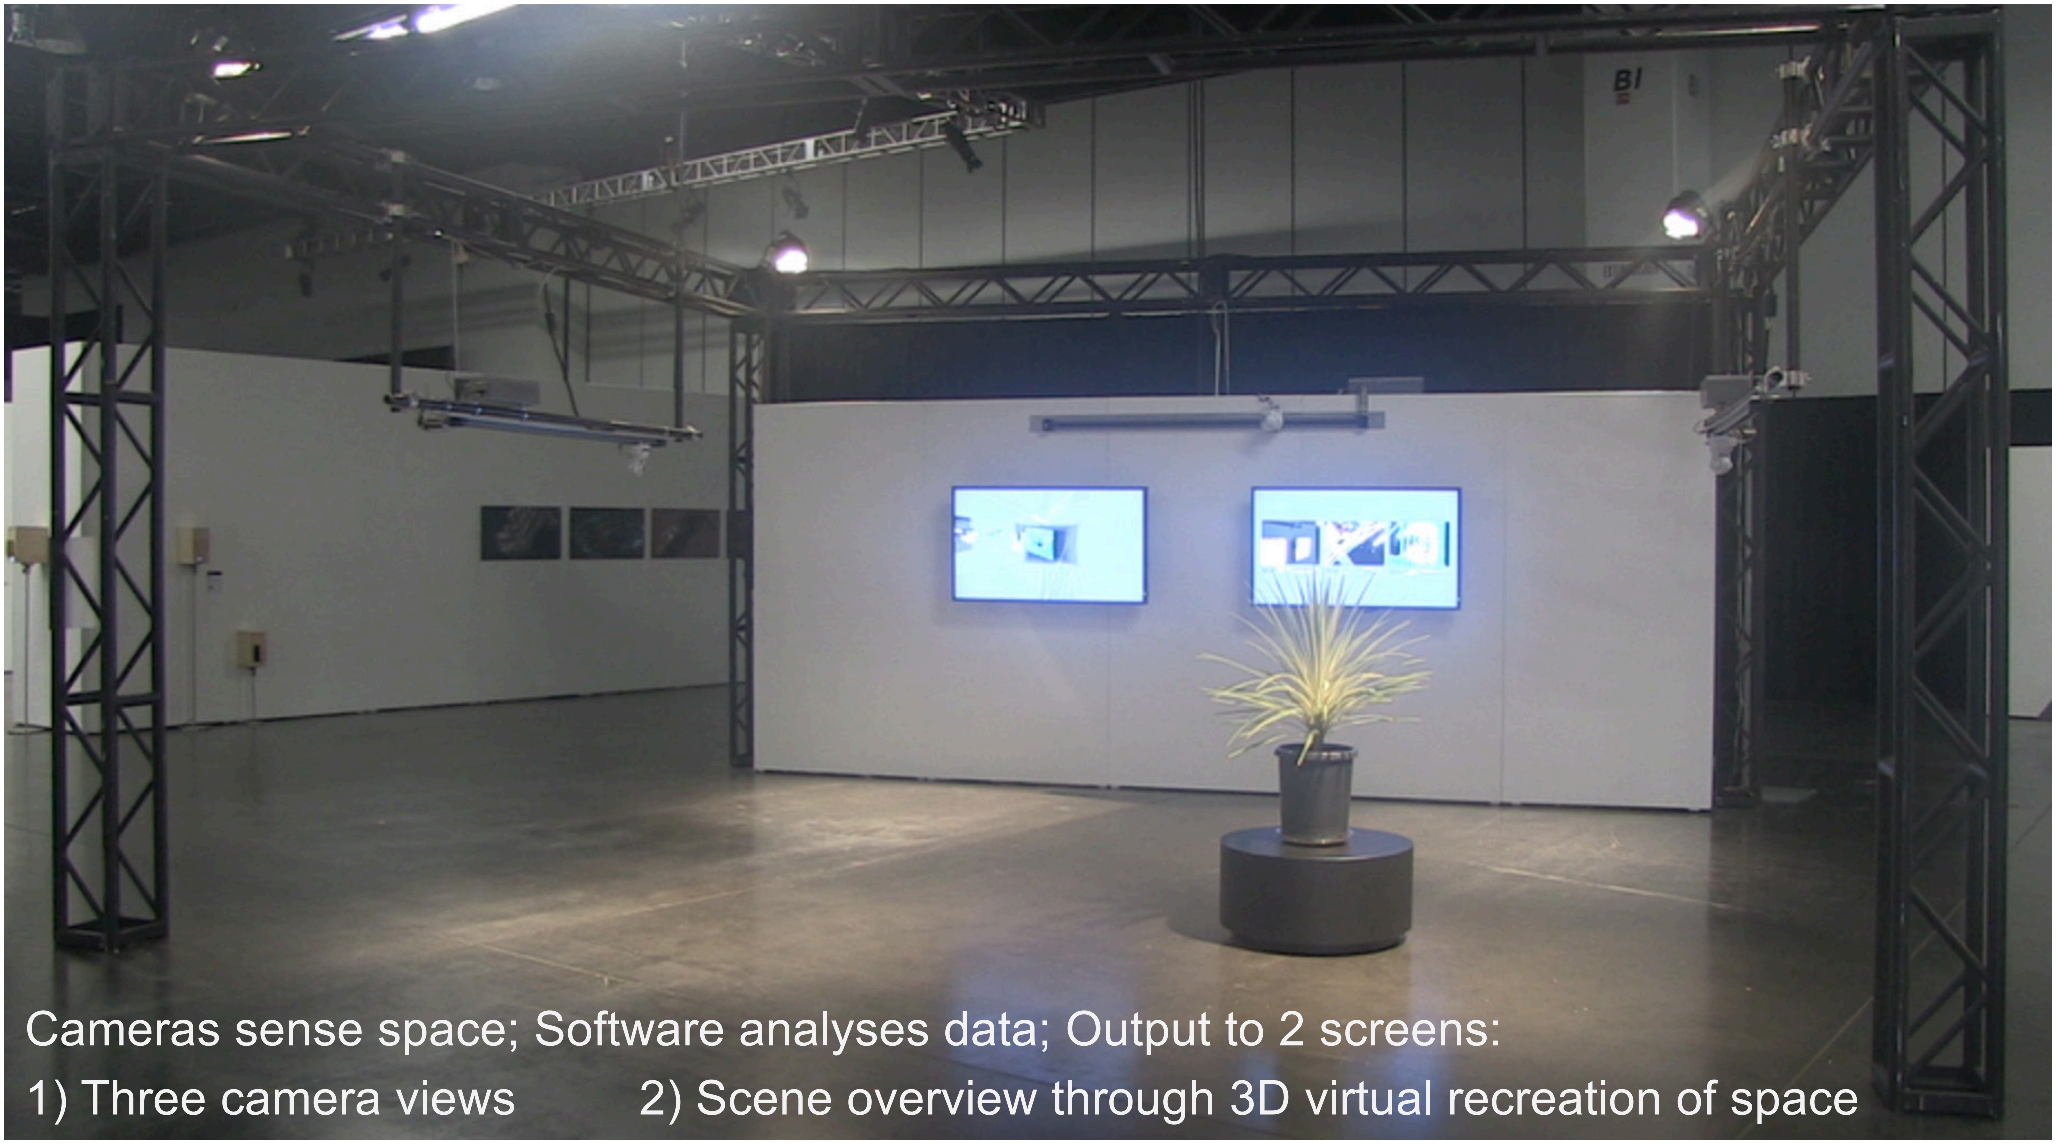

Cameras sense space; Software analyses data; Output to 2 screens:

1) Three camera views 2) Scene overview through 3D virtual recreation of space

SwarmVision: Autonomous Aesthetic Multi-Camera Interaction

Danny Bazo, George Legrady, Marco Pinter
Experimental Visualization Lab, Media Arts & Technology,
University of California Santa Barbara, Santa Barbara, CA 93106, USA

ABSTRACT

Initiated by research in autonomous swarm robotic camera behavior, SwarmVision is an installation consisting of multiple Pan-Tilt-Zoom cameras on rails positioned above spectators in an exhibition space, where each camera behaves autonomously based on its own rules of computer vision and control. Each of the cameras is programmed to detect visual information of interest based on a different algorithm, and each negotiates with the other two, influencing what subject matter to study in a collective way.

SCREEN VISUALIZATIONS

In the installation, four visualizations are featured on two screens or projections. The first screen features what each of the three cameras "see" - a depiction of what their vision algorithms are currently processing. Informational labels under each of the livefeed describe the incoming data and analysis, giving a sense of each of the camera's performance.

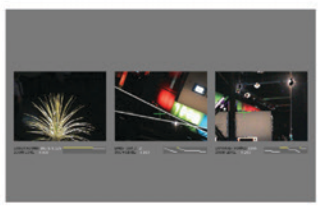

The second screen shows an overview in a 3D reconstruction of the gallery environment showing the location of the cameras, where they point, and the video they capture transformed into still images. Each camera continuously produces 10 still frames per second, and fills the 3D space with up to a hundred images per camera resulting in a volumetric form of layered stacked photographs that continuously change as images fade away.

