CULTURAL AUTOMATION WITH MACHINE LEARNING

SESSION 01: INTRODUCTION

Parag K. Mital
UCLA DMA

LOGISTICS

Logistics

EDI
TA - ????
Technical tools
Learning outcomes
Course structure
Assignments
Course grading

Logistics

EDI

TA - ????

Technical tools

Learning outcomes

Course structure

Assignments

Course grading

Logistics

EDI

TA - ????

Technical Tools

Learning outcomes

Course structure

Assignments

Course grading

Logistics

EDI
TA - ????
Technical tools
Learning outcomes
Course structure
Assignments
Course grading

Logistics

EDI
TA - ????
Technical tools
Learning outcomes
Course structure
Assignments
Course grading

Logistics

EDI
TA - ????
Technical tools
Learning outcomes
Course structure
Assignments
Course grading

WHAT IS MACHINE LEARNING?

HOW IS IT INVOLVED IN OUR SOCIETY/DAILY LIVES?

WHAT ARE SOME CONCERNS WITH ITS USE?

CULTURAL APPROPRIATION

What's the connection between Cultural Appropriation and Machine Learning?

- Deep fakes / misrepresentation / slander
- Job automation
- Forgery
- Spam
- Cultural appropriation
- Security / privacy

What is cultural appropriation?

What is cultural appropriation?

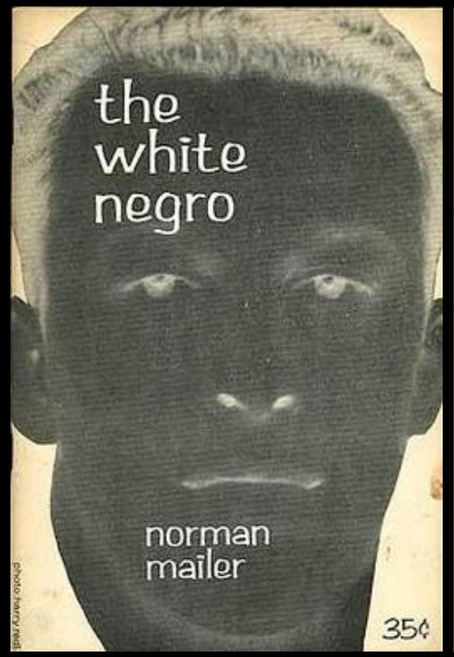

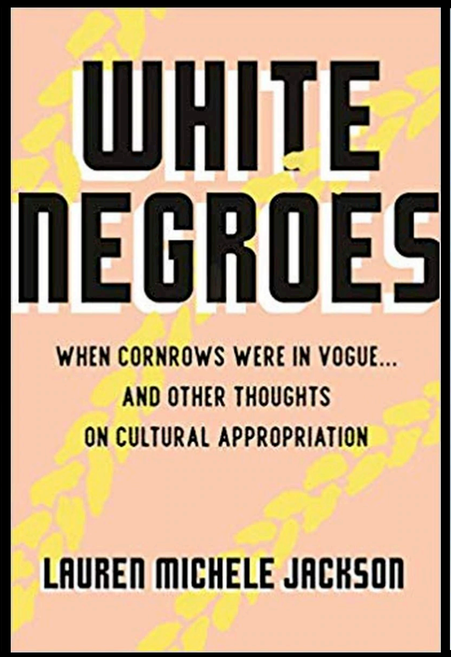

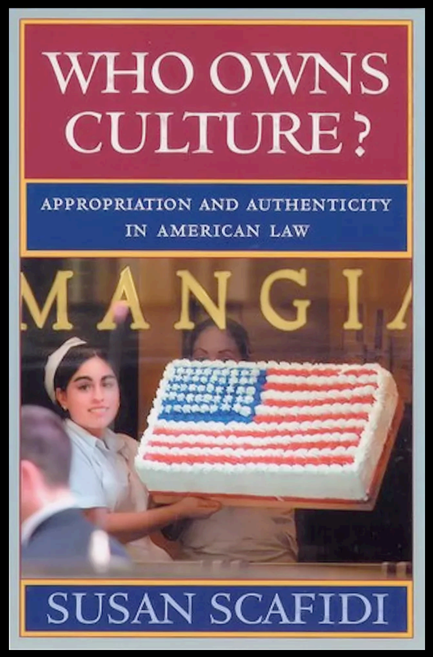

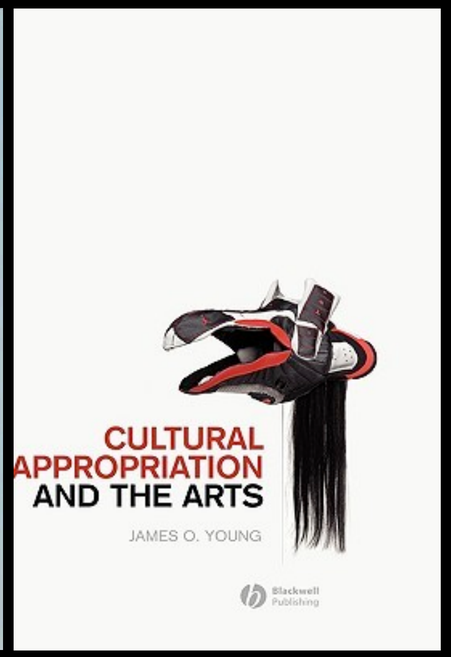

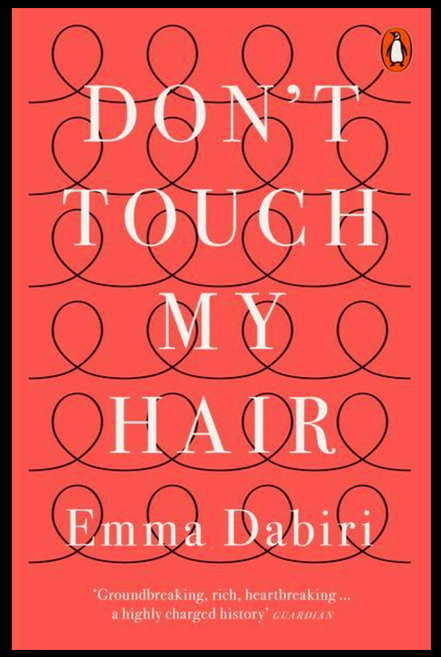

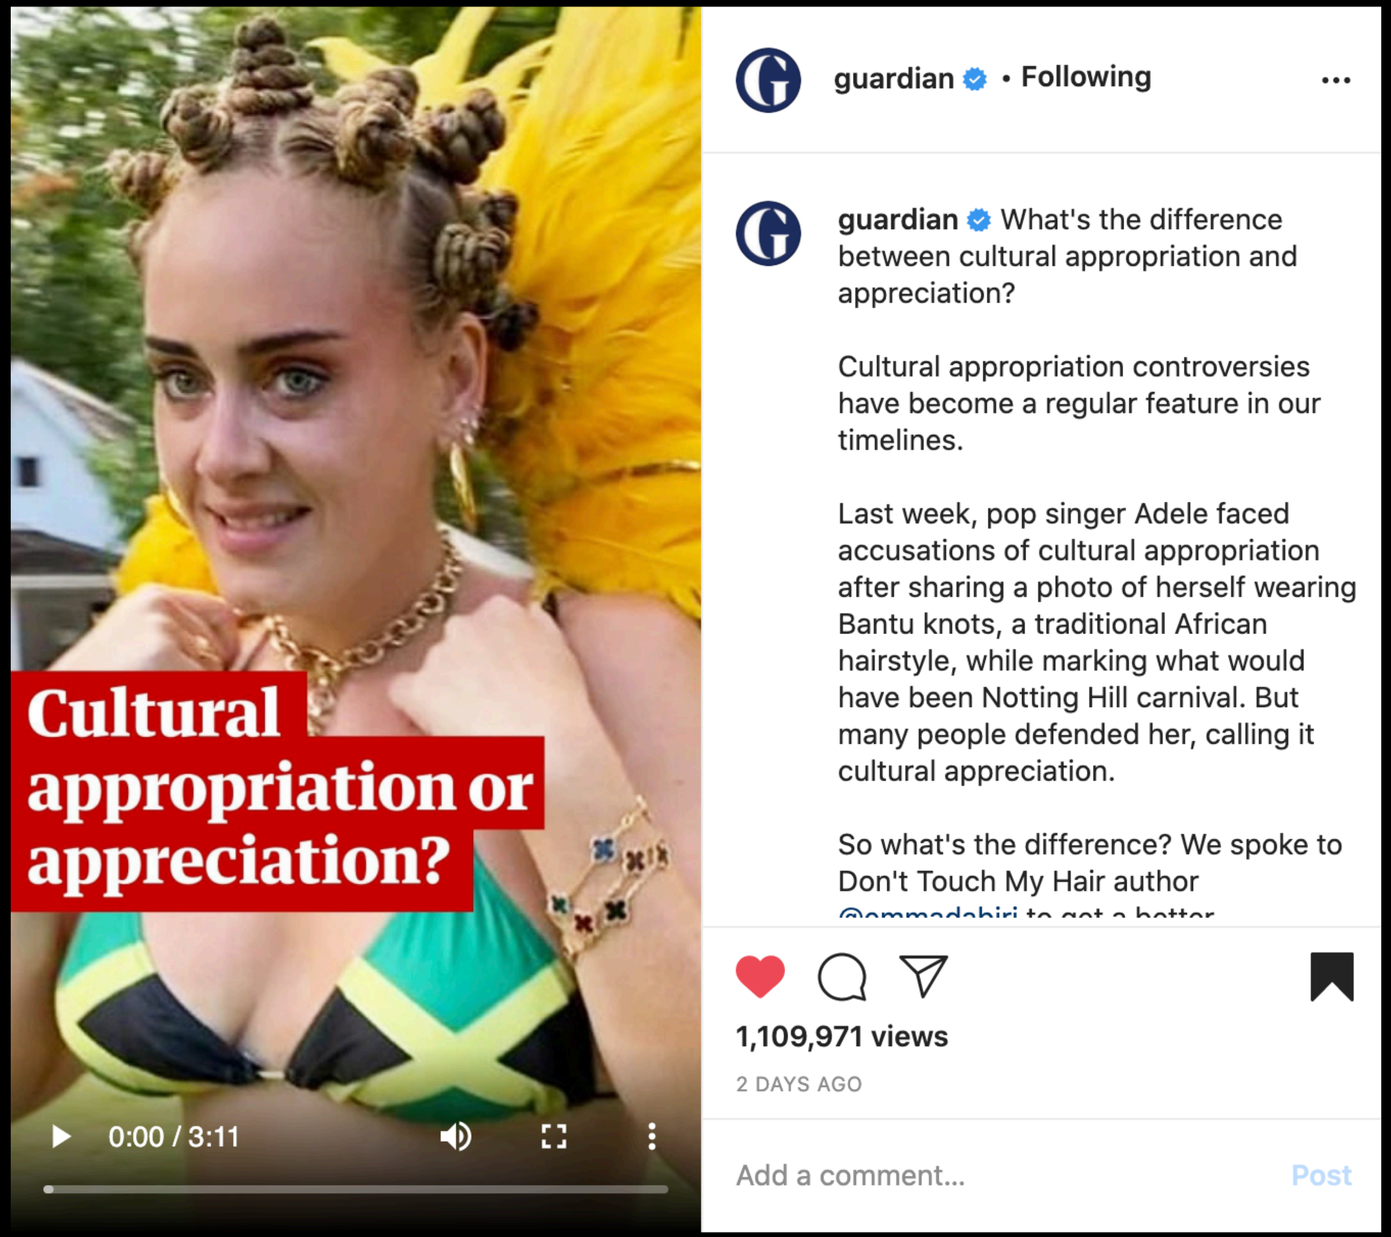
Cultural appropriation or appreciation?
0:00 / 3:11
guardian • Following
guardian What's the difference between cultural appropriation and appreciation?
Cultural appropriation controversies have become a regular feature in our timelines.
Last week, pop singer Adele faced accusations of cultural appropriation after sharing a photo of herself wearing Bantu knots, a traditional African hairstyle, while marking what would have been Notting Hill carnival. But many people defended her, calling it cultural appreciation.
So what's the difference? We spoke to Don't Touch My Hair author
1,109,971 views
2 DAYS AGO
Add a comment...
Post

What is cultural appropriation?

Madison Hummer, The Appropriation of African Objects in Pablo Picasso's Les Demoiselles d'Avignon. Trinity College, Hartford Connecticut, madison.hummer@trincoll.edu
https://digitalrepository.trincoll.edu/cgi/viewcontent.cgi?article=1760&context=theses

Michael Ochs Archives/Getty Images

https://www.instagram.com/shudu.gram/

https://www.instagram.com/lilmiquela/

https://www.instagram.com/shudu.gram/

CULTURAL FRAGMENTS

Still-Life with Chair Caning, 1912 by Pablo Picasso

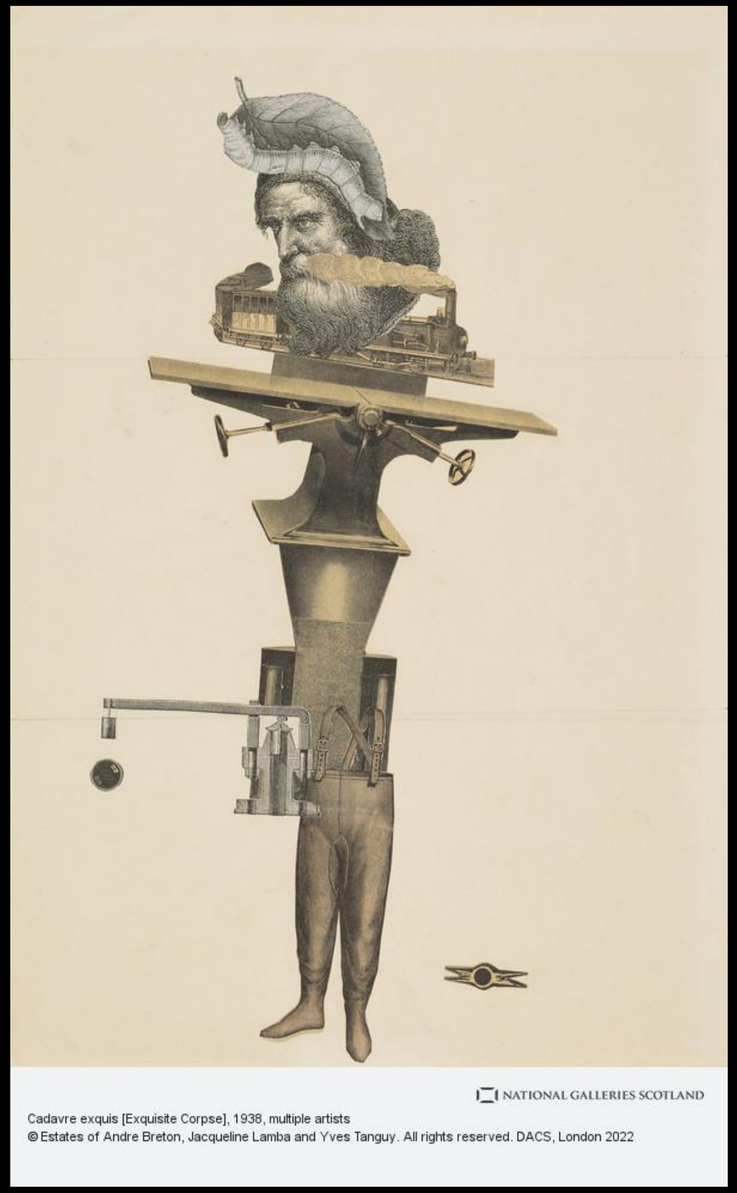

"Exquisite Corpse" - Andre Breton, Jacqueline Lamba and Yves Tanguy

To Make A Dadaist Poem

Take a newspaper.

Take some scissors.

Choose from this paper an article of the length you want to make your poem.

Cut out the article.

Next carefully cut out each of the words that make up this article and

put them all in a bag.

Shake gently.

Next take out each cutting one after the other.

Copy conscientiously in the order in which they left the bag.

The poem will resemble you.

And there you are — an infinitely original author of charming sensibility

even though unappreciated by the vulgar herd.

http://www.remixthebook.com/

Daphne Oram Trust / Goldsmiths College, University of London

Pierre Schaeffer - "Musique concrète"

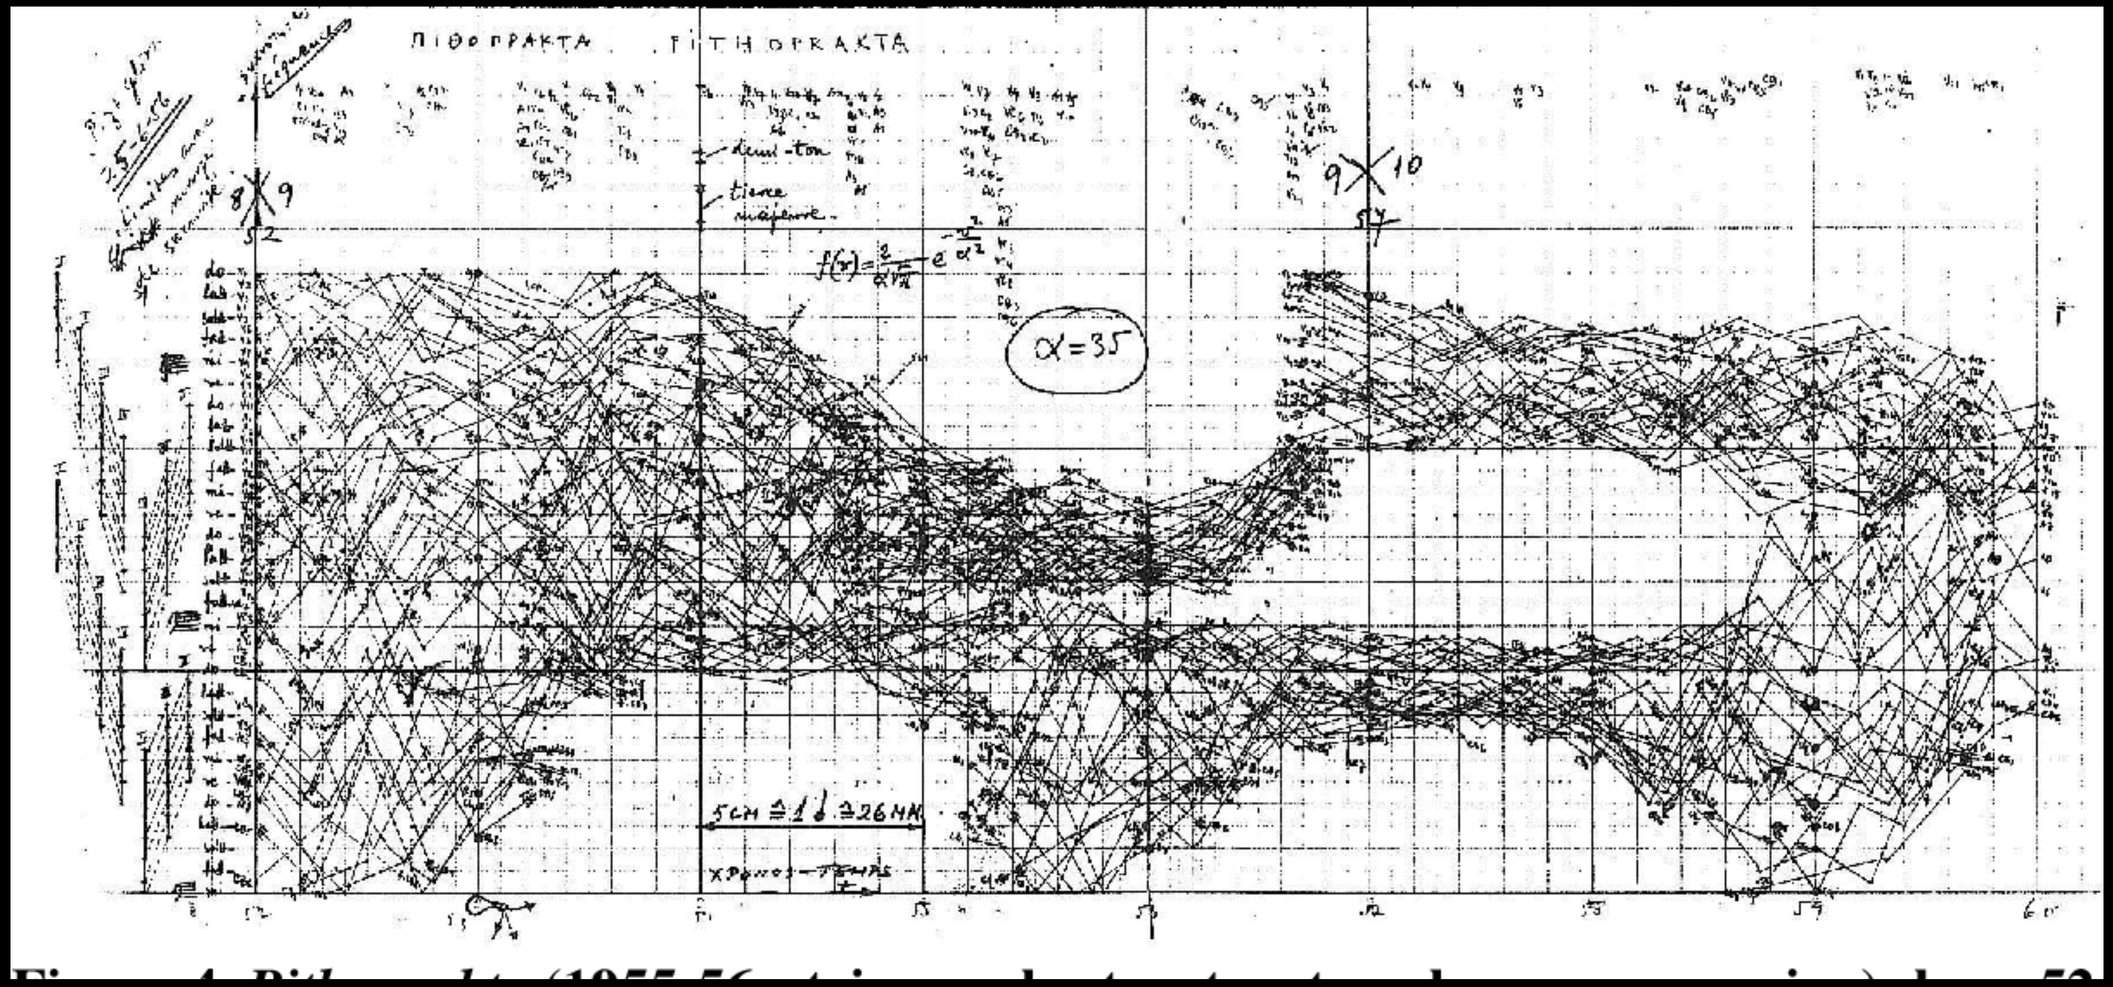

Iannis Xenakis - Granular Synthesis

Where can we see "cultural copies" or "borrowing"?