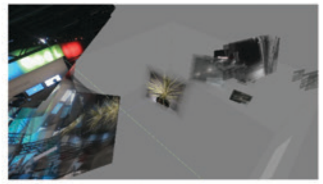

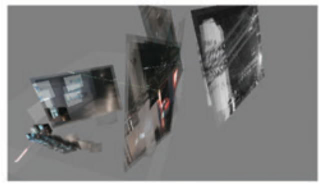

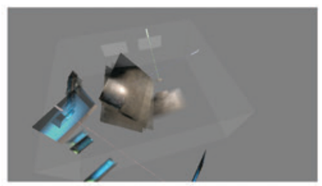

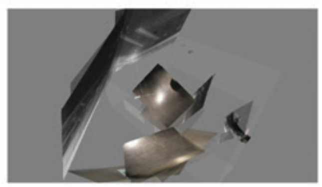

COMPUTER VISION

Each robot has a built-in computer vision algorithm such as: a contrast-detecting Pyramid of Gaussians distance in pixel space, a straight-line-detecting Hough line filter, or a saturated-pixel-seeking algorithm. These were chosen for their dissimilarity and their relationship to the parallel, heterogeneous feature detectors operating in the visual systems of living creatures.

COMMUNICATION BETWEEN CAMERAS

Over time, the robots converge on and explore the features of regions of high visual saliency in the environment together. This is enabled by a simple target acquisition and communication protocol. When a robot detects a local maximum in the saliency of its visual environment, it tells the other robots the location of its current gaze. All robots then converge on that location and proceed with their saliency-driven behavior.

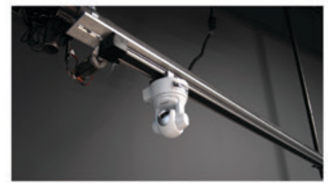

RAIL MOVEMENT

Commands from the main computer drive the robots in the direction of their gaze. They dolly towards the subject of their attention in order to improve resolution of their image, if possible.

PUBLIC INTERACTION

Viewers watching the cameras become instantly aware that the robots are studying and analyzing the scene. Looking at the large lens of a camera, a viewer would know exactly where in the room a robot was "looking." Just as humans are drawn to look at an area of space indicated by the direction of gaze of another human, so are viewers drawn to look where the robots are watching. Viewers in the space also desire attention from the robots, and often move in front of the camera eye, and experience satisfaction when having been designated a subject of interest. During moments when the robot cameras appear to be looking at each other, viewers experience a sense of non-verbal communication between those autonomous agents.

MAT Media Arts and Technology Graduate Program

This material is based upon work supported by the Robert W. Deutsch Foundation and the National Science Foundation under Grant #1149001. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Swarm Vision, 2 screen installation in Mois de la Photo, Festival Montreal (2013)

← BACK PAST LIST
CURRENT
UPCOMING
PAST

George Legrady: Voice of Sisyphus丨Slice

2015-6-3

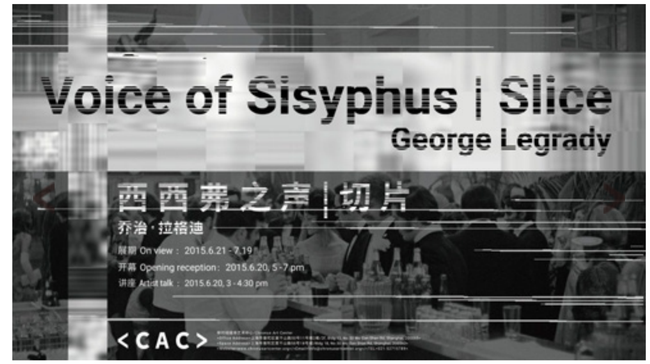

Voice of Sisyphus丨Slice
On View: 2015.6.21-7.19
Opening Reception: 2015.6.20, 5-7pm
Artist Talk: 2015.6.20, 3-4:30 pm
Venue: Chronus Art Center

Chronus Art Center (CAC) is pleased to present Voice of Sisyphus丨Slice, an exhibition by the Hungarian born American artist George Legrady.

Voice of Sisyphus丨Slice is composed of two works - Voice of Sisyphus and Slice that both integrate algorithmic processes as a means to data visualization and sonification by which new forms of aesthetic representations and socio-cultural narrative experiences are created. While the translation of sampled data from an image to sound in Voice of Sisyphus proposes and questions the potentiality of an aesthetic equivalence from one medium to another through simulated synaesthetic modeling, the transitional movement from a state of legibility of the photograph to an abstract and photographically unrecognizable state in Slice investigates into our perceptual experience of an image and examines the ways through which culturally identifiable meaning is constructed and could be deconstructed.

The repetitive and cycling feature of both works easily bring to mind the Greek myth of king Sisyphus' dilemma who was compelled to ceaselessly roll an immense boulder up a hill, only to watch it roll back down repeatedly, whereas their experimental nature with fluctuating parameter values provides variations within the constraints of their settings to allow for aesthetic breadth and constant fluidity of perceptual and aesthetic experience within the epic-cycle.