Where can we see "cultural copies" or "borrowing"?

- Collage
- Cosmetics
- Dummies
- Design
- Impersonations
- Montage
- Re-enactments
- Simulations
- Quotation
- Plagiarism
- Parody
- Fashion
- Forgery
- Sampling
- Pastiche
- Remix
- Art
- Music
- Poetry
- Books
- Scripts
- Genes
- Memes
- Machine Learning...?

MACHINE LEARNING

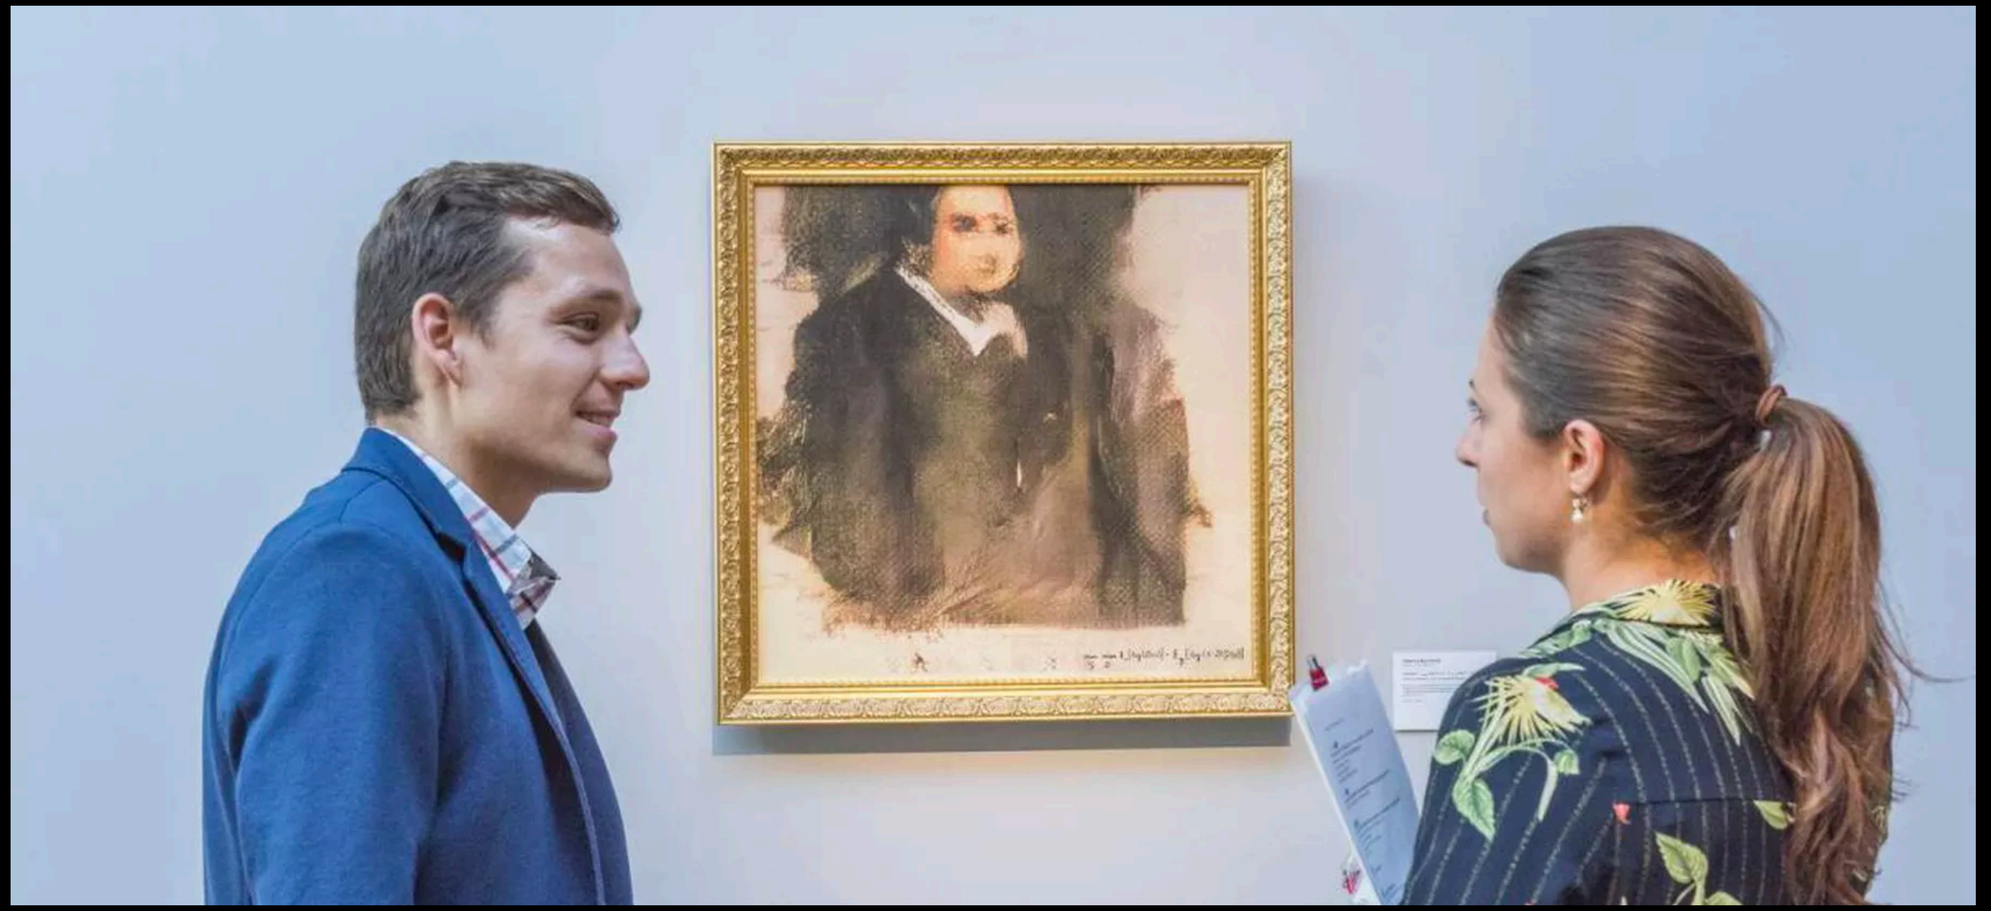

https://www.theverge.com/2018/10/25/18023266/ai-art-portrait-christies-obvious-sold

https://www.barbican.org.uk/whats-on/2019/event/ai-more-than-human

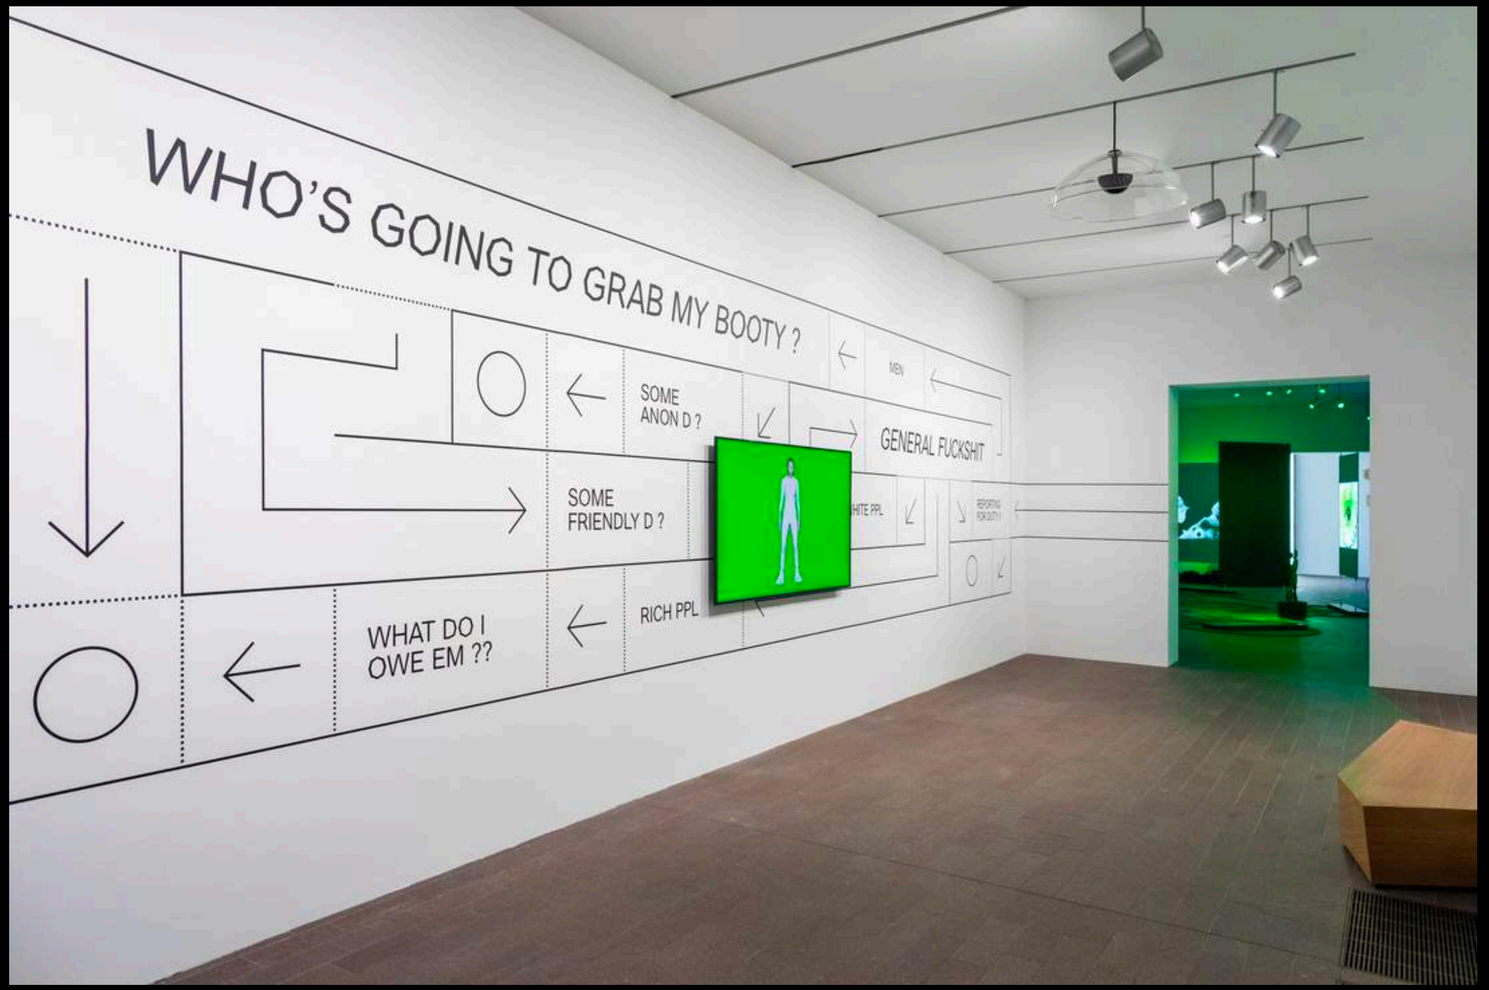

https://deyoung.famsf.org/martine-syms-threat-model-mythiccbeing

https://deyoung.famsf.org/uncanny-valley-being-human-age-ai

- Big datasets (Everything is digital media these days)
- Big computers (Server farms; GPU computing; TPUs)
- Lots of money and time to compute (Massive industry backing and adoption)

Where does machine learning make decisions today?

- Mortgage risk evaluation
- Insurance risk
- Profiling / policing
- Facial surveillance and recognition
- Recommendation / Personalization
- Driving / navigation
- Factory sorting
- Games / Agents
- Drug discovery / sequencing
- Prostheses
- Indexing / Search
- Optimization / Routing
- Industrial / Military robotics
- Voice assistants
- Advertising / Targeting
- Bots / Spam / Marketing
- Translation
- Art?
- Language / Poetry?
- Sound / Music?

How does Machine Learning make decisions?

- Supervised learning: learn to classify this thing into some box/label
- Unsupervised learning: learn the density of the data, primarily what we will be looking at in this course
- Semi-supervised learning: combination of the two
- Reinforcement learning: agency, robots, control, agency / automation in games e.g.
- Other kinds: various optimization methods

ARTIFICIAL INTELLIGENCE

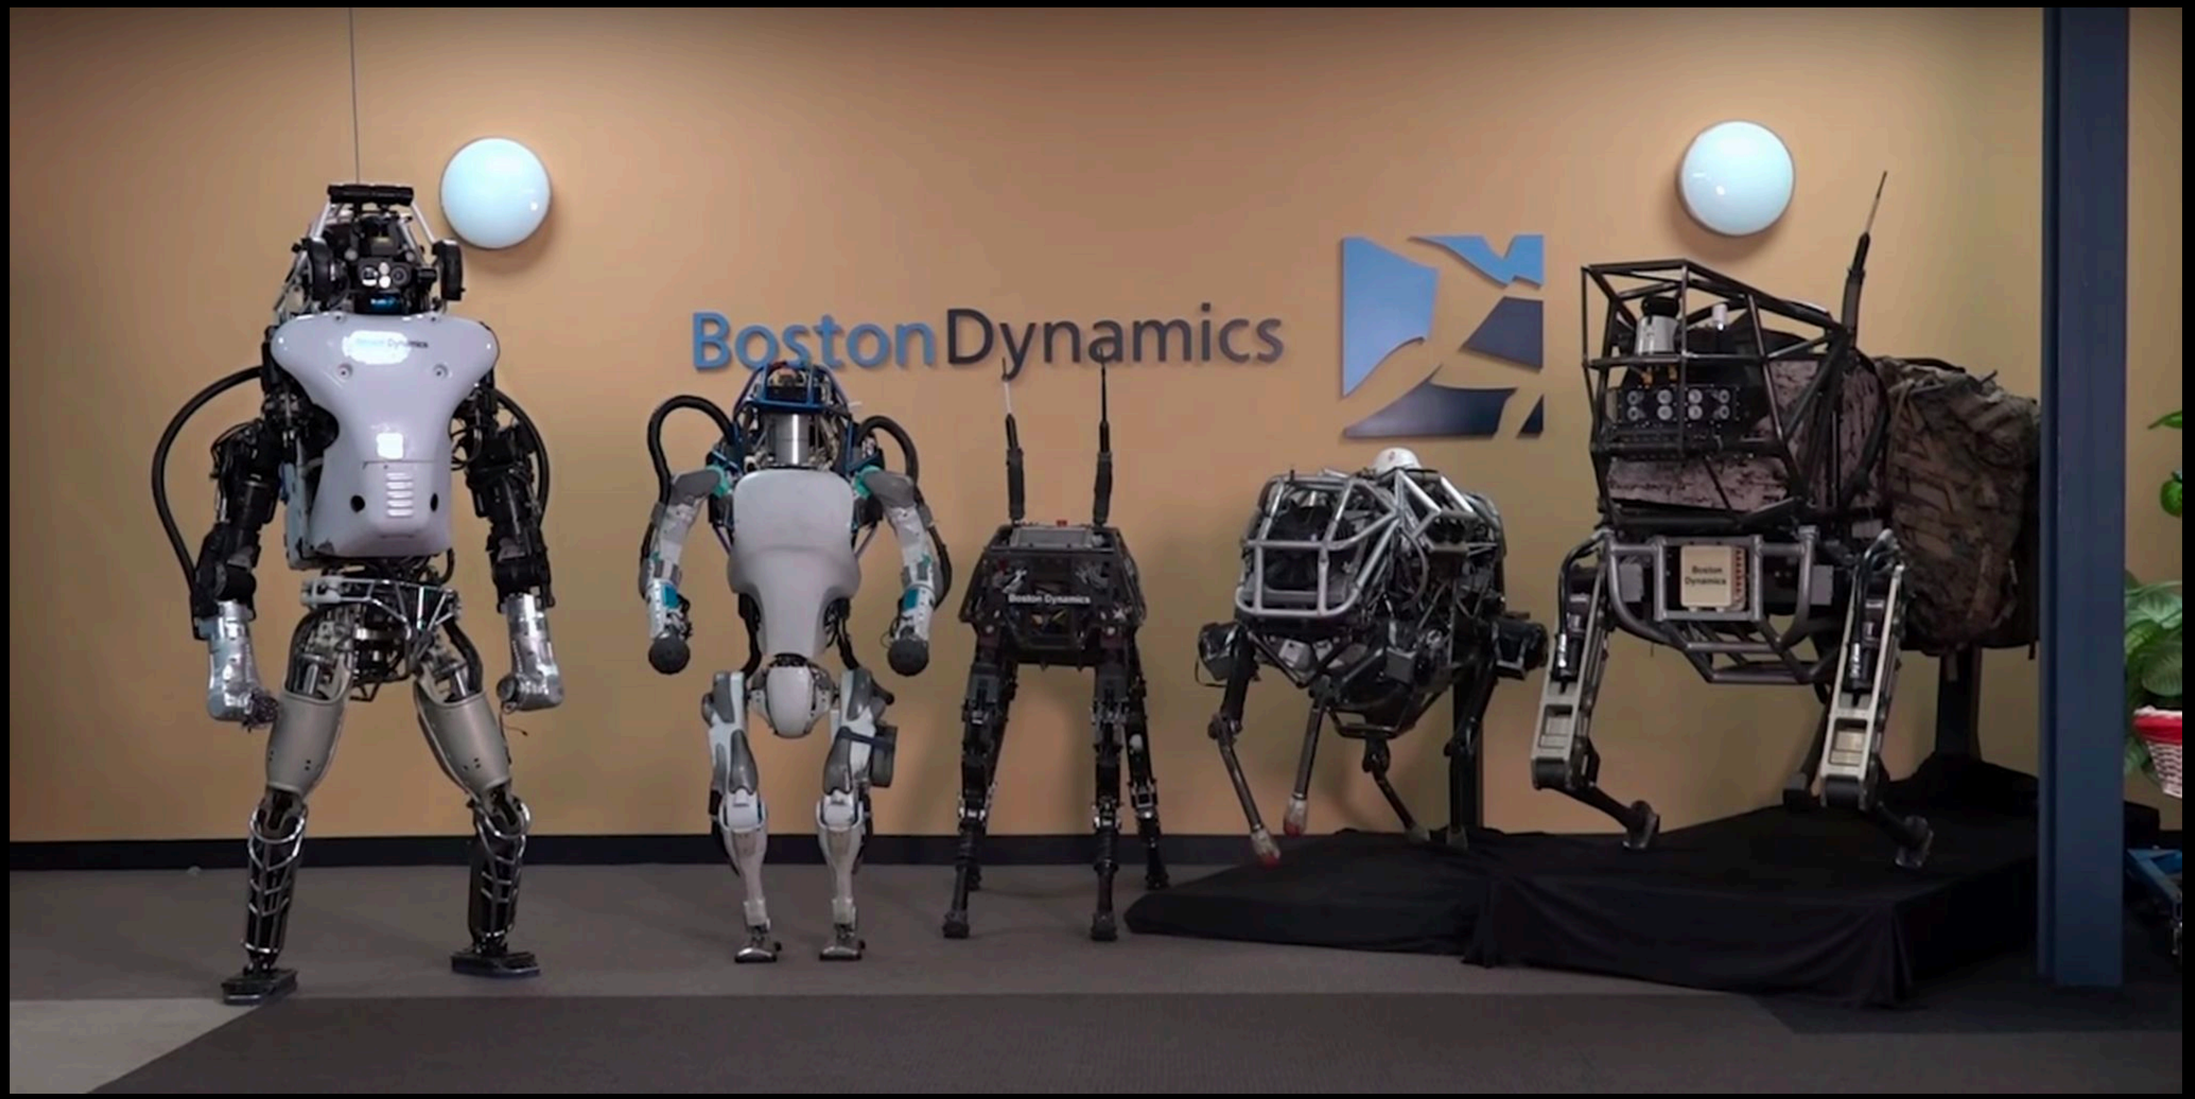
BostonDynamics

naturalmotion

facebook Ad Creative

.SKETCH

Mobile News Feed

Suggested App

Your App for Business
Sponsored ·

Explore stories from friends and the world around you.

Paper
Available on the App Store

Install Now

10 1 Comment

Like Like Share

Desktop Right Column

Your App for Business
Business
Use our app and get delicious recipes for Jasper's Market's latest produce. It's fun, easy...