Voice of Sisyphus
2013 | George Legrady | Multimedia Installation
Image analysis – George Legrady
Audio and spatialization software development- Ryan McGee
Audio composition software development - Joshua Dickinson

Voice of Sisyphus is a multimedia projection with 4 channel spatialized sound installation in which a black and white image is sonified by a computer program which synthesizes image segments and produces sounds resulting in a continuously evolving composition.

Voice of Sisyphus is a time-based study of a single photograph, realized as a continuous performing audio-visual composition. It is presented as a multimedia installation with a large cinematic projection and 4 channel audio, spatializing sounds by speakers positioned in each of the four corners of the exhibition space. The sound composition is created out of the analysis of visual regions in the photograph through the sampling of pixel clusters as the software "reads" and translates areas of the image at 30 frames per second in four ways: 1) Stationary 2) Smooth scanning, 3) Rectangular divisions and 4) Regions of interest in the image such as faces, clusters of people, windows, glasses, lines, mirrors, plants, decorations, etc. within the image.

Voice of Sisyphus was exhibited at Edward Cella Gallery, Los Angeles (2011); Nature Morte Gallery, Berlin (2012); "Sights & Sound", Beall Center for Art + Technology, UC Irvine (2013)

SLICE
2011| George Legrady | Custom software animation |3840 x 1080 pixels | Image dimensions variable
Visualization software - Yun Teng

Eight black and white and color tinted photographs of a formal social event are repeatedly sliced in half, until reduced to abstract visual slivers that are no longer photographically recognizable. Once arrived to that abstract state, the slices systematically re-assemble, doubling in size into a different image.

The sequence consists of eight images in total and each is color coded with one of four colors. This results in a situation where the viewer only perceives the color changes every second image. The transition between images

MORE U
SPECIAL
2015-06-
Artist Tal
the Medi
COME T
Opening
Wednes
Visit Us:
No.76, N
Station
Hanzhon
801 to C
Zhongta
exchang
Aomen R

VOICE OF SISYPHUS: AN IMAGE SONIFICATION MULTIMEDIA INSTALLATION

Ryan McGee, Joshua Dickinson, and George Legrady

Experimental Visualization Lab
Media Arts and Technology
University of California, Santa Barbara
ryan@mat.ucsb.edu, dickinson@mat.ucsb.edu, legrady@arts.ucsb.edu

ABSTRACT

Voice of Sisyphus is a multimedia installation consisting of a projection of a black and white image sonified and spatialized through a 4 channel audio system. The audio-visual composition unfolds as several regions within the image are filtered, subdivided, and repositioned over time. Unlike the spectrograph approach used by most graphical synthesis programs, our synthesis technique is derived from raster scanning of pixel data. We innovate upon previous raster scanning image to sound techniques by adding frequency domain filters, polyphony within a single image, sound spatialization, and complete external control via network messaging. We discuss the custom software used to realize the project as well as the process of composing a multimodal artwork.

1. INTRODUCTION

Voice of Sisyphus relies on an Eisensteinian process of *montage* [1], the assembling of phrases with contrasting visual and tonal qualities as a way to activate change that can be considered as a narrative unfolding. Whereas cinematic montage involves the contrast of discontinuous audio-visual sequences as a way to build complexity in meaning, in this work, the referent photograph that is processed does not ever change, except through the filtering that generates the tonal and visual changes. As the composition evolves but then returns to where it began, the event brings to mind the Greek myth of king Sisyphus who was compelled to ceaselessly roll an immense boulder up a hill, only to watch it roll back down repeatedly. The intent is to have a continuously generated visual and sound composition that will keep the spectator engaged at the perceptual, conceptual, and aesthetic levels even though the referent visual source is always present to some degree.

Voice of Sisyphus was partially inspired by the overlay of image processing techniques in Peter Greenaways 2009 film, *Wedding at Cana*[1], a multimedia installation that digitally parses details of the 1563 painting by the late Renaissance artist Pablo Veronese in a 50 minute video. The filmmaker skillfully uses computer vision techniques to highlight, isolate, and transform visual details to explore the meaning of the visual elements in the original painting.

The project evokes two early digital works by Legrady. *Noise-To-Signal*[2] (1986) is an installation artwork that uses digital processing to explore the potential of image analysis, noise, and Information Theory's definition of noise to signal. *Equivalents II*[3], realized in 1992, is another interactive digital media artwork that implements 2D midpoint fractal synthesis as a way to create organic-looking abstract images whose abstract cloud-like visual forms were defined by textual input provided by viewers. Both artworks integrated synthesis algorithms to generate cultural content through computational creation of images.