Install Now · 37,235 persons using this

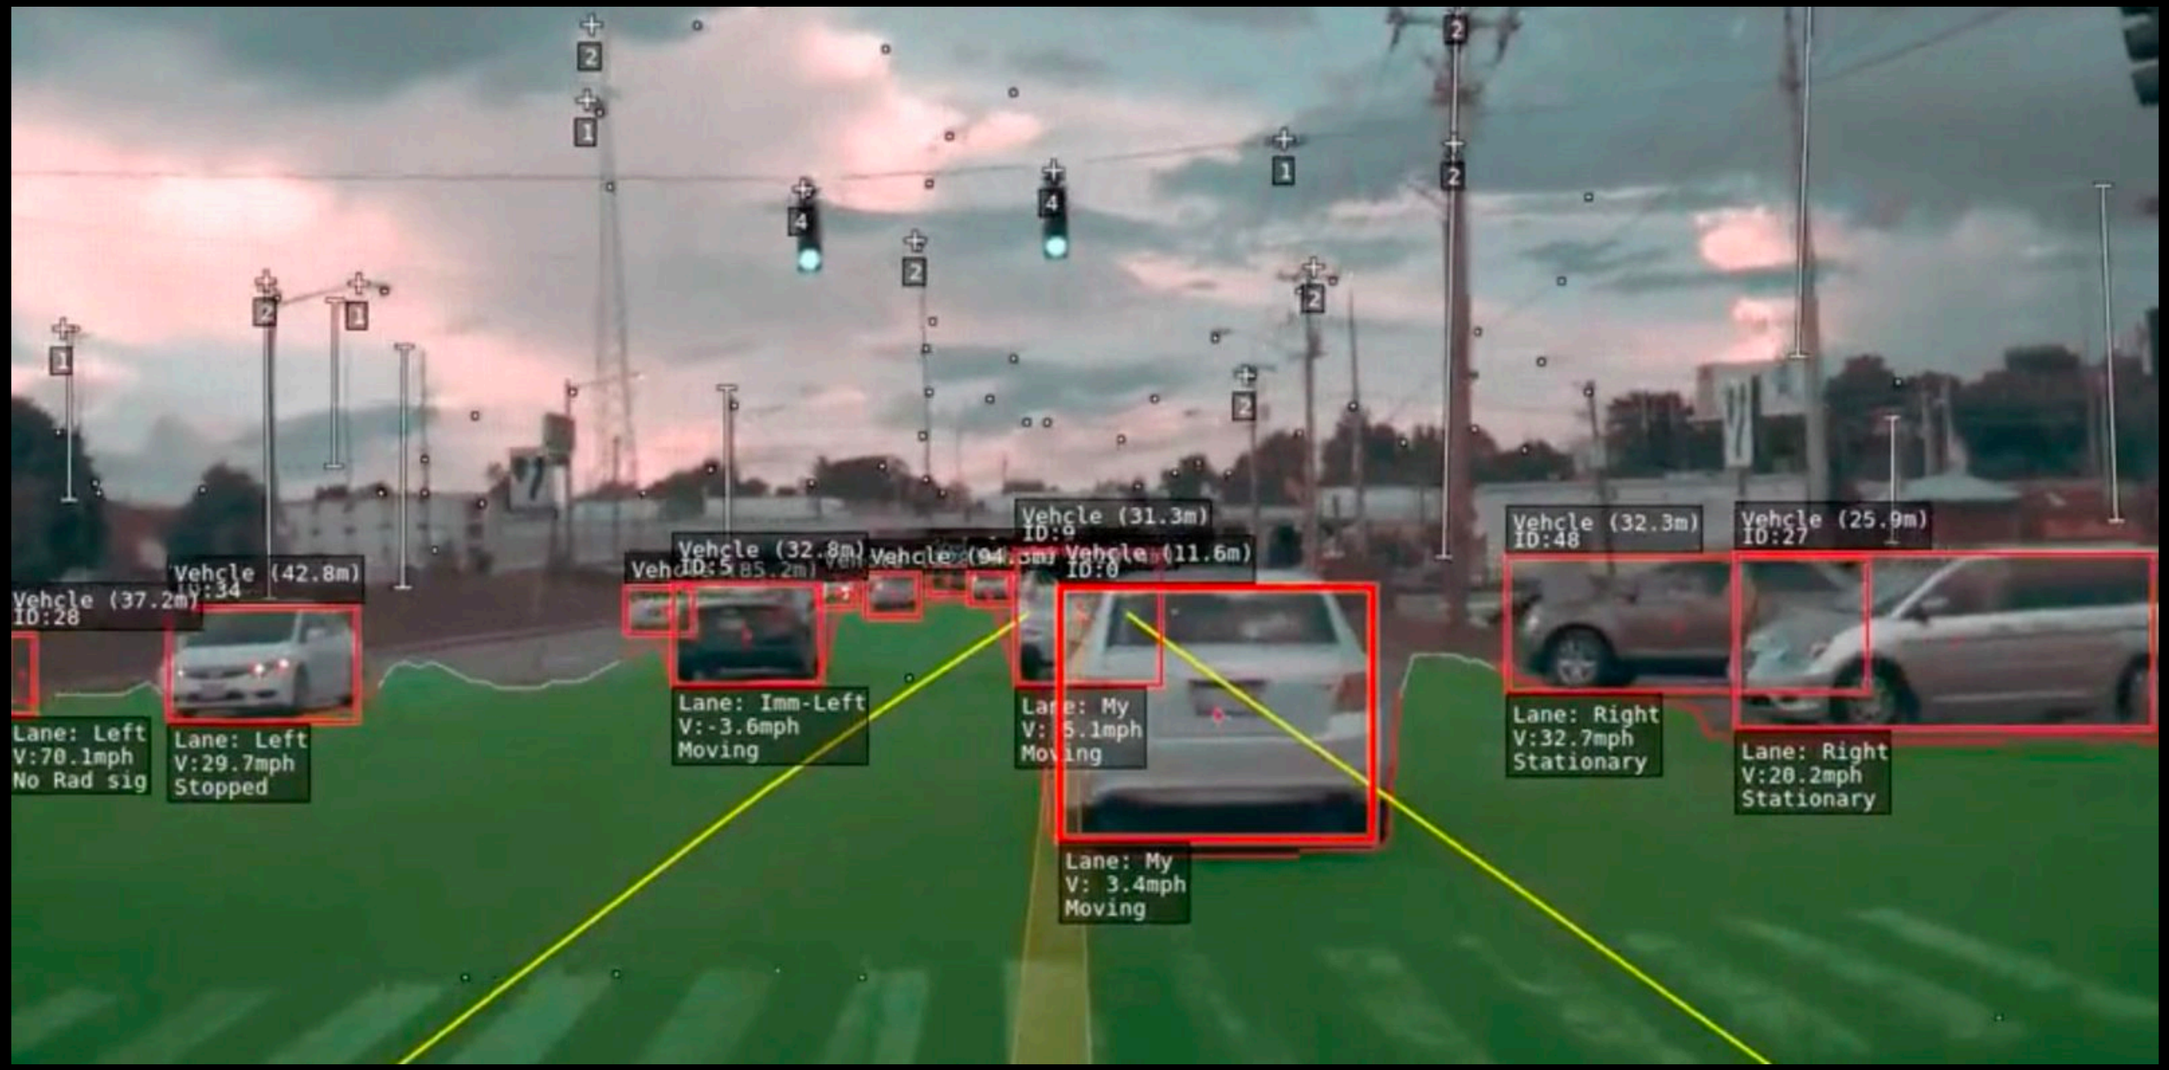
Vehcle (37.2m)
ID:28
Lane: Left
V:70.1mph
No Rad sig
Vehcle (42.8m)
ID:34
Lane: Left
V:29.7mph
Stopped
Vehcle (32.8m)
ID:5
Lane: Imm-Left
V:-3.6mph
Moving
Vehcle (31.3m)
ID:9
Vehcle (11.6m)
ID:0
Lane: My
V: 3.4mph
Moving
Vehcle (32.3m)
ID:48
Lane: Right
V:32.7mph
Stationary
Vehcle (25.9m)
ID:27
Lane: Right
V:20.2mph
Stationary

VOICE ASSISTANTS	GOOGLE ASSISTANT	AMAZON Alexa	SAMSUNG Bixby	MICROSOFT Cortana	APPLE Siri
DATA SOURCE The database where the voice assistant gets its data from	Answer Boxes (Google Search), Google "Knowledge Graph"	Bing Answer Boxes, Bing "Knowledge Graph"	"Powered by Google"	Bing Answer Boxes, Bing "Knowledge Graph"	Wikipedia, Wolfram, Alpha, Google Search

https://arteric.com/our-thoughts/structured-data-voice-driven-search

DEEP LEARNING

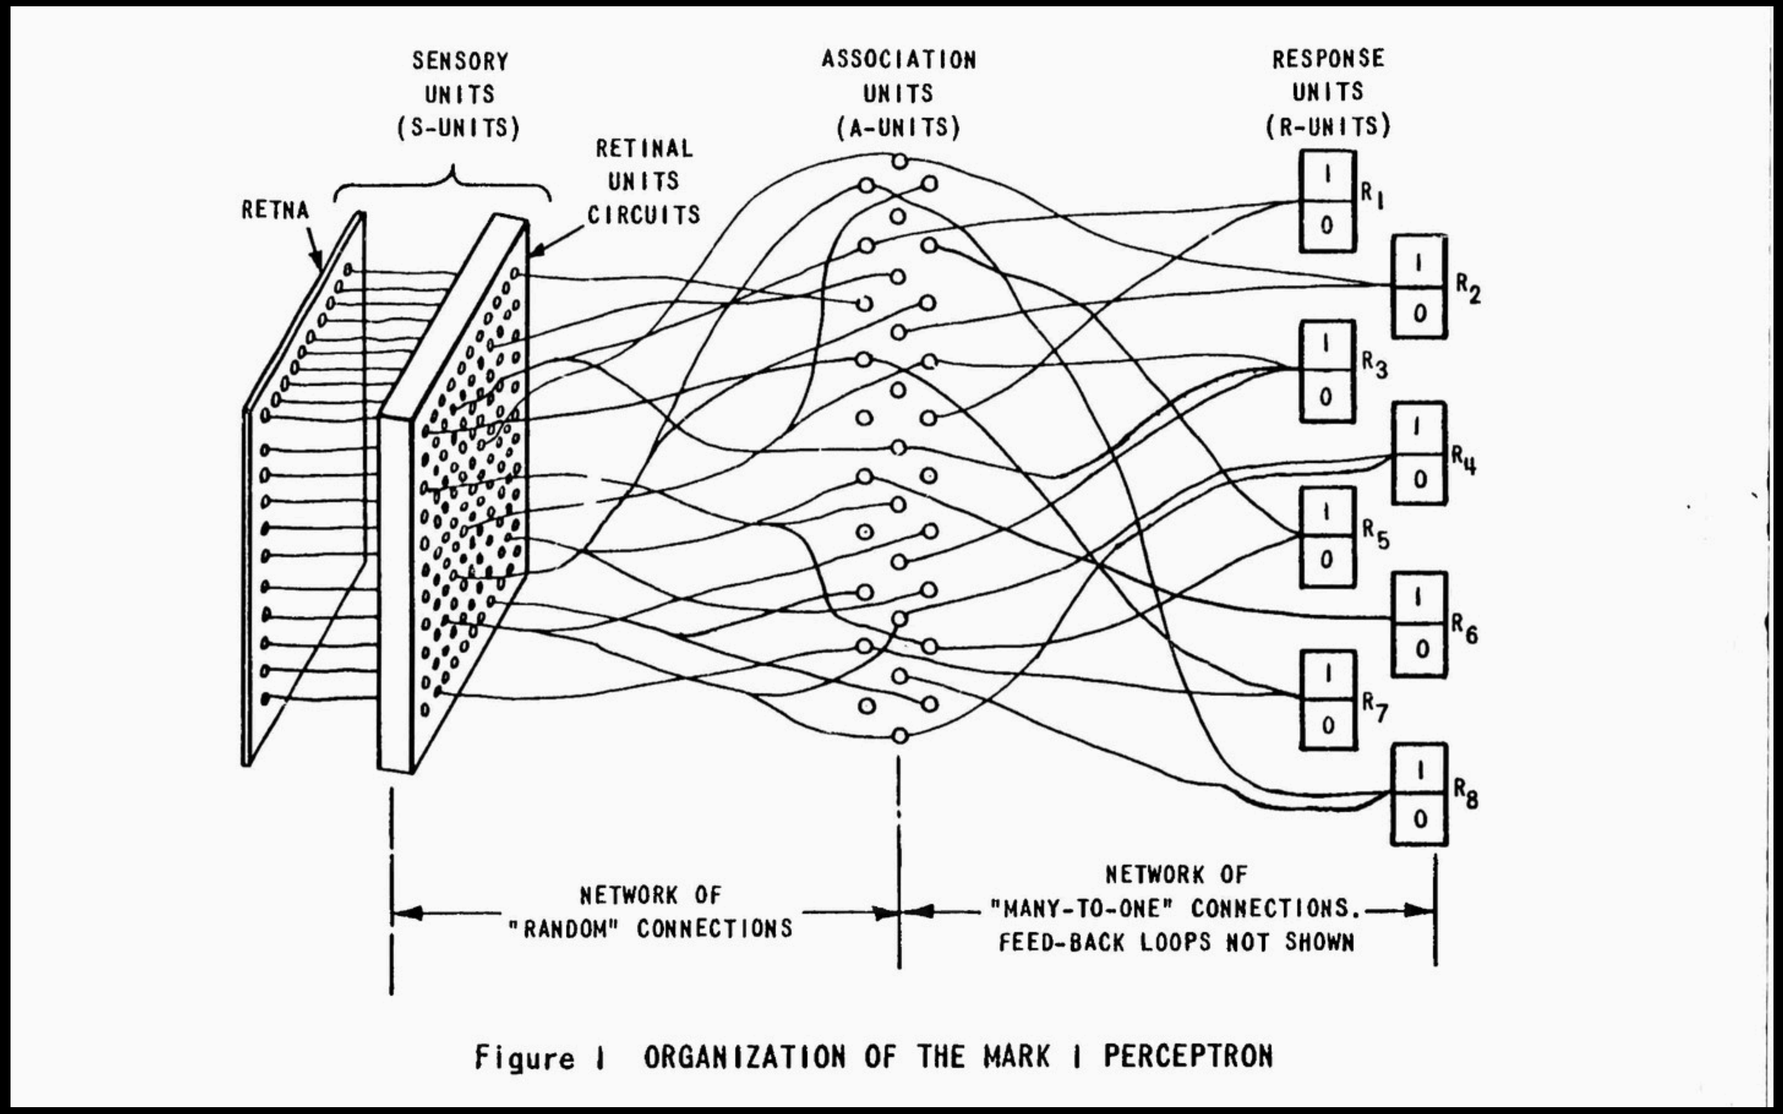

Illustration from Frank Rosenblatt, *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms,* (Cornell Aeronautical Laboratory, Buffalo NY, 1961). Retrieved from *https://www.e-flux.com/journal/101/273221/three-thousand-years-of-algorithmic-rituals-the-emergence-of-ai-from-the-computation-of-space/*

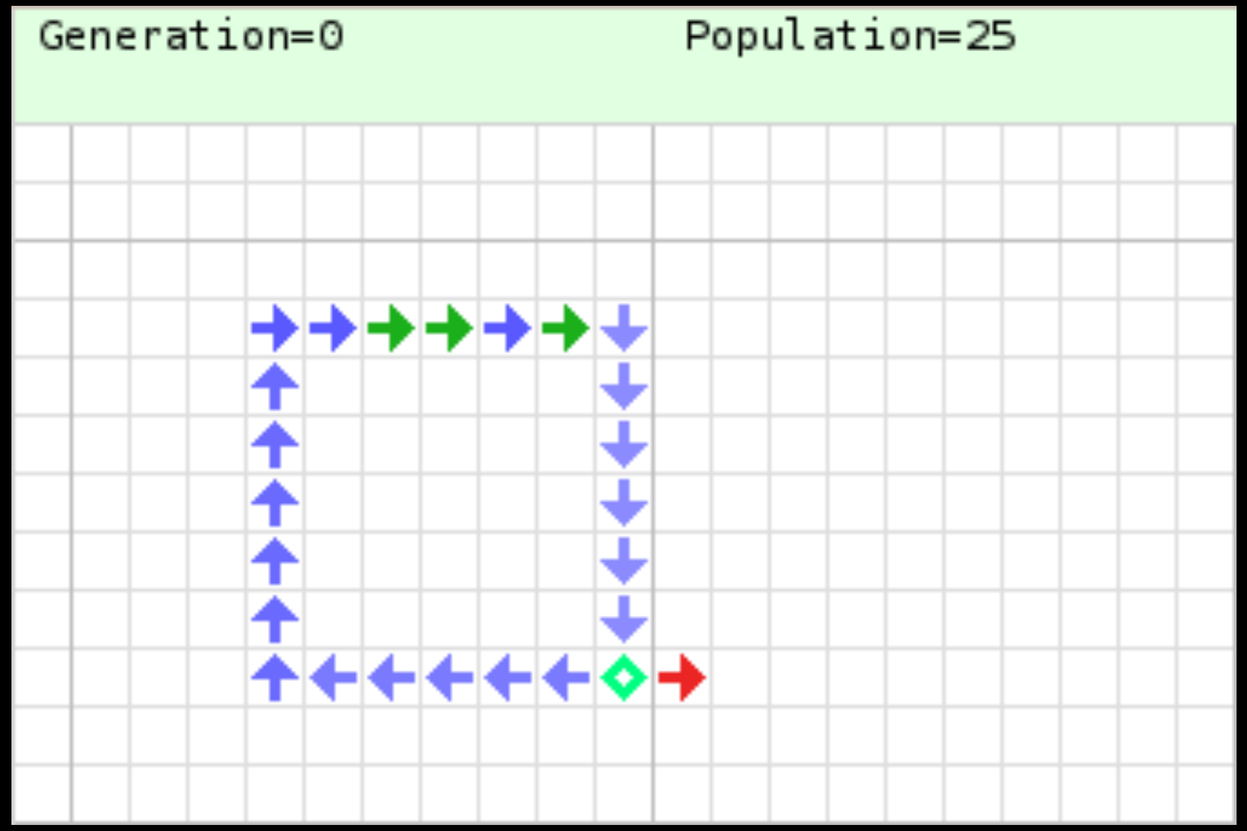

John von Neumann's cellular automata (wikipedia)

Konrad Zuse - Rechnender Raum painting

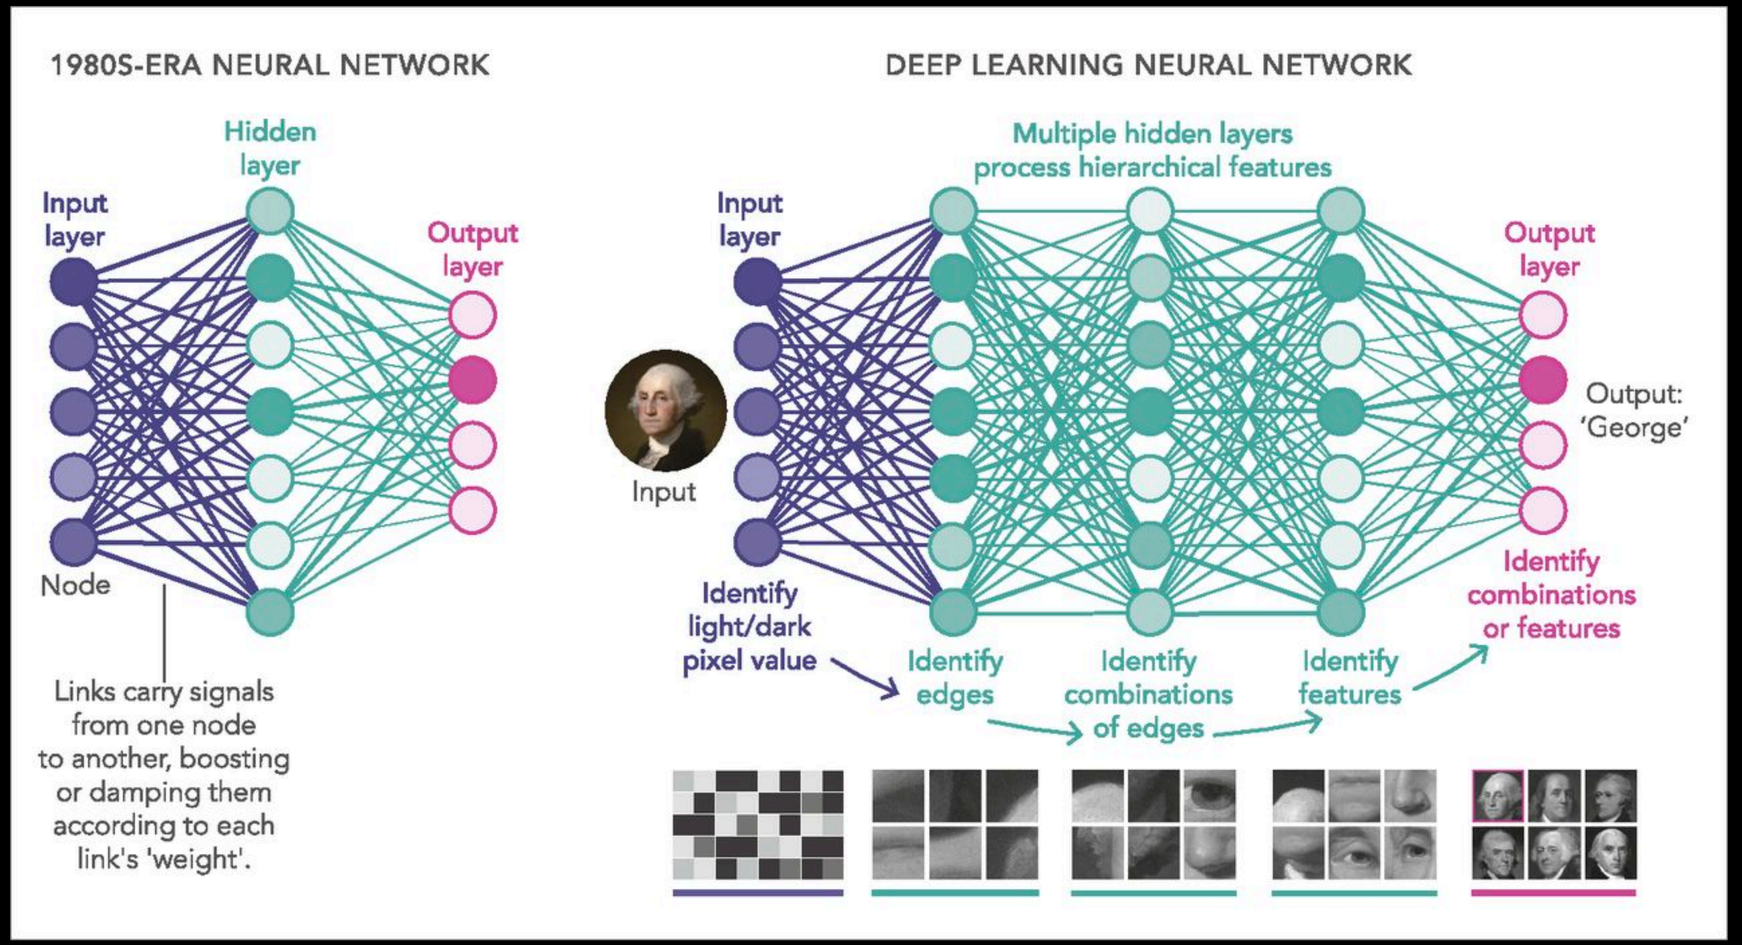

https://www.pnas.org/content/116/4/1074

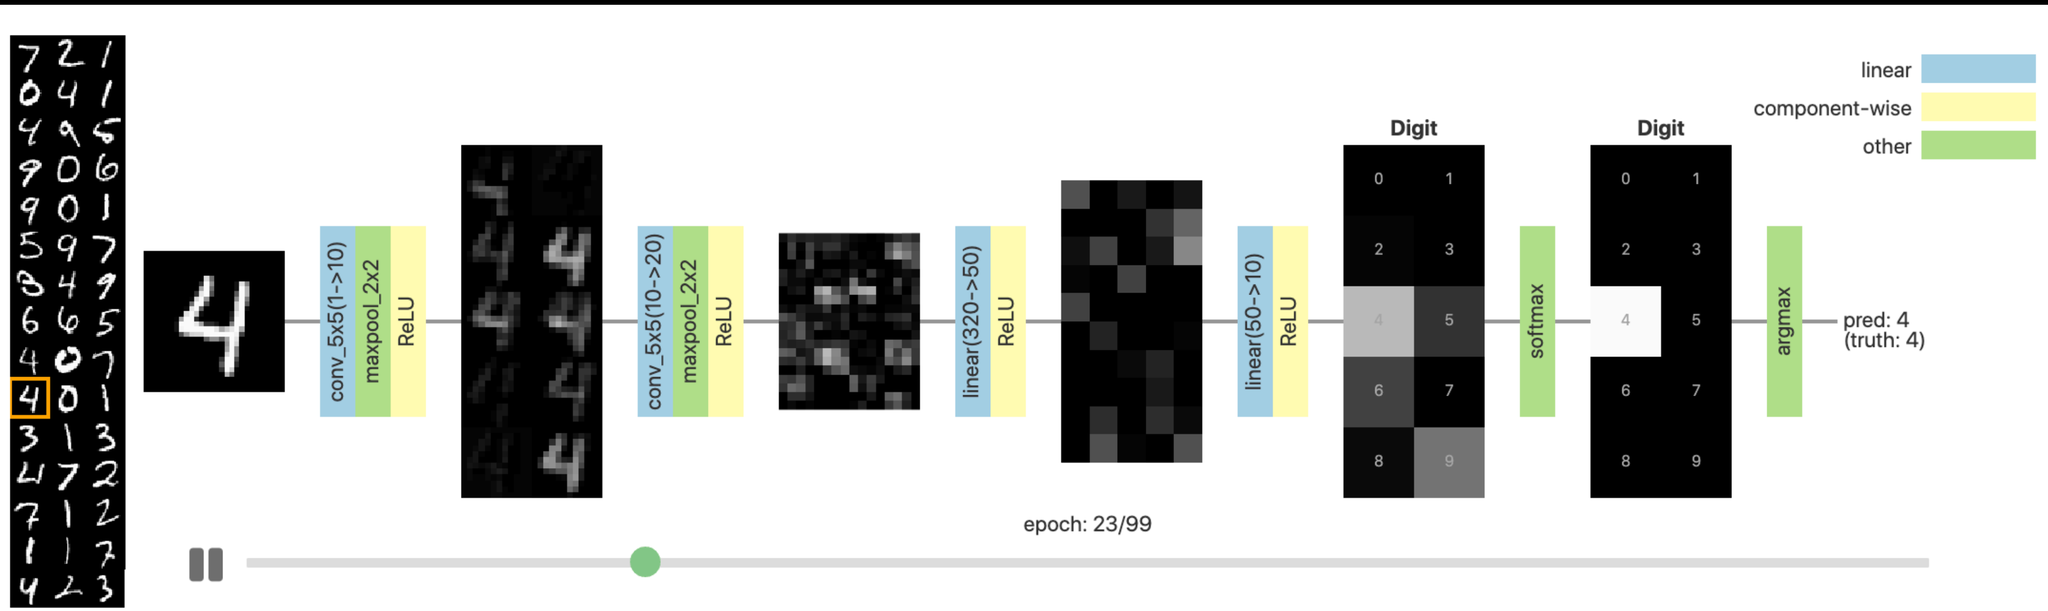

Neural network opened. The colored blocks are building-block functions (i.e. neural network layers), the gray-scale heatmaps are either the input image or intermediate activation vectors after some layers.

https://distill.pub/2020/grand-tour/

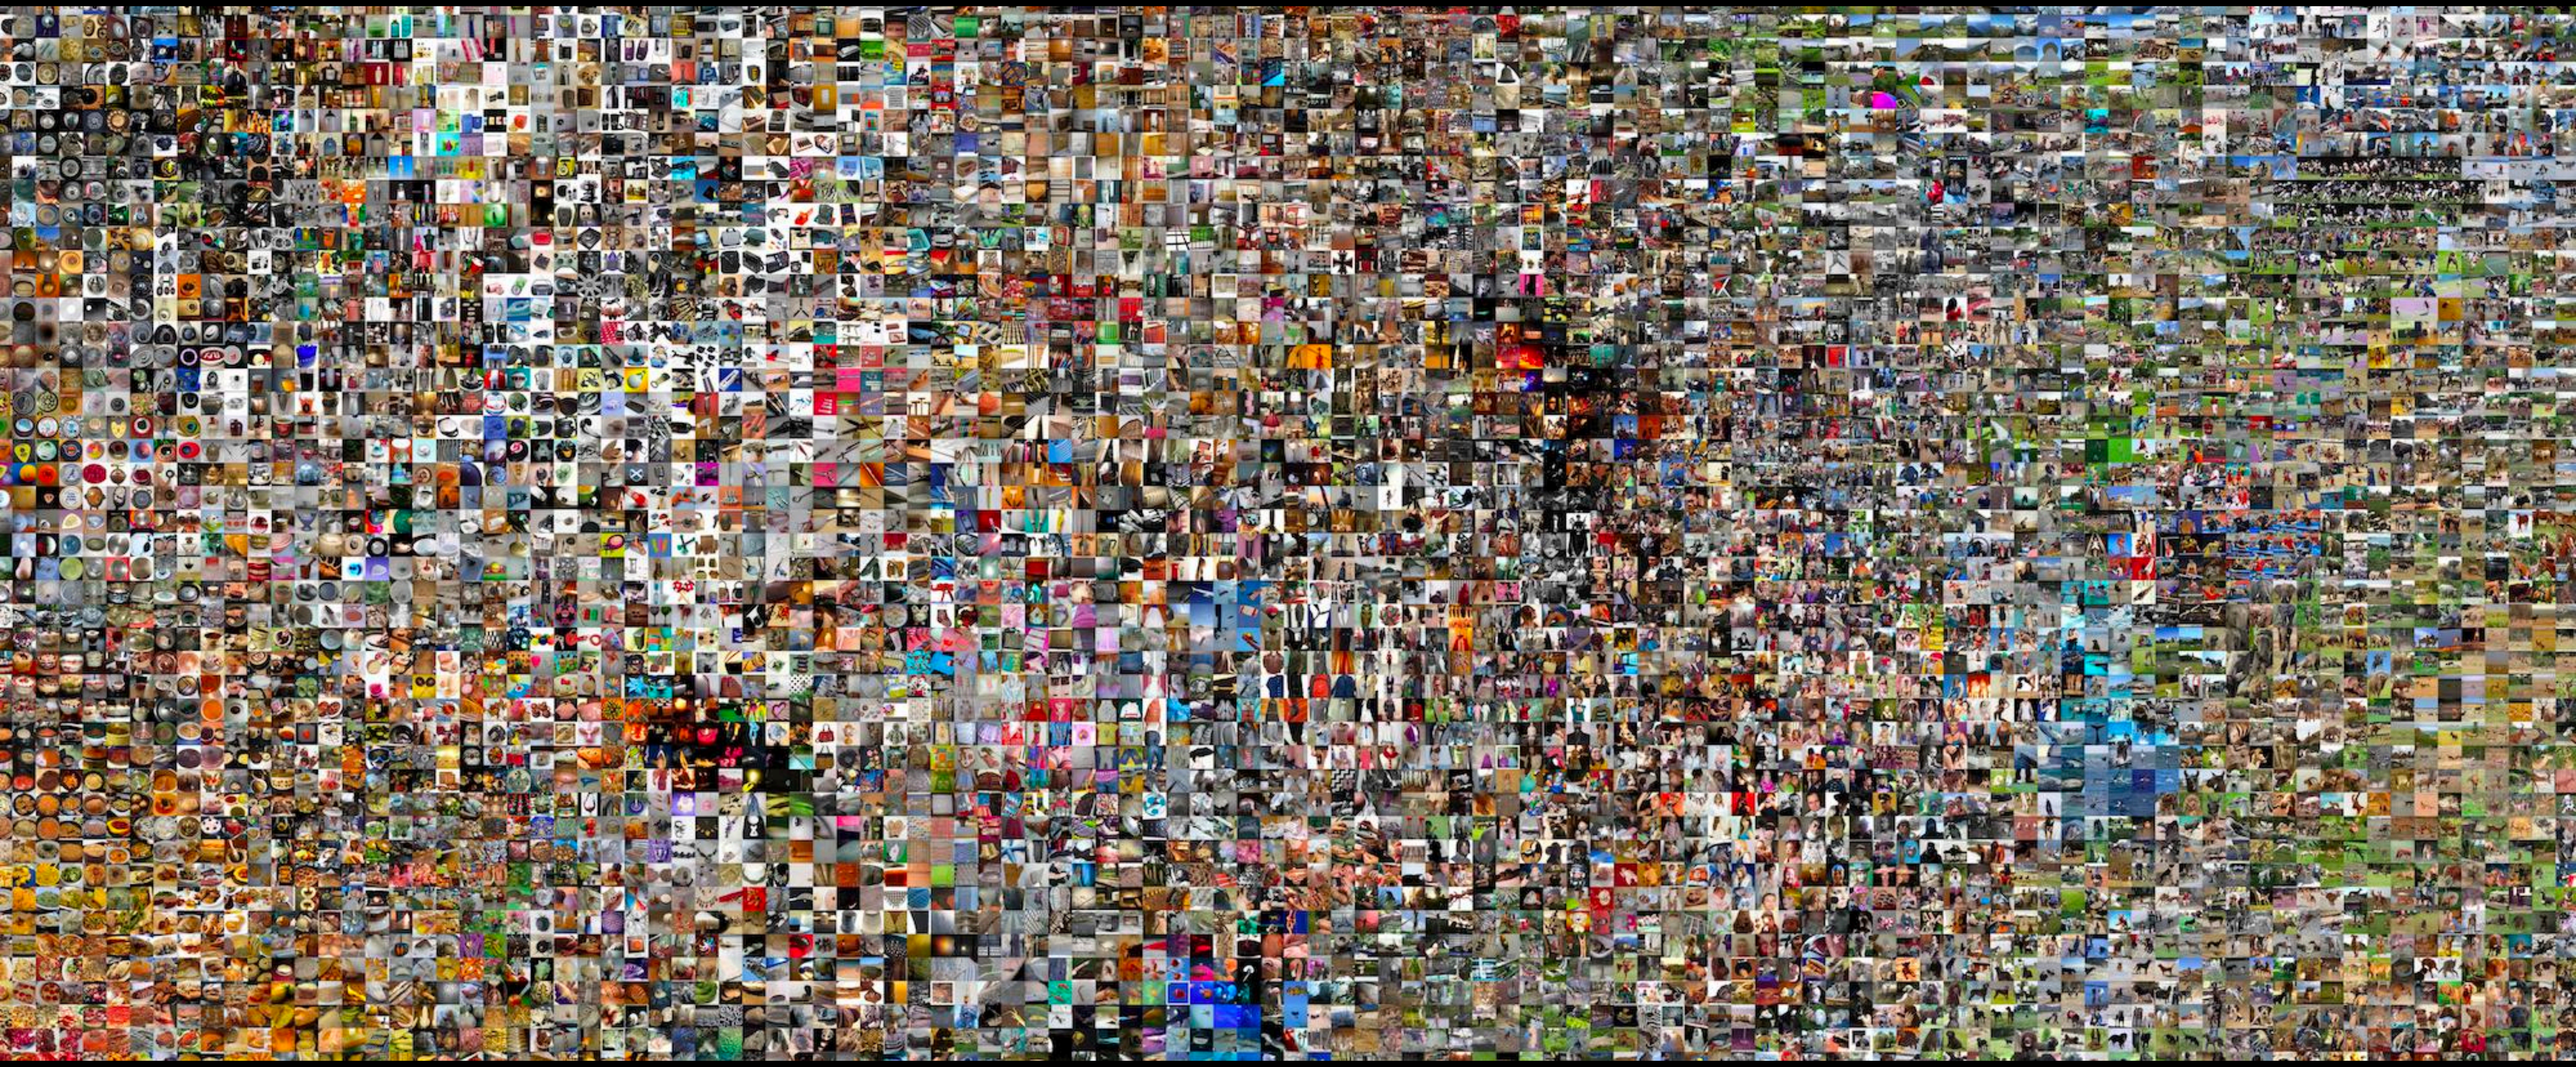

ImageNet

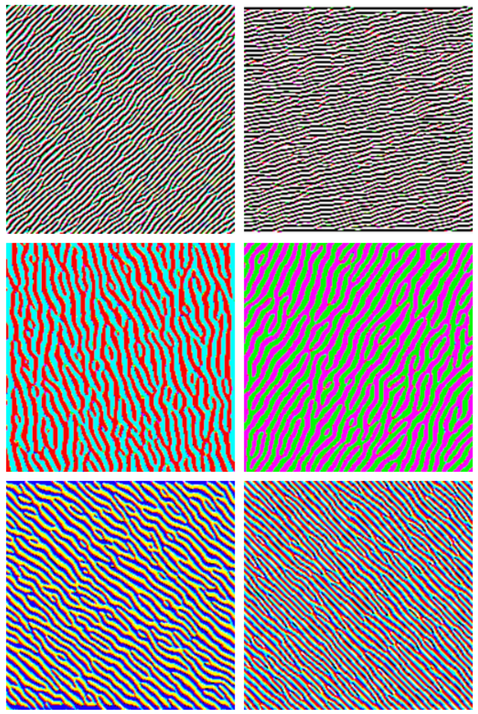

Edges (layer conv2d0)

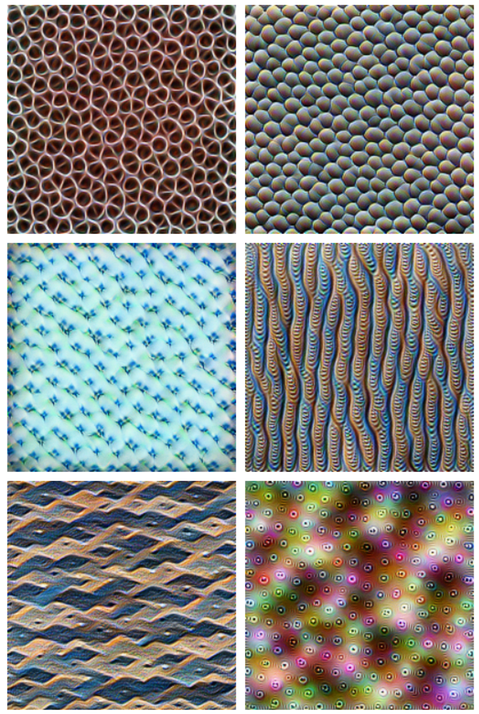

Textures (layer mixed3a)

Patterns (layer mixed4a)

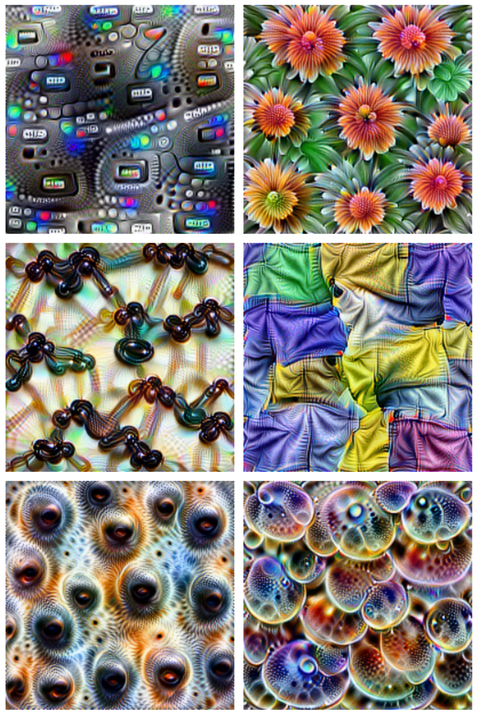

Parts (layers mixed4b & mixed4c)

Objects (layers mixed4d & mixed4e)

https://distill.pub/2017/feature-visualization/

Dataset Examples show us what neurons respond to in practice

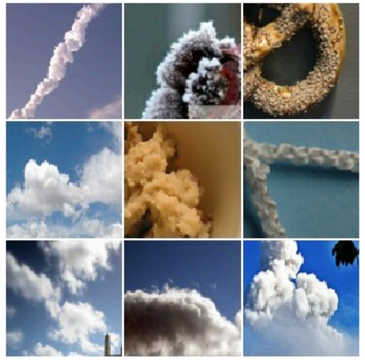

Optimization isolates the causes of behavior from mere correlations. A neuron may not be detecting what you initially thought.

Baseball—or stripes?
mixed4a, Unit 6

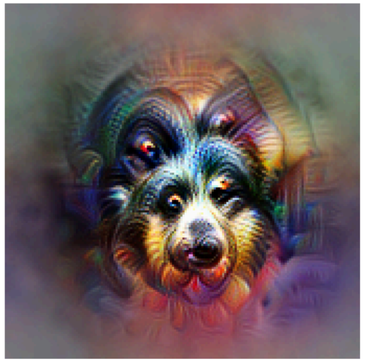

Animal faces—or snouts?
mixed4a, Unit 240

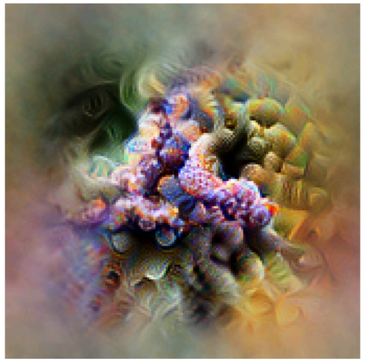

Clouds—or fluffiness?
mixed4a, Unit 453

Buildings—or sky?
mixed4a, Unit 492

https://distill.pub/2017/feature-visualization/

ROBOTICS

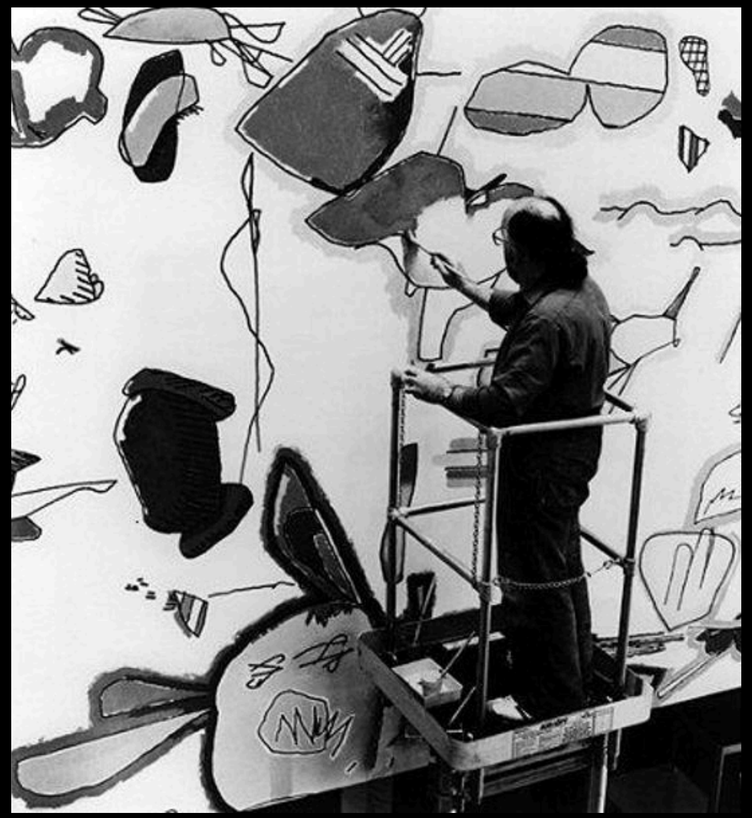

Harold Cohen coloring the forms produced by the AARON drawing "Turtle" at the Computer Museum, Boston, MA, ca. 1982. Collection of the Computer History Museum, 102627459.

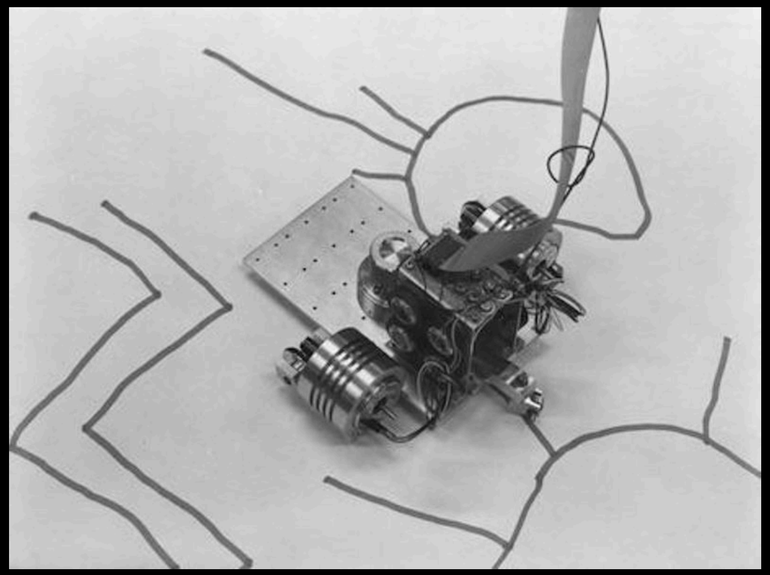

The 1979 exhibition, Drawings, at SFMOMA, featured this "turtle" robot creating drawings in the gallery. Collection of the Computer History Museum, 102627449.