Most experiments examining the relationships between sound and image begin with sounds or music that influence the visuals. Chladni's famous 18th century "sound figures" experiment involves visual patterns generated by playing a violin bow against a plate of glass covered in sand[2]. 20th century visual music artists often worked by tediously synchronizing visuals to preexisting music. Though, in some cases, the sounds and visuals were composed together as in *Tarantella* by Mary Ellen Bute. Today, visual artists often use sound as input to produce audio-reactive visualizations of music in real-time.

Less common are technical methodologies requiring images as input to generate sound. However, in 1929 Fritz Winckel conducted an experiment in which he was able to receive and listen to television signals over a radio[2], thus resulting in an early form of image audification. Rudolph Pfenninger's *Tnende Handschrift* (Sounding Handwriting), Oskar Fischinger's *Ornament Sound Experiments*, and Norman McLaren's *Synchromy* utilized a technique of drawing on film soundtracks by hand to synthesize sounds. *Voice of Sisyphus* continues in the tradition of the aforementioned works by using visual information to produce sound.

2. SOFTWARE

Custom software was developed to realize the artist's vision of translating an image into a sonic composition. Although *Voice of Sisyphus* is based on a particular photograph, the software was designed to be used with any image. Once an image file is imported one may select any number of rectangular regions within the image as well as the entire image itself to sonify. Greyscale pixel values within a region are read into an array, filtered, output as a new image, and read as an audio wavetable. The wavetables of multiple regions are summed to produce polyphonic sound. Consideration was taken for real-time manipulation of region locations and sizes during a performance or installation without introducing unwanted audio artifacts.

[1] http://www.factum-arte.com/eng/artistas/greenaway/veronese_cana.asp
[2] http://www.mat.ucsb.edu/g.legrady/glWeb/Projects/noise/noise.html
[3] http://www.mat.ucsb.edu/g.legrady/glWeb/Projects/equivalents/Equi.html

Fabian Offert's doctoral research is at the interface of computer science, media studies, and continental visual studies (Bildwissenschaft)

Problem: CNNs (convolutional neural networks) are very good at object recognition tasks but it is difficult to know why: we need **interpretable machine learning**

--

One technique of interpretable machine learning is *Feature Visualization where* a noise image is iteratively tuned to maximally activate one or multiple internal components ("neurons") of a CNN

Eventually, the image shows what these components have learned (patterns, objects, colors, etc.) which allows a human observer to better understand "how the machine perceives the world"