Patrick Tresset - "Paul"

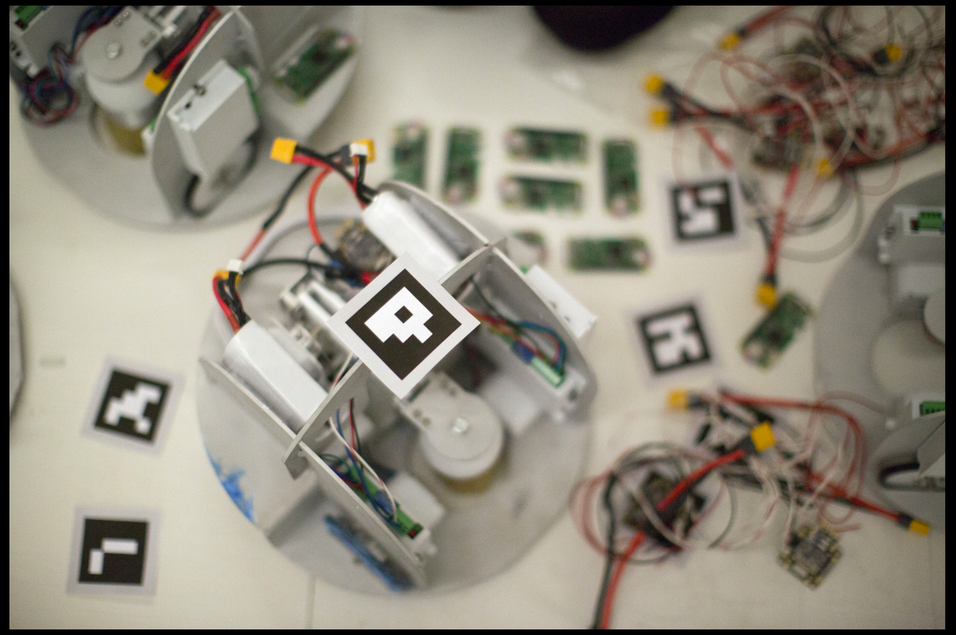

Sougwen Chung

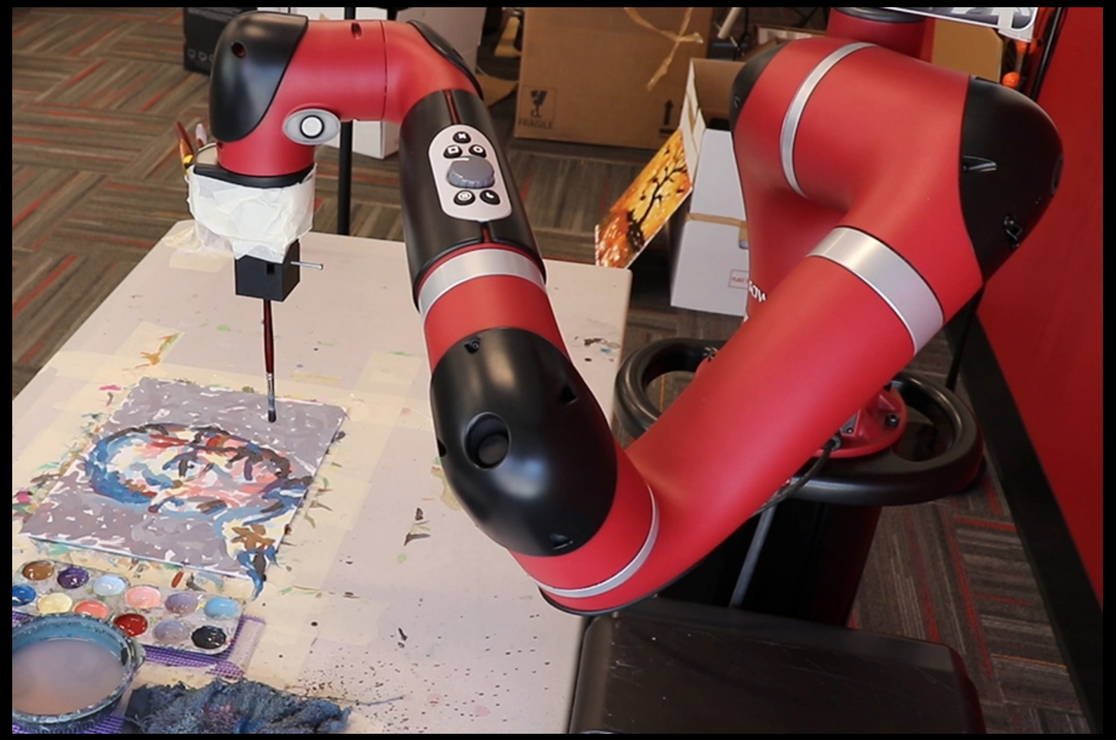

Carnegie Mellon University's Robotics Institute - "Frida"

DATA

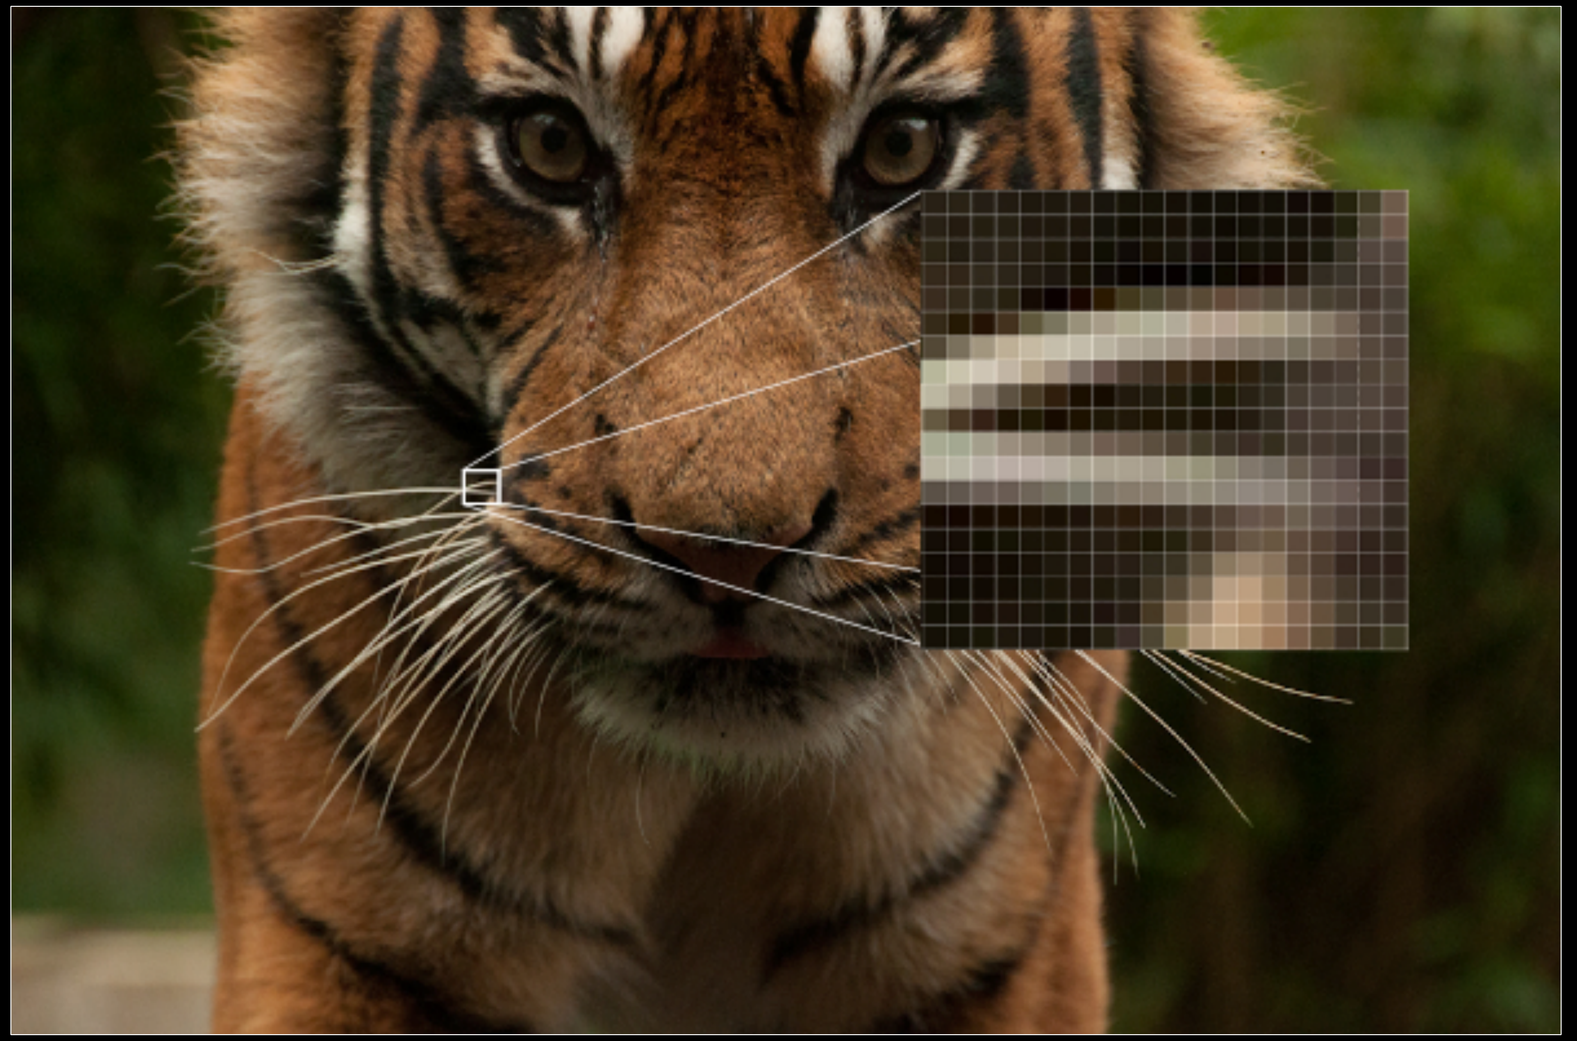

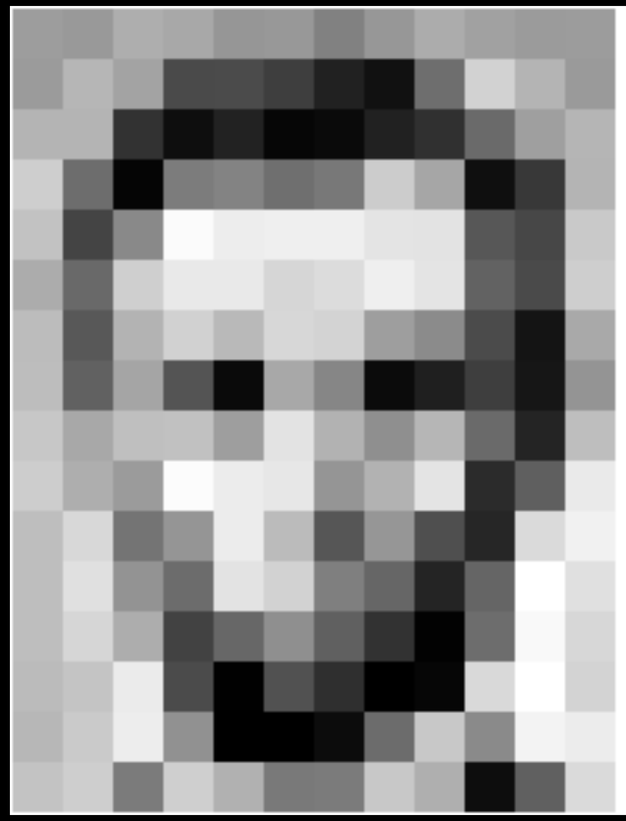

157	153	174	168	150	152	129	151	172	161	155	156
155	182	163	74	75	62	33	17	110	210	180	154
180	180	50	14	34	6	10	33	48	106	159	181
206	109	5	124	131	111	120	204	166	15	56	180
194	68	137	251	237	239	239	228	227	87	71	201
172	105	207	233	233	214	220	239	228	98	74	206
188	88	179	209	185	215	211	158	139	75	20	169
189	97	165	84	10	168	134	11	31	62	22	148
199	168	191	193	158	227	178	143	182	106	36	190
206	174	155	252	236	231	149	178	228	43	95	234
190	216	116	149	236	187	86	150	79	38	218	241
190	224	147	108	227	210	127	102	36	101	255	224
190	214	173	66	103	143	96	50	2	109	249	215
187	196	235	75	1	81	47	0	6	217	255	211
183	202	237	145	0	0	12	108	200	138	243	236
195	206	123	207	177	121	123	200	175	13	96	218

157	153	174	168	150	152	129	151	172	161	155	156
155	182	163	74	75	62	33	17	110	210	180	154
180	180	50	14	34	6	10	33	48	106	159	181
206	109	5	124	131	111	120	204	166	15	56	180
194	68	137	251	237	239	239	228	227	87	71	201
172	105	207	233	233	214	220	239	228	98	74	206
188	88	179	209	185	215	211	158	139	75	20	169
189	97	165	84	10	168	134	11	31	62	22	148
199	168	191	193	158	227	178	143	182	106	36	190
206	174	155	252	236	231	149	178	228	43	95	234
190	216	116	149	236	187	86	150	79	38	218	241
190	224	147	108	227	210	127	102	36	101	255	224
190	214	173	66	103	143	96	50	2	109	249	215
187	196	235	75	1	81	47	0	6	217	255	211
183	202	237	145	0	0	12	108	200	138	243	236
195	206	123	207	177	121	123	200	175	13	96	218

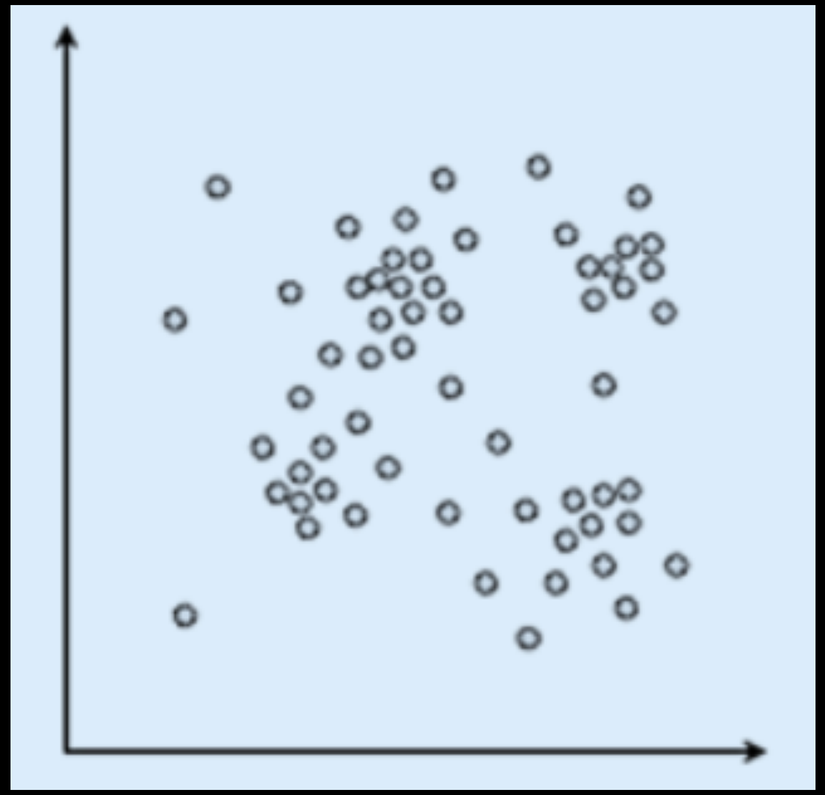

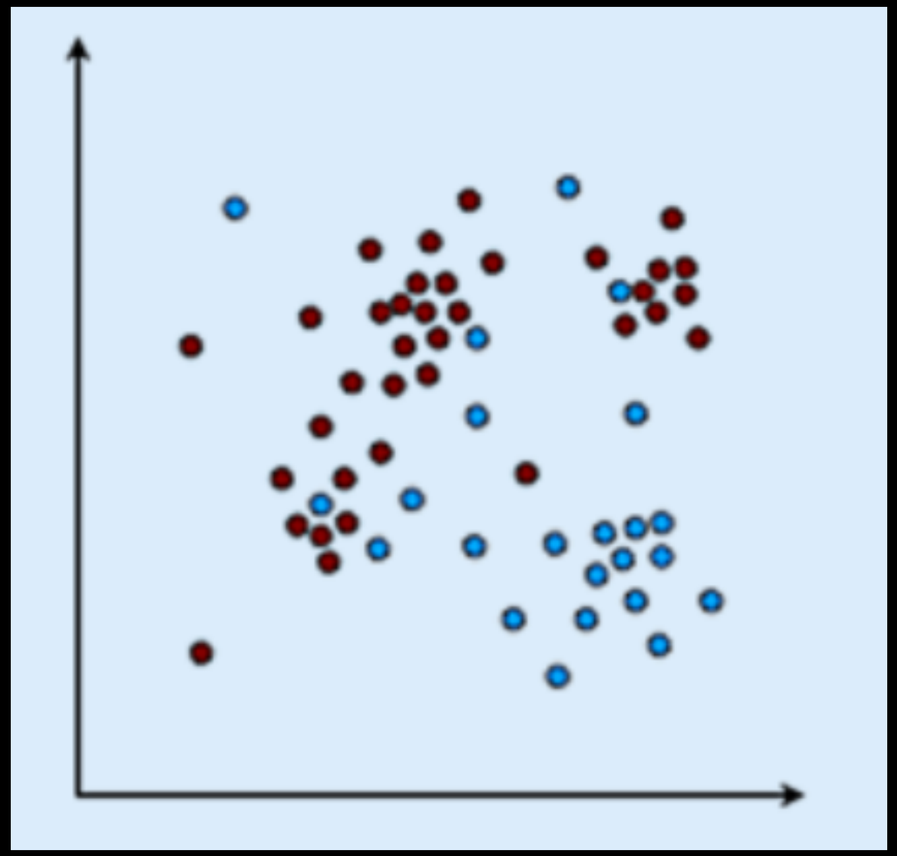

Supervised learning: each training example has a ground truth label. The model learns a decision boundary and replicates the labeling on new data.

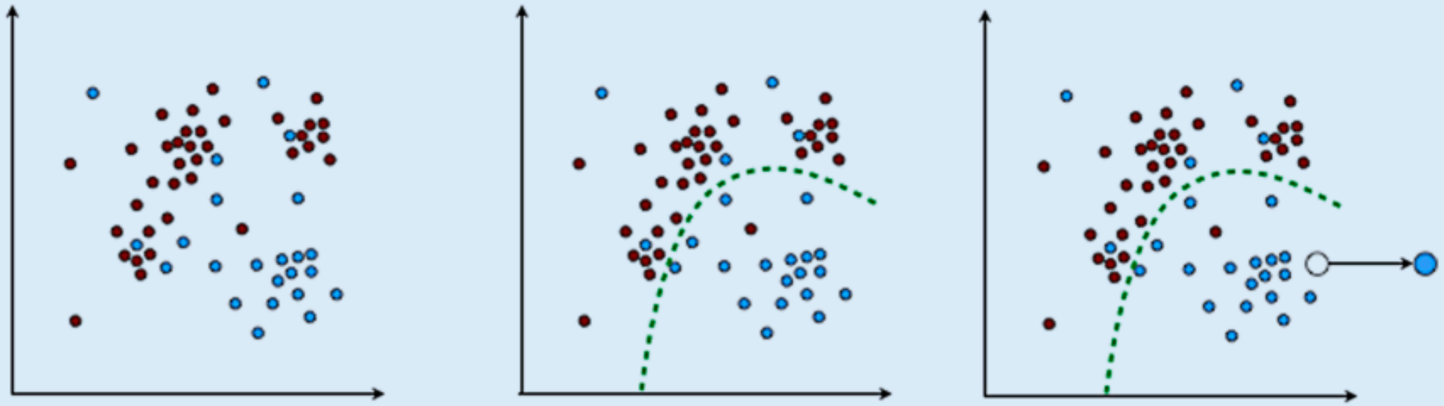

https://link.springer.com/article/10.1007/s00117-018-0407-3

Unsupervised learning: training examples do not have ground truth labels. The model identifies structure such as clusters. New data can be assigned to clusters.

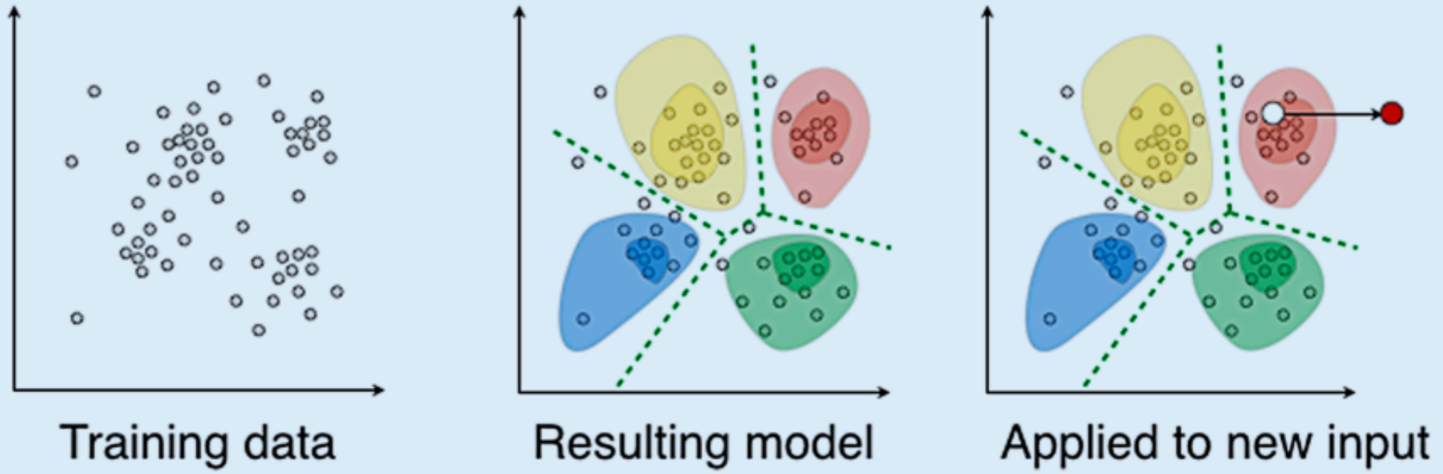

https://link.springer.com/article/10.1007/s00117-018-0407-3

SUPERVISED LEARNING

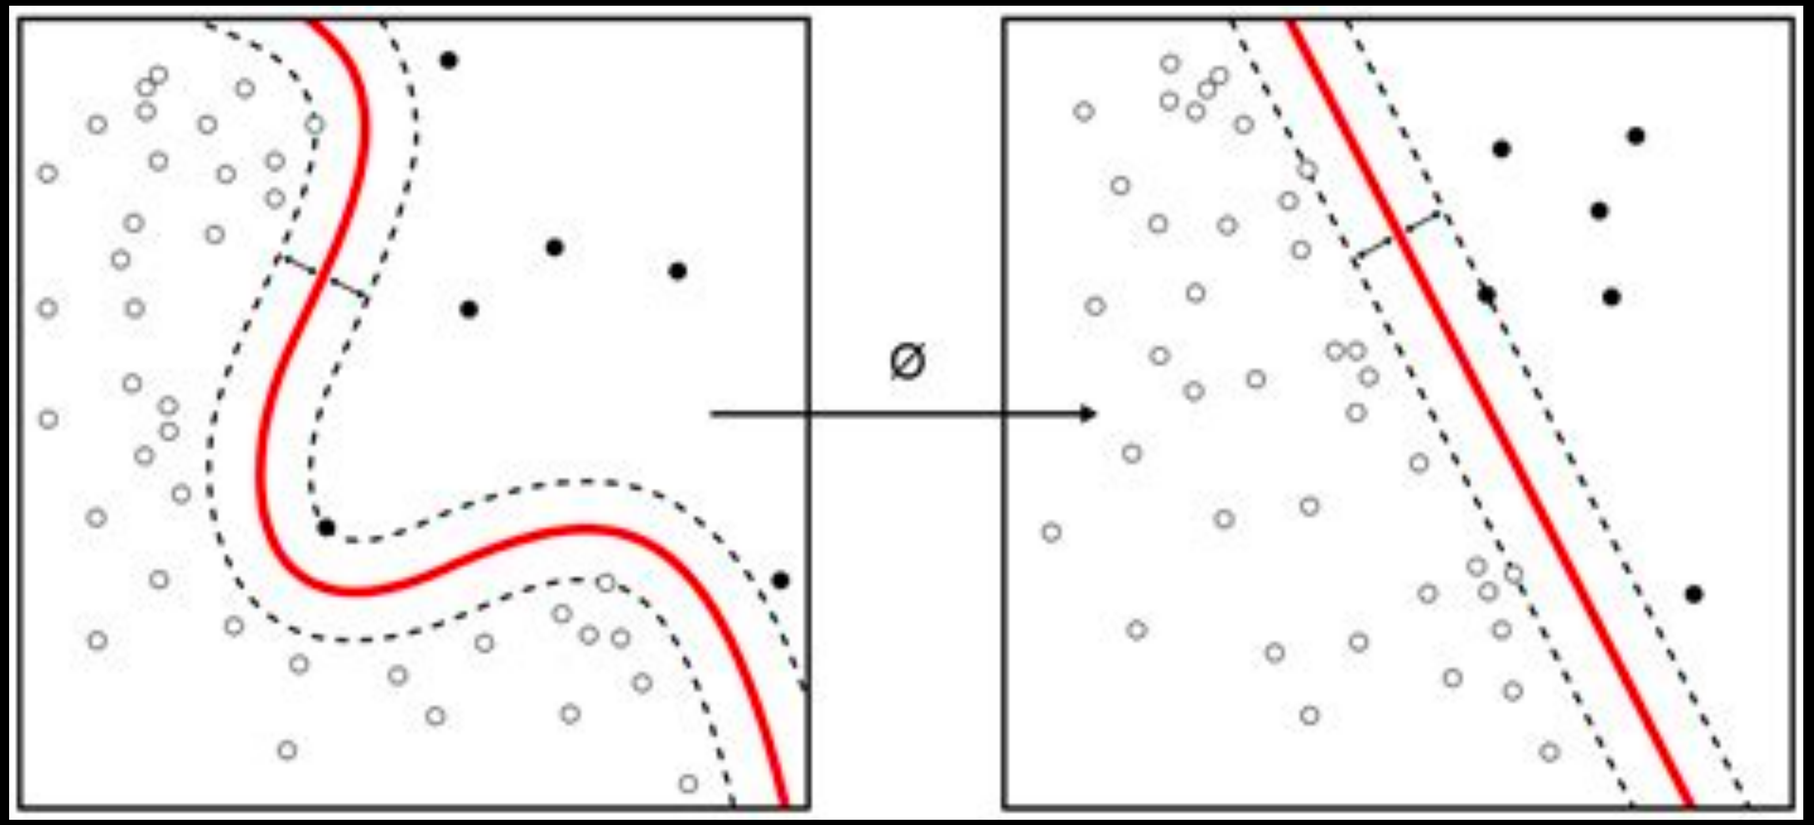

UNSUPERVISED LEARNING

VISION

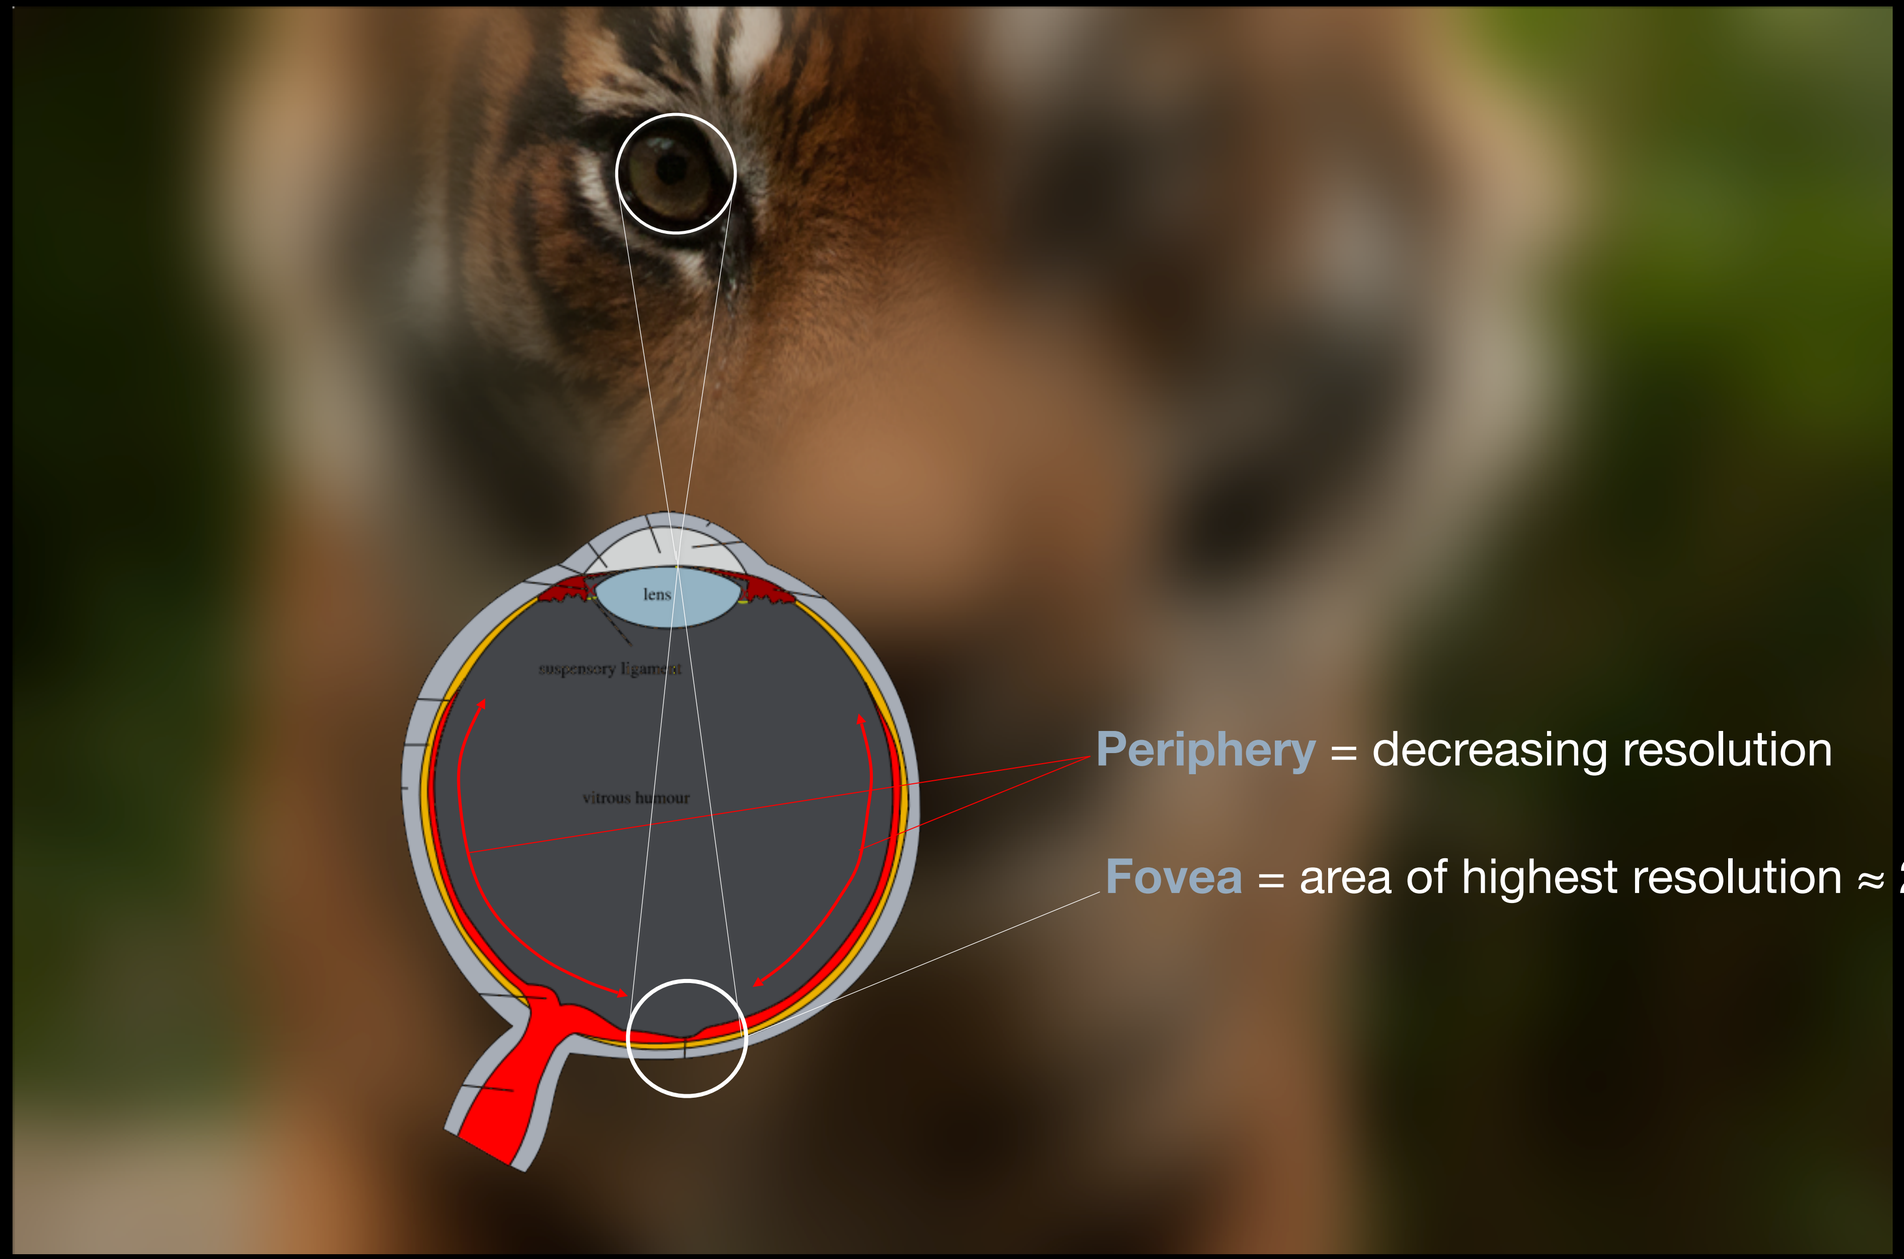

Periphery = decreasing resolution

Fovea = area of highest resolution ≈ 2°

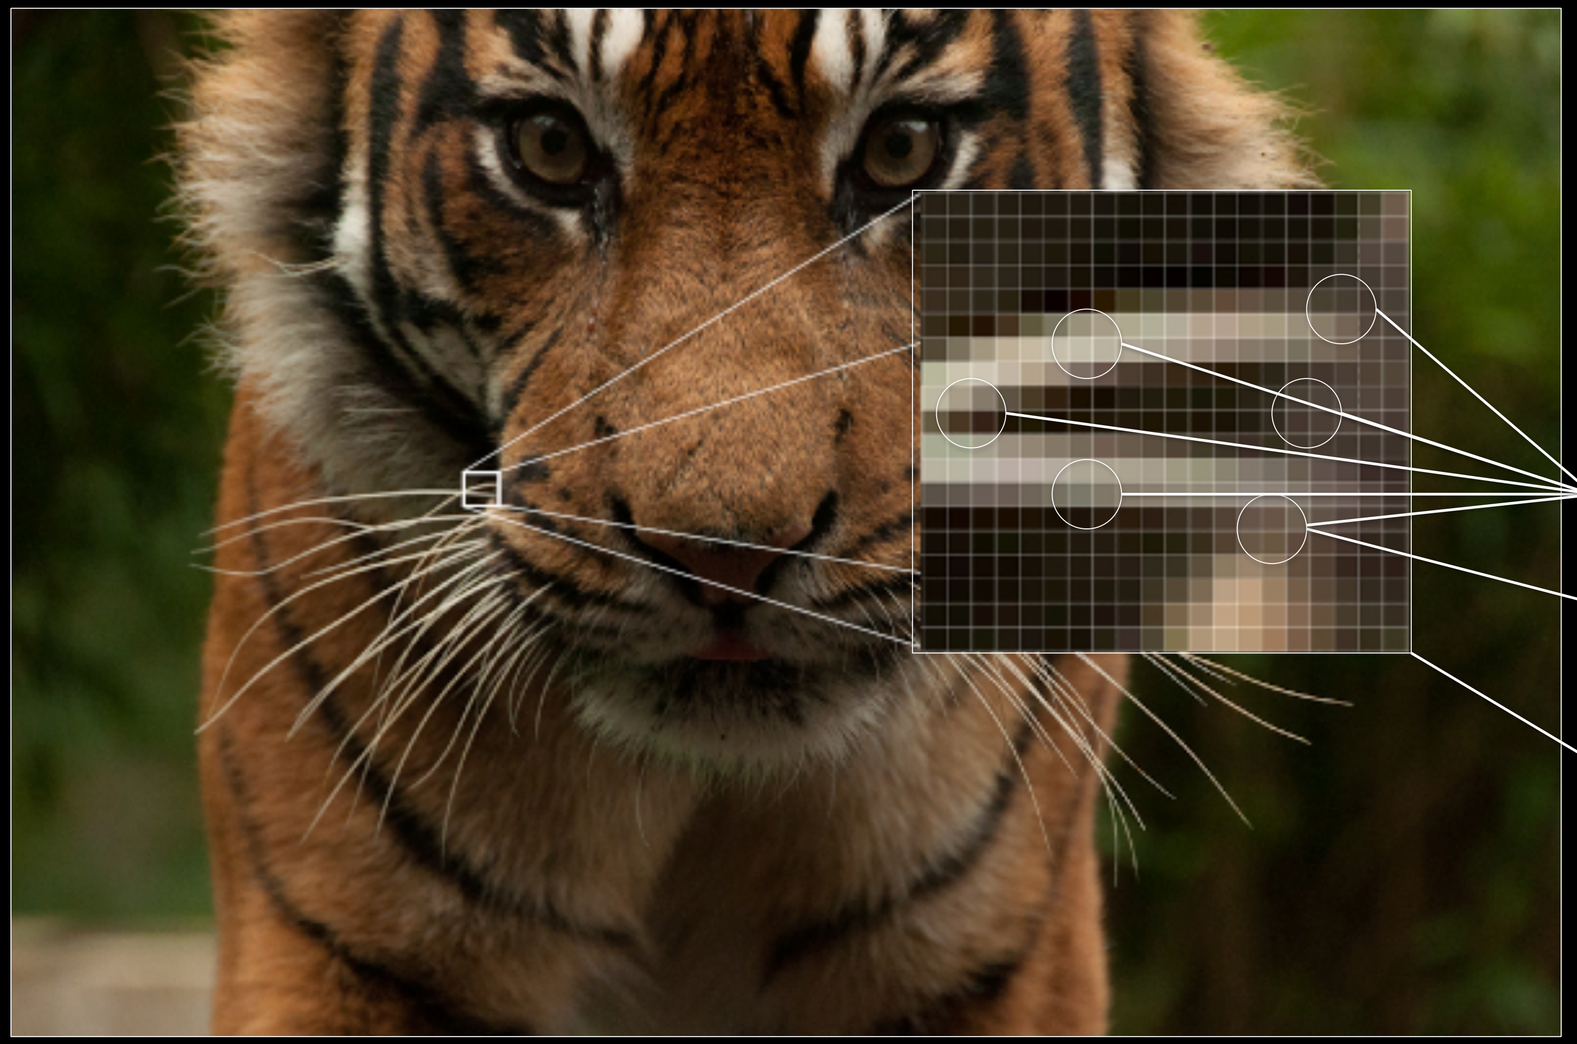

Semantic label =
High-level description

Grouping of Features =
Mid-level description

Single feature =
Low-level description

Pixels =
Low-level description

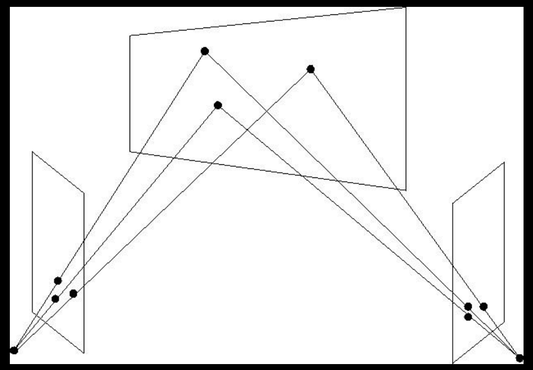

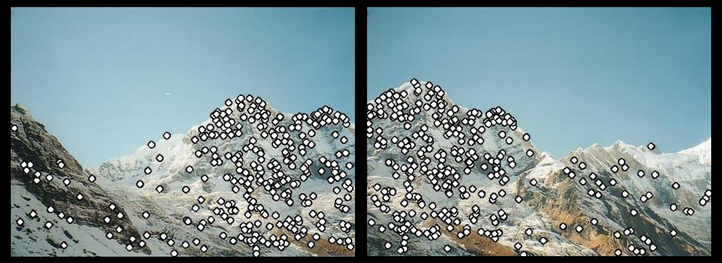

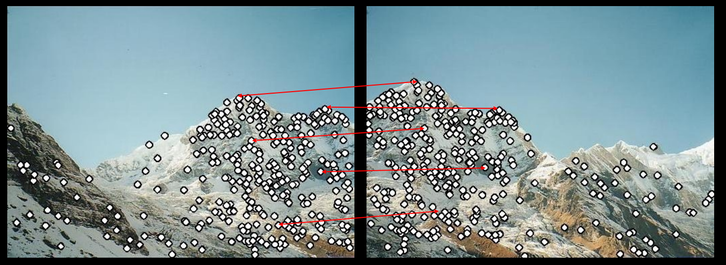

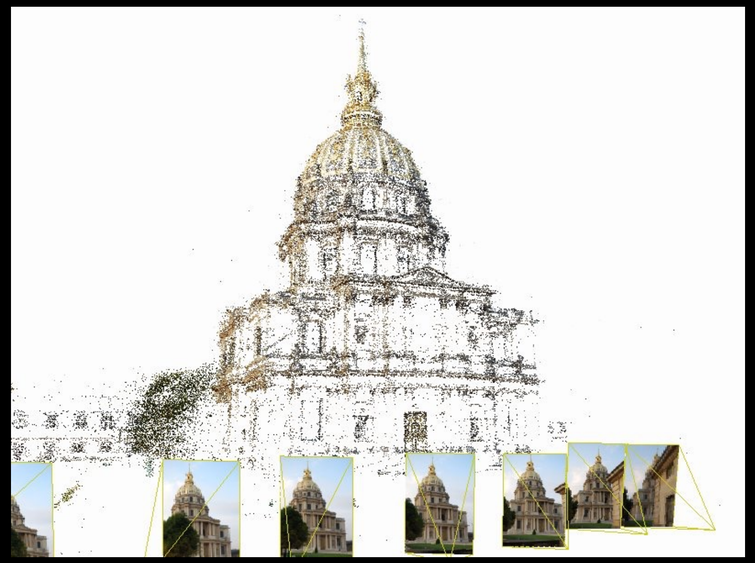

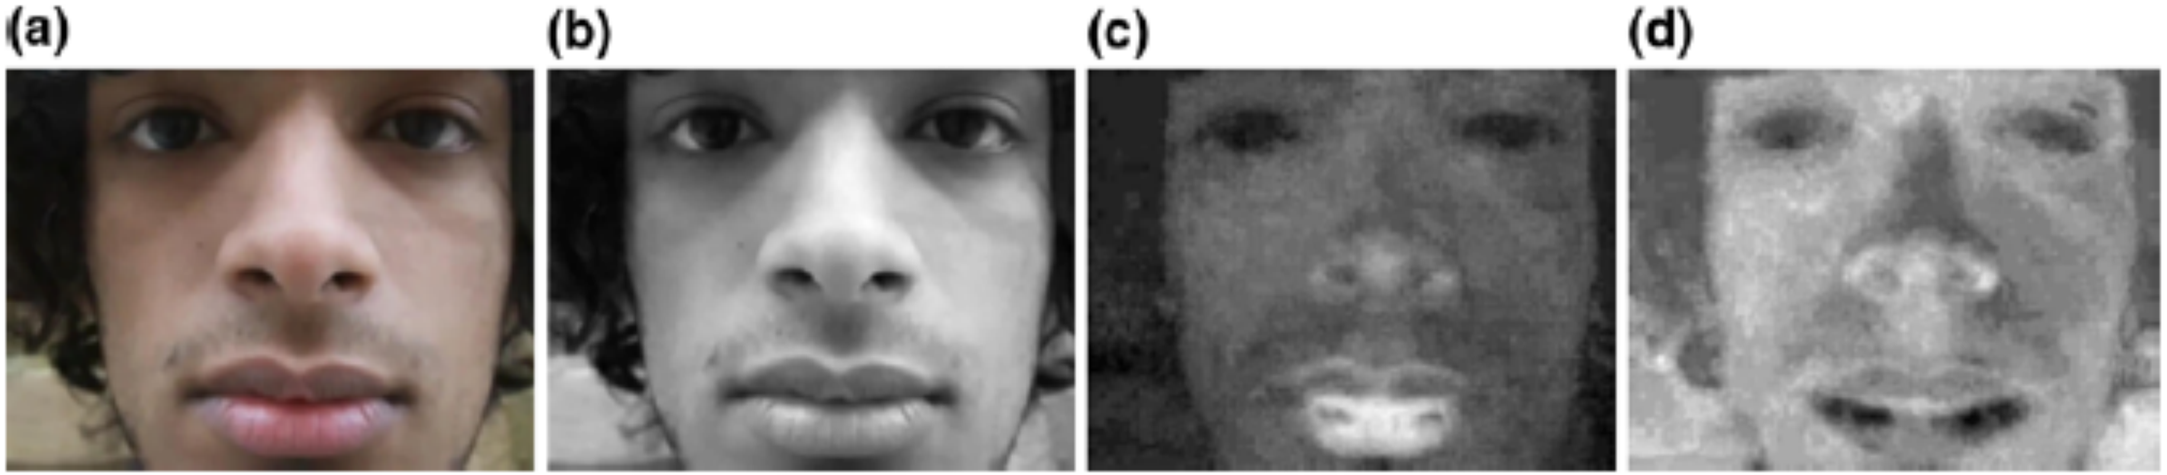

Fig. 1 **a** Original image of frame 1975 of video 24 ('Video Republic' http://www.demos.co.uk/publications/videorepublic); **b** L^* image depicting luminance (Lum); **c** a^* image depicting red/green opponent colors (RG); **d** b^* image depicting blue/yellow opponent colors (BY)

Parag K. Mital, Tim J. Smith, Robin Hill, John M. Henderson.
"Clustering of Gaze during Dynamic Scene Viewing is Predicted by Motion"
Cognitive Computation, Volume 3, Issue 1, pp 5-24, March 2011.