https://www.zentralwerkstatt.org/

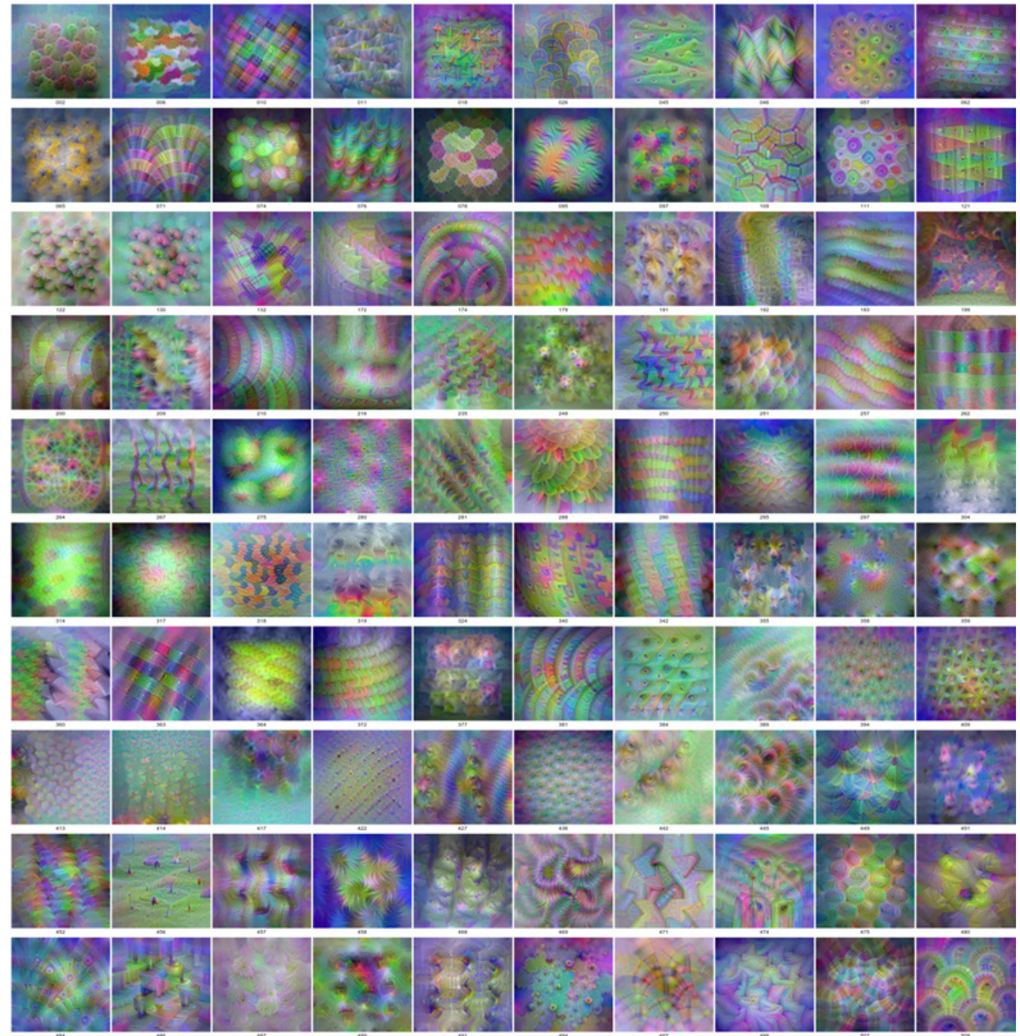

Feature visualizations of an Inception v1 (detection and classification) CNN trained on ImageNet patterns

Visual Diagnostics for Deep Reinforcement Learning Policy Development

Jieliang Luo, Sam Green, Peter Feghali, George Legrady, and Çetin Kaya Koç
University of California, Santa Barbara
https://arxiv.org/abs/1809.06781, jieliang@ucsb.edu

UC SANTA BARBARA

Background

Reinforcement learning (RL) is a family of methods aimed at training an **agent** to collect rewards from an environment through trial-and-error approaches. Since the deep Q-network (DQN) algorithm was introduced in 2013, there has been a surge of interest in using convolutional neural networks (CNN) in vision-based RL algorithms. In the context of cyber-physical systems, vision-based RL has exciting potential to provide high levels of autonomy in applications like robotics, self-driving cars, and infrastructure inspection. However, CNNs are known to be opaque to debugging and RL's emphasis on trial-and-error demands rigorous behavioral verification before they may be allowed control over safety-critical cyber-physical systems. This work adapts CNN visualization techniques to the domain of RL.

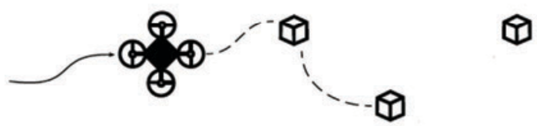

Vision-based RL algorithmic advancements are typically introduced in the context of simple games, e.g. benchmarking with the Arcade Learning Environment. In this work we modified Microsoft's AirSim, which is a photorealistic 3D simulator based on the Unreal Engine. We modified AirSim to support RL[1]. Our agent was a drone with three available actions: forward, left, and right. These actions had deterministic effects. At the beginning of each episode, the drone would be reset and cubes would be randomly distributed in front of it.

Reinforcement learning approach

The agent received a reward of +1 for each cube that it "collected". The goal for our agent was to collect as many rewards as possible each episode. The agent's policy was a CNN parameterized by weights θ, so the goal was to solve the following optimization problem:

$$\theta^* = \arg\max_\theta \sum_{t=0}^{T-1} r(s_t, a_t),$$

where s_t and a_t are the states and actions at time t and $r(s_t, a_t)$ is the reward obtained from the environment for the given state and action. For this simple problem, the REINFORCE algorithm was adequate.

Visualization methods and results

CNN visualizations are useful for identifying strengths and weaknesses in a trained network. For example, the *class visualization* for GoogLeNet's "saxophone" class indeed extracts a saxophone shaped object from the network. That is, the method generates an image, which, when input into the trained GoogLeNet CNN, will maximize the output probability of the saxophone class. However, when looking at the generated image, one can clearly see that the outline of a man has also been extracted from the network! Thus, CNN visualizations provide insight into what a network has learned to pay attention to, and a human can then determine if modifications are required.

In the context of RL, existing CNN visualization techniques attempt to cluster inputs according to their resulting action, provide decision attribution, or visualize canonical actions. Techniques considered here include:

- t-SNE maps – Clusters similar inputs by the actions they trigger.
- Attribution visualization – Identifies image regions most responsible for an action decision.
- Action visualization – Generates inputs which trigger specified actions.

t-Distributed Stochastic Neighbor Embedding (t-SNE) is a dimensionality reduction algorithm developed by [1]. It is well suited for visualizing high-dimensional datasets. The method positions each high-dimensional datapoint (e.g. image) in a two or three-dimensional map in a way that similar datapoints are nearby and dissimilar ones are distant. The most recent use of t-SNE is to use a trained convolutional neural network (CNN) to extract features from each image, feed the features to t-SNE to get the position of each image, and arrange the images on a 2D or 3D space based on the given positions.