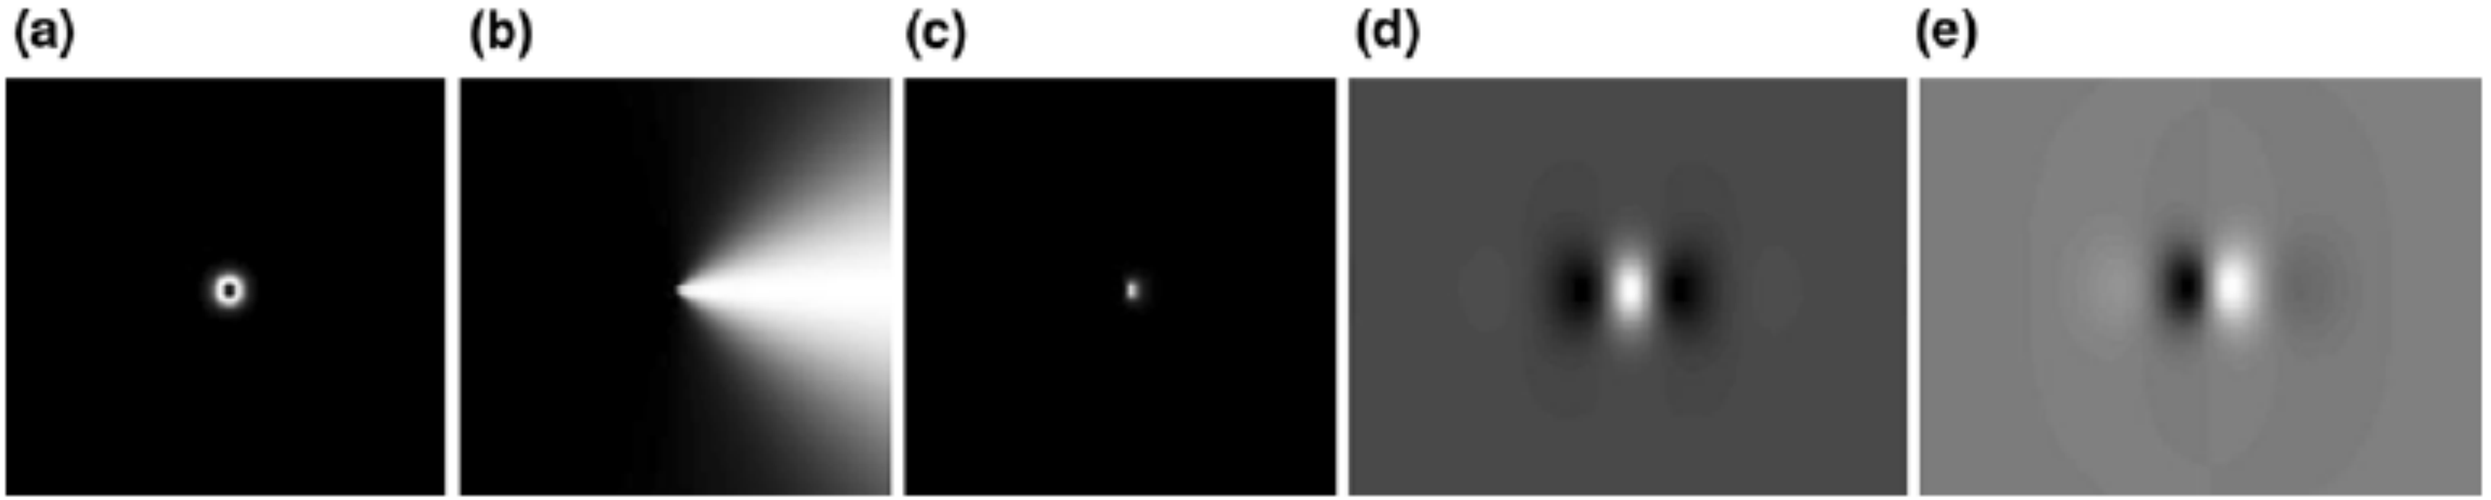

Fig. 3 The process for creating a log-Gabor kernel for 0° (*left* to *right*): **a** the radial map computed from multiplying a sinusoid with a Gaussian kernel; **b** the orientation of the kernel set for 0°; **c** the result of multiplying the radial (**a**) and orientation (**b**) maps; **d** the even symmetric component of the log-Gabor filter taken from the real part of the inverse fourier transform of the kernel; **e** the corresponding odd symmetric component taken from the imaginary component of the kernel

Parag K. Mital, Tim J. Smith, Robin Hill, John M. Henderson.
"Clustering of Gaze during Dynamic Scene Viewing is Predicted by Motion"
Cognitive Computation, Volume 3, Issue 1, pp 5-24, March 2011.

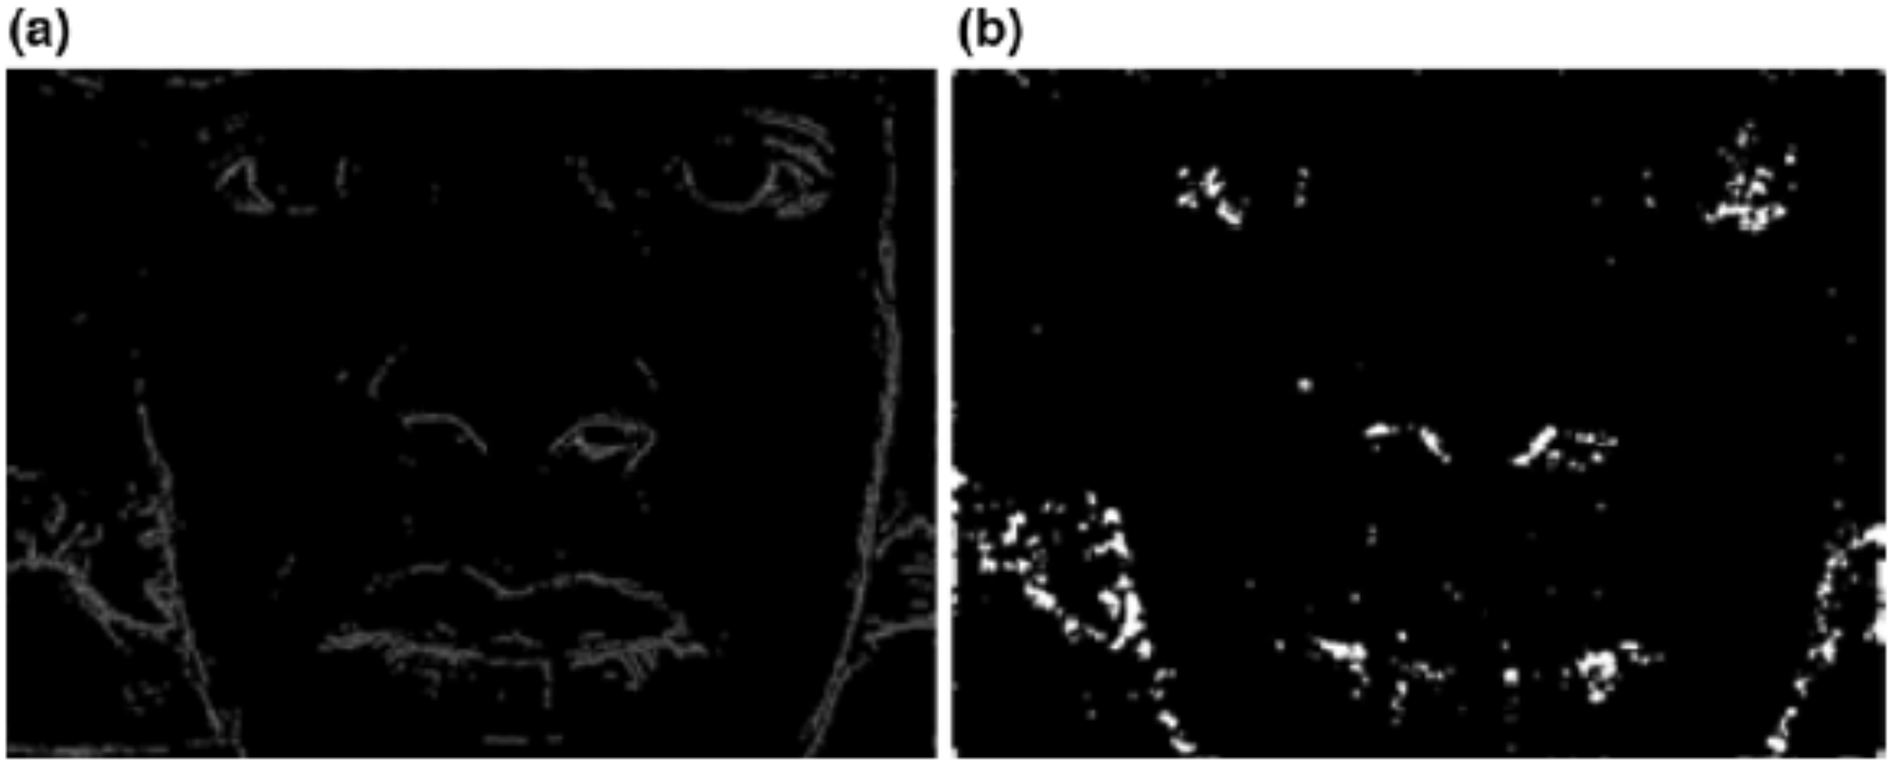

Fig. 2 a Edge map using Sobel edge detection; and b Harris corner detection map for the luminance image in Fig. 1b

Parag K. Mital, Tim J. Smith, Robin Hill, John M. Henderson.
"Clustering of Gaze during Dynamic Scene Viewing is Predicted by Motion"
Cognitive Computation, Volume 3, Issue 1, pp 5-24, March 2011.

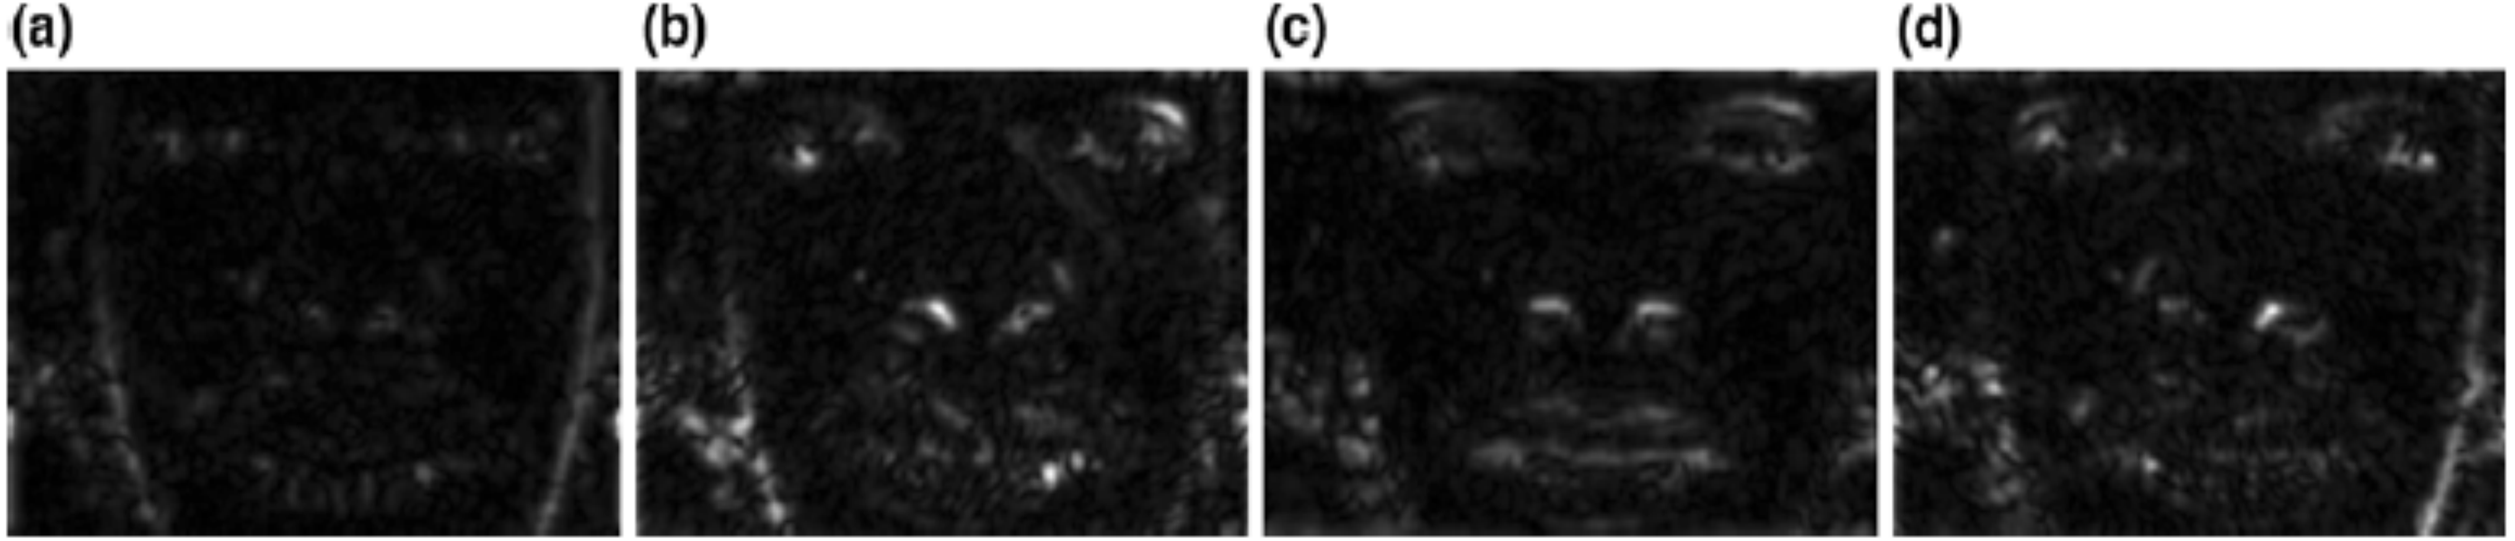

Fig. 4 Gabor-oriented maps for **a** 0°, **b** 45°, **c** 90°, and **d** 135° for the luminance image in Fig. 1b

Parag K. Mital, Tim J. Smith, Robin Hill, John M. Henderson.
"Clustering of Gaze during Dynamic Scene Viewing is Predicted by Motion"
Cognitive Computation, Volume 3, Issue 1, pp 5-24, March 2011.

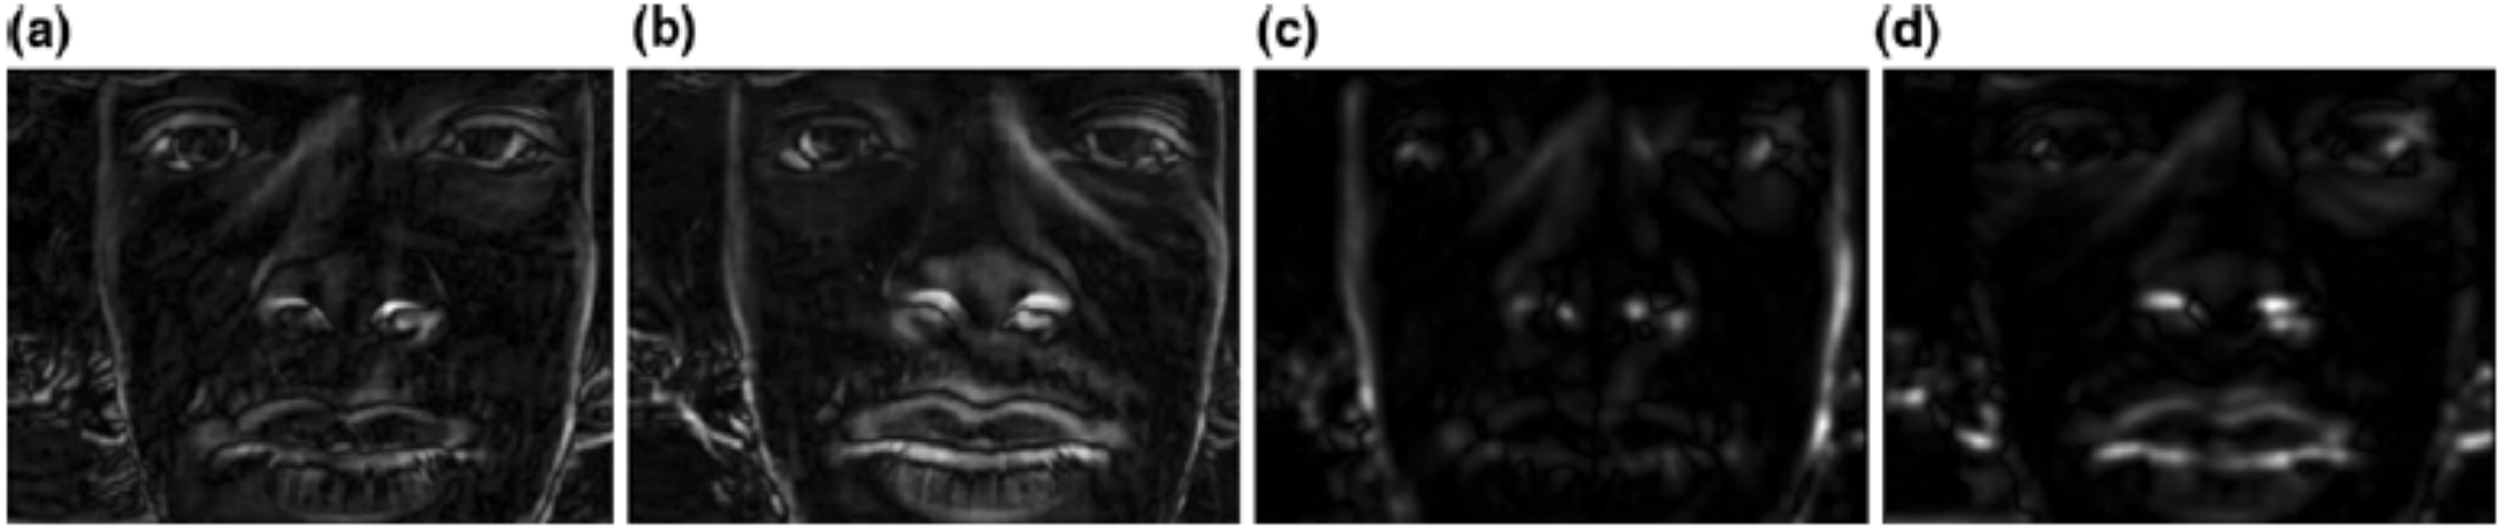

Fig. 5 **a** High-pass flicker (Flicker); **b** low-pass flicker (Flicker-N); **c** horizontal optical flow (*U*-Flow); **d** vertical optical flow (*V*-Flow) for the frame in Fig. 1a

Parag K. Mital, Tim J. Smith, Robin Hill, John M. Henderson.
"Clustering of Gaze during Dynamic Scene Viewing is Predicted by Motion"
Cognitive Computation, Volume 3, Issue 1, pp 5-24, March 2011.