The figures below are t-SNE visualizations of three policies: a high-performance, poor-performance, and right-and-forward-only policy. The tinting of each patch is based on the action taken by the policy, given the drone's observation: red indicates "forward", green indicates "left", and blue indicates "right". By studying the t-SNE outputs it is possible to identify inputs which cause correct or incorrect behavior.

Figure 1. High-performance policy

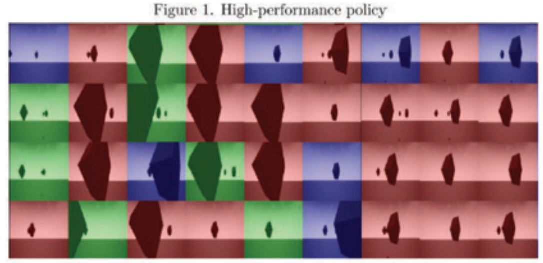

Figure 2. Poor-performance policy

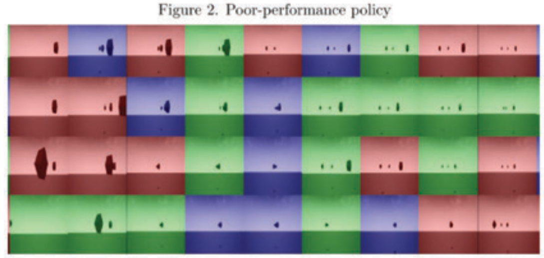

Figure 3. Forward-and-right-only policy

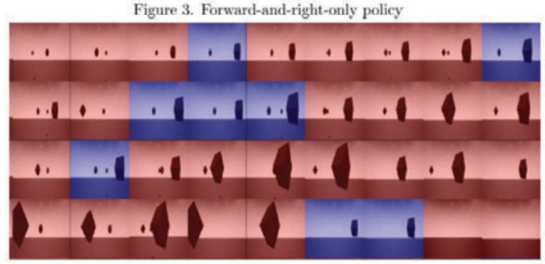

Attribution visualization techniques highlight regions in an input which are most responsible for a particular action in a CNN-based policy. We will use an attribution visualization technique called **Gradient-weighted Class Activation Mapping** (Grad-CAM) [2], which highlights regions in the input image most responsible for an action probability.

1. New versions of AirSim now have native RL support.

K Ç L A B

The high-performance policy's class visualization in Fig. 4 clearly explains what the policy is looking for, where the bias toward the "left", "forward", or "right" depends on the position of the cube. Similarly, Figs. 5 and 6 provide insights into the poor and forward-and-right-only policies.

Figure 4. High-performance policy

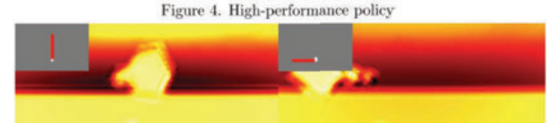

Figure 5. Poor-performance policy

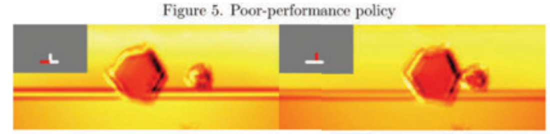

Figure 6. Forward-and-right-only policy

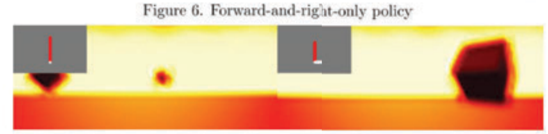

Action visualization methods generate visual inputs which activate a particular output in a *trained* neural network. This approach allows for a high-level of human comprehension about the behavior of a network, rather than treating the network like a black-box model. For our specific feature visualization approach, we use **Class Model Visualization** (CMV) [3].

Action visualizations for the forward-and-right-only policy Fig. 9 highlight one of the challenges in reinforcement learning. In this case, the drone experienced an early success by moving "right" and "forward", which resulted in the elimination of "left" action probabilities. The remedy for this was to lower the learning-rate of the policy updates.

Figure 7. High-performance policy

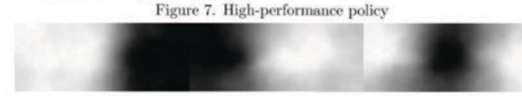

Figure 8. Poor-performance policy

Figure 9. Forward-and-right-only policy

[1] L. V. D. Maaten and G. Hinton, "Visualizing data using t-sne," *Journal of Machine Learning Research*, vol. 9, no. Nov, pp. 2579–2605, 2008.
[2] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-cam: Visual explanations from deep networks via gradientbased localization," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 618–626
[3] K. Simonyan, A. Vedaldi, and A. Zisserman, "Deep inside convolutional networks: Visualising image classification models and saliency maps," arXiv preprint arXiv:1312.6034, 2013

Arxiv paper:

Alumni

- Enrica Dias Lovaglio (MA 2002) - faculty of architecture, CalPoly State University, San Luis Obispo
- Ayoub Sarouphim (MA 2003) - Senior Design Associate, Klai Juba Architects, Las Vegas, Nevada. Practicing architect responsible for a number of designs in Las Vegas (Hard Rock Hotel, Ellis Resort, E Hotel/Casino, Gateway, and others)
- Andreas Schlegel (MS 2004) - Faculty of Digital Design, Lasalle University, Singapore. Internationally featured interactive works. Contributor to the processing.org software development. Previously worked with Art + Com, Berlin
- Eunsu Kang (MA 2005) - professor of New Media/Machine Learning – Computer Science Carnegie Mellon, Internationally exhibiting artist in Asia, US and Europe
- August Black (MA 2005) - PhD 2011, Media Arts and Technology, UCSB. Internationally exhibiting artist
- Ethan Kaplan (MFA 2005, Art dept) - Director of Technology, Fender Instruments, formerly Warner Brothers, Burbank, California
- Aldo Figueroa (MA 2006) - Instructor, 3D Animation, CSU Channel Islands
- Wes Smith (MA 2008) - PhD student, Media Arts and Technology, UCSB. Exhibiting artist, software designer, consultant at cycling74
- Rama Hoetzlein (MA 2008) - PhD 2010, Media Arts and Technology, UCSB. Formerly Assistant Professor,, Aalborg University, Copenhagen, Denmark. Currently Senior Engineer, NVidia, Exhibited artist and published papers in ACM, SIGGRAPH, Leonardo.
- Jeungah Kim (MA 2009) - PhD in Education, Boston University
- Pehr Hovey (MS 2010) - Media artist/engineer, Disney Imagineering, Los Angeles. Fiber optics suits featured at the Guggenheim
- Syed Reza Ali (MS 2010) - Designer/engineer. Google Research, Has worked at Motion Theory, Nokia Research
- Javier Villegas (PhD 2012) - Media artist/engineer, Los Angeles
- Qian Liu (MSc 2012) – Data visualization
- Andres Burbano (PhD 2013) - Media artist/historian, Professor, University of de los Andes, Bogota, Columbia
- Angus Forbes (PhD 2014) - Media artist/engineer. UCSC Computational Media, formerly University of Illinois, Chicago; Faculty at School of Information: Science, Technology and Arts, University of Arizona
- Yoon Chung Han (2016), Assistant Professor, San Jose State University
- Mohit Hingorani (MSc 2016), Dropbox
- Danny Bazo (PhD 2016) – media artist/engineer, software engineer, Meow Wolf, Santa Fe
- Ambika Yadav (MSc 2017) Technical staff, Illumino, Sunnyvale
- Weihao Qiu (MSc 2018), UCSB

Current Affiliations

- UCSB Art Department
- UCSB Data Science Initiative
- UCSB Center for Digital Games Research
- UCSB Center for Information Technology & Society

--

- Institut des mines, Telecom, Paristech, Paris, France (Visiting Fellow 2017 - open)
- iCinema, University of New South Wales, Sydney, Australia (Visiting Professional Fellow 2017 - open)

Funding Support to-Date

- Robert W. Deutsch Foundation (1 GSR each year, 2012-2016)
- National Science Foundation EAGER Intelligence & Information Systems ($124,644, 2012-2016)
- National Science Foundation EAGER Arctic Social Science ($211,121, 2013-2016)
- Center for Nanotechnology in Society at UCSB ($60,000, 2013-2014)
- UCIRA - University of California Institute for Research in the Arts ($10,000)
- UCSB Faculty Senate Research Grants ($11,400, 2016-2017)

COE-HFA Situation Impact on Lab

1) MAT arts-engineering collaboration may be a major opportunity. Currently there is significant funding and resource imbalance between the Engineering and HFA areas which results in a 2-class system. How to resolve this?

- Greater effort to develop joint COE & HFA project and funding proposals
- Administration to follow capital funding model to invest in the MAT interdisciplinary research work through GSR support (most students are not funded!) to offset the imbalance

--

2) High-Level Issues: Both Colleges consider the program to belong to the other. How to get better support and buy-in?

- COE is supportive but has engineering priority and is waiting for MAT to realize contribution
- HFA seems to have lost the vision of how MAT can promote the HFA - may not realize that MAT is a potential pipeline to research resources that could advance HFA

--

3) COE - HFA research activities & cultural differences:

- COE faculty dependent on campus-based GSR/students to advance their work, career and advancement
- HFA faculty's advancement dependent on external activities. Parallel to COE "research as teaching", student involvement in project development becomes a significant learning experience but this is not captured by the academic evaluation system

Experimental Visualization Lab

http://www.vislab.mat.ucsb.edu