AUDITION

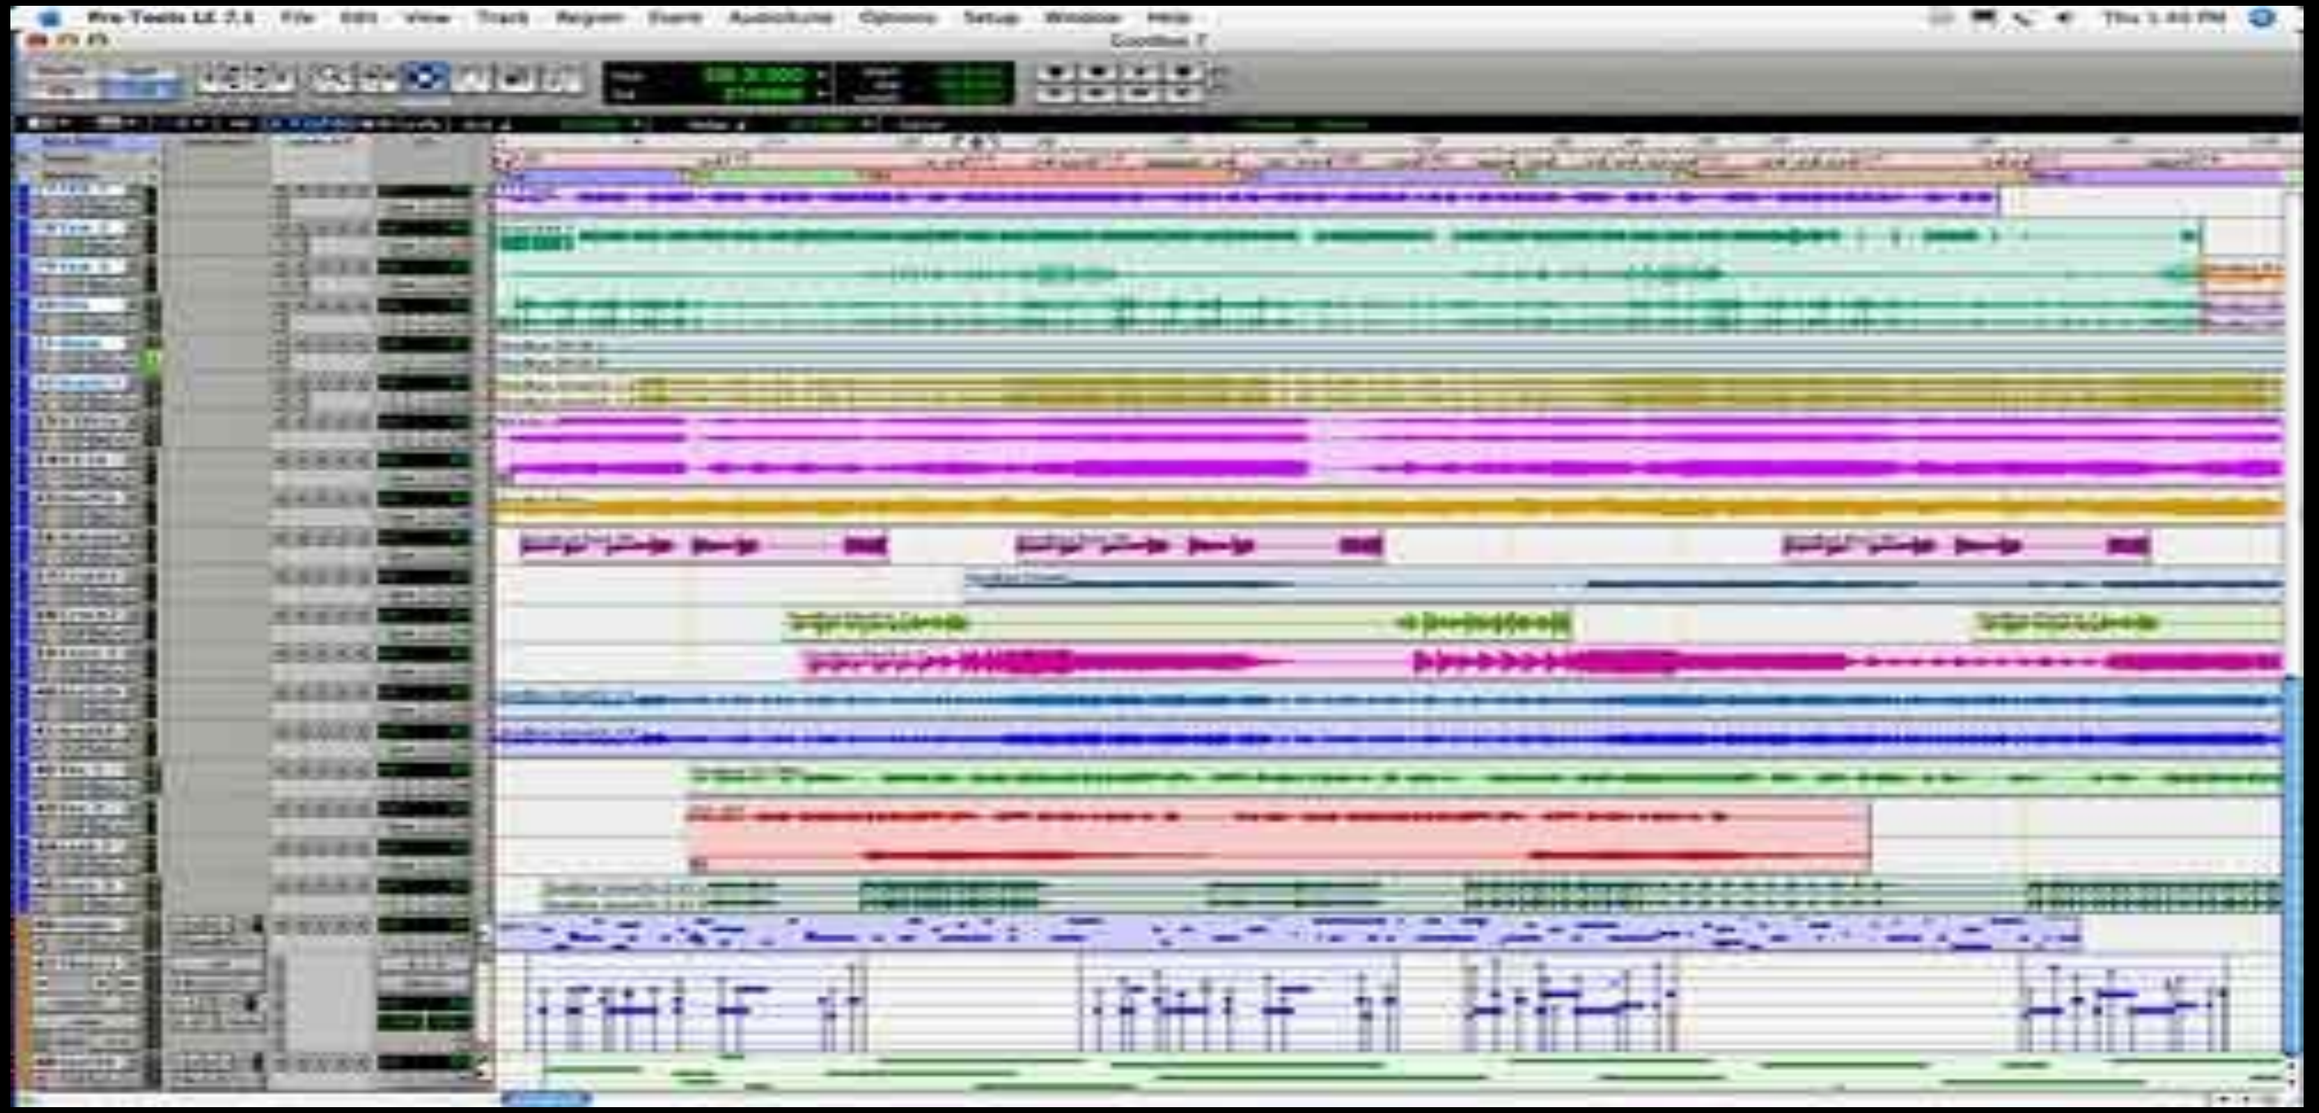

Parag K. Mital

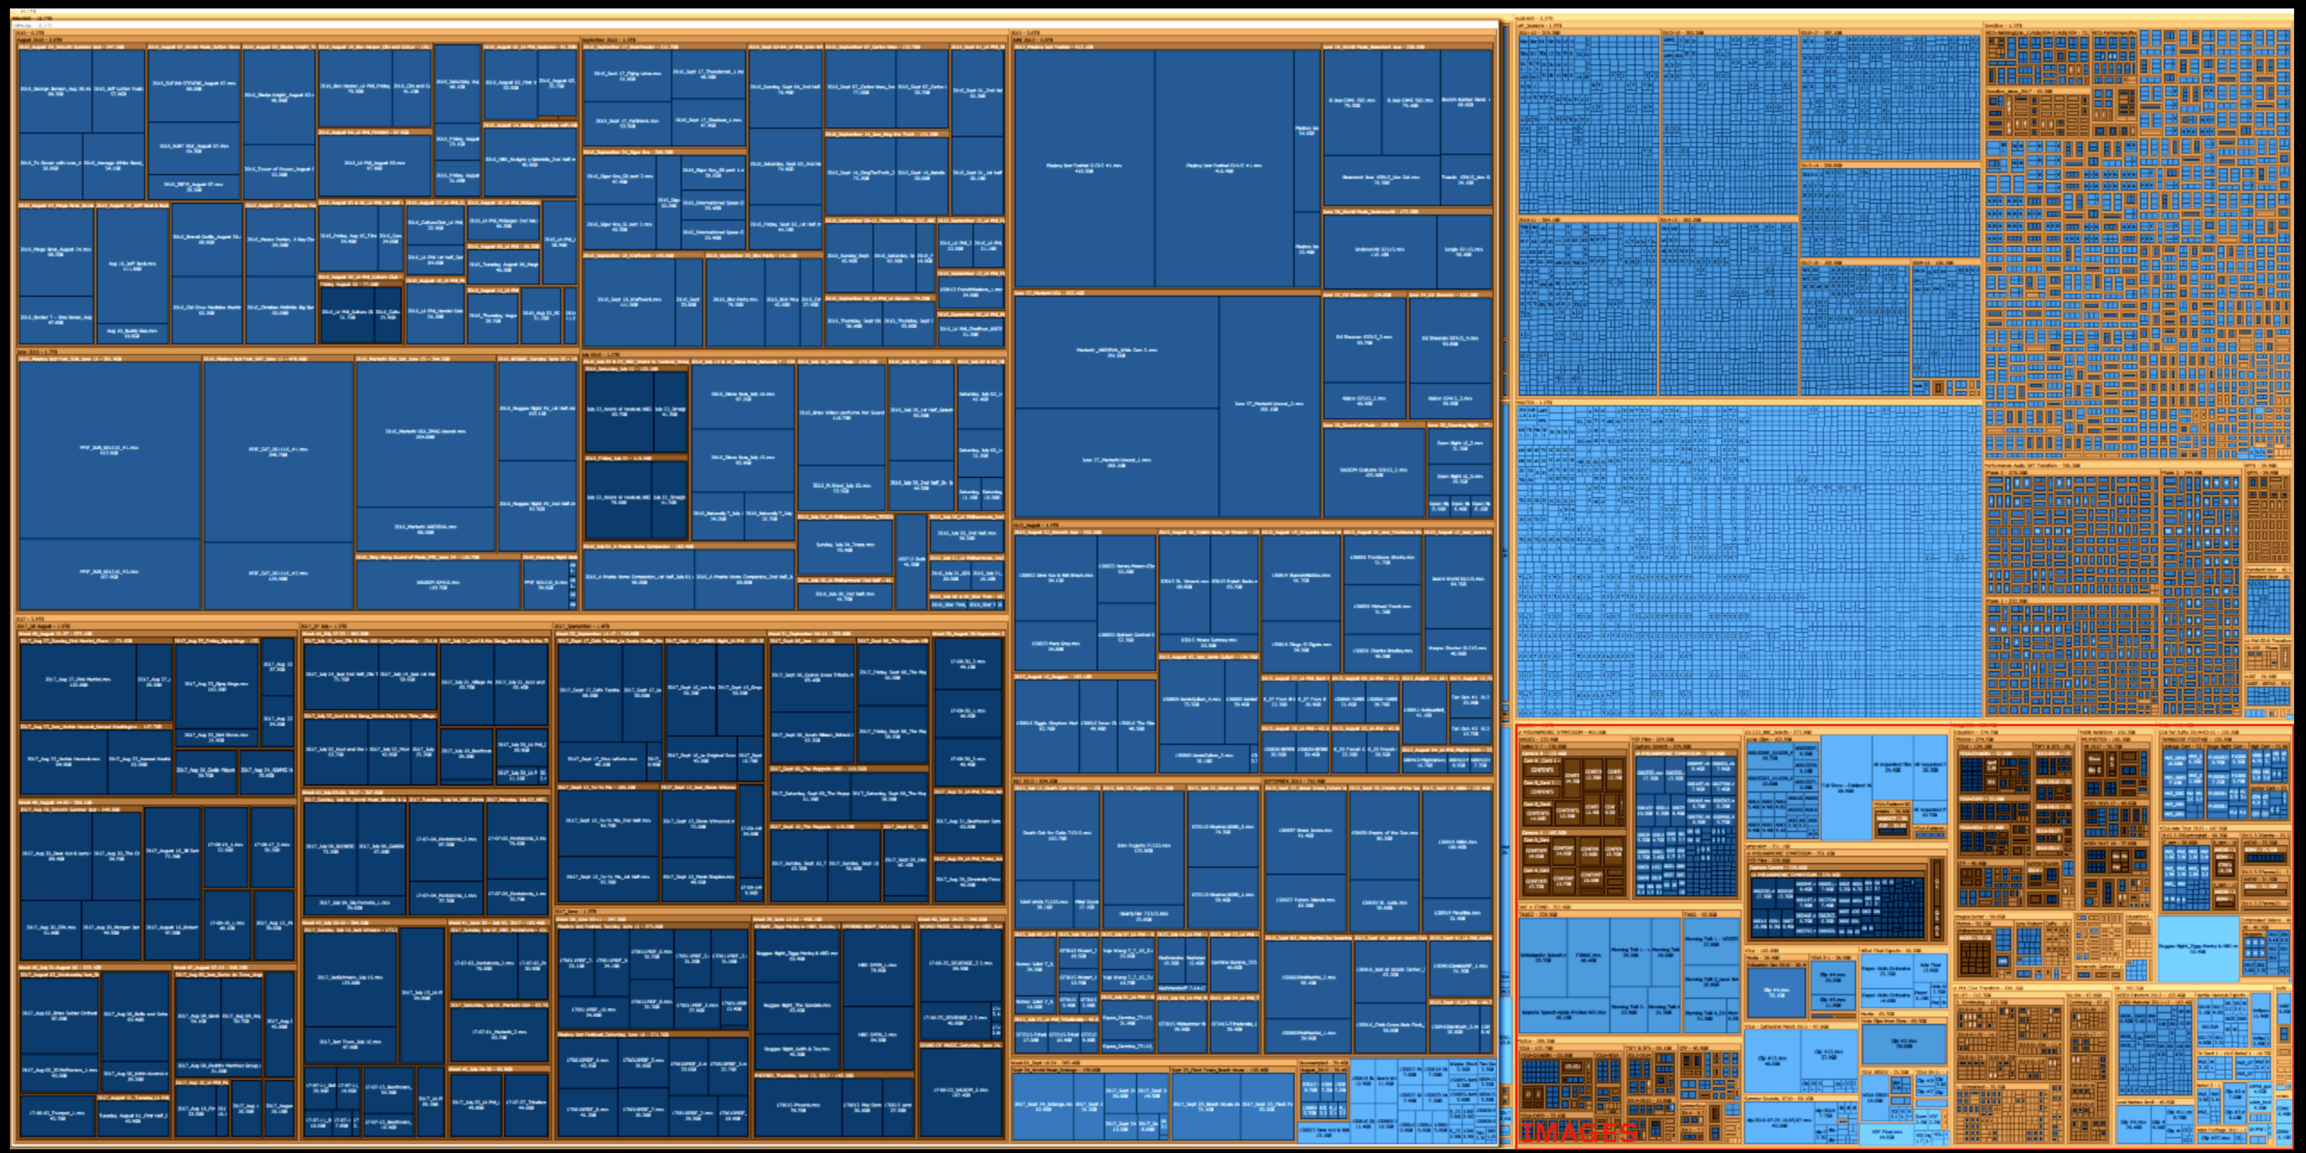

Parag K. Mital + Refik Anadol Studio

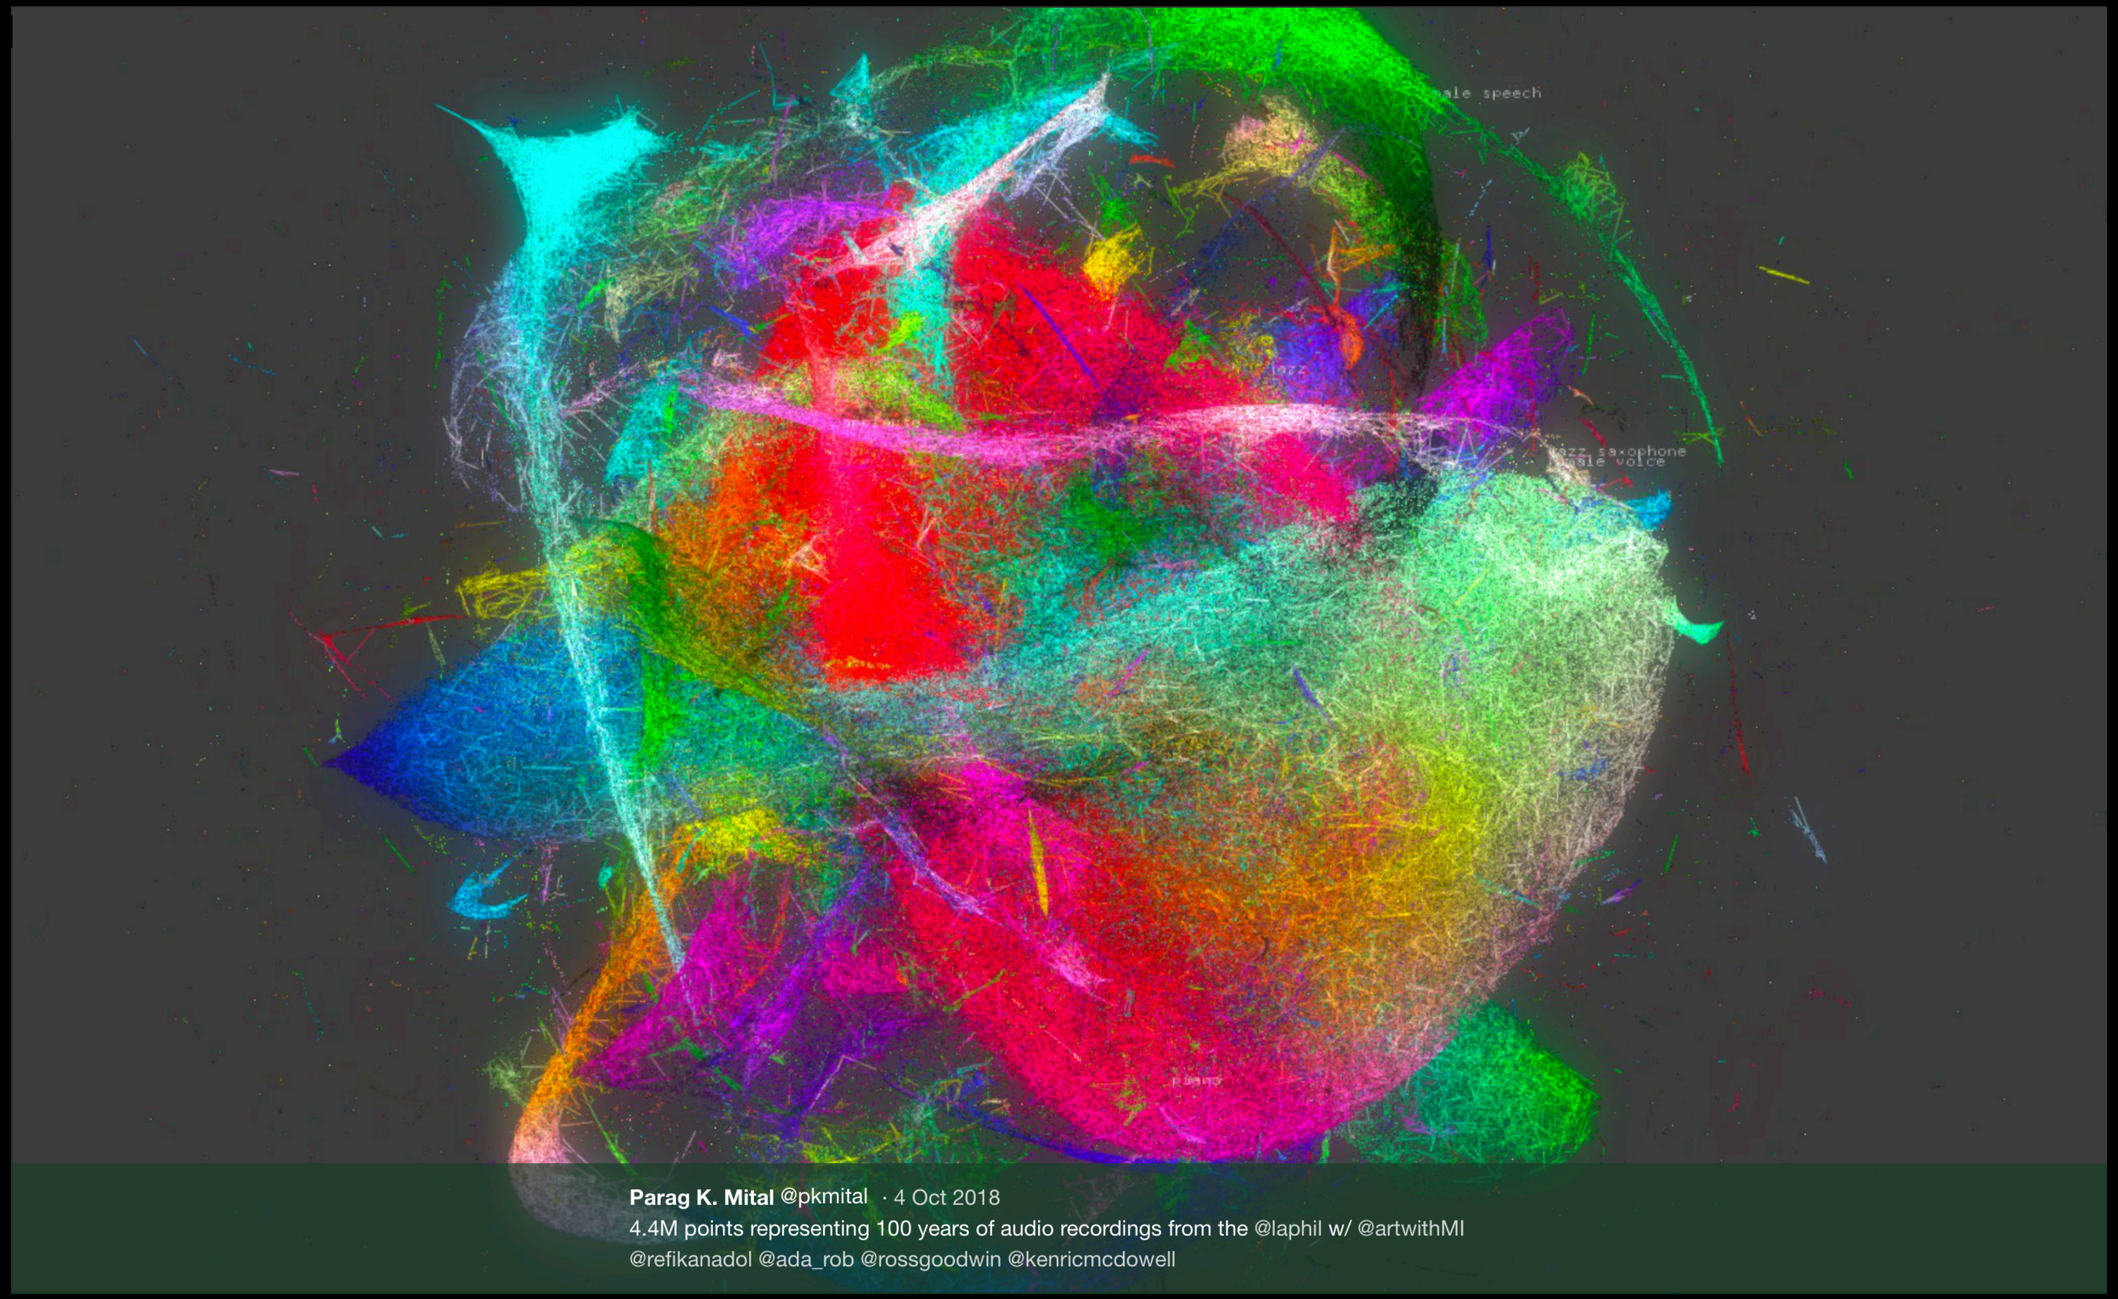

Parag K. Mital @pkmital · 4 Oct 2018
4.4M points representing 100 years of audio recordings from the @laphil w/ @artwithMI
@refikanadol @ada_rob @rossgoodwin @kenricmcdowell

AUTOMATION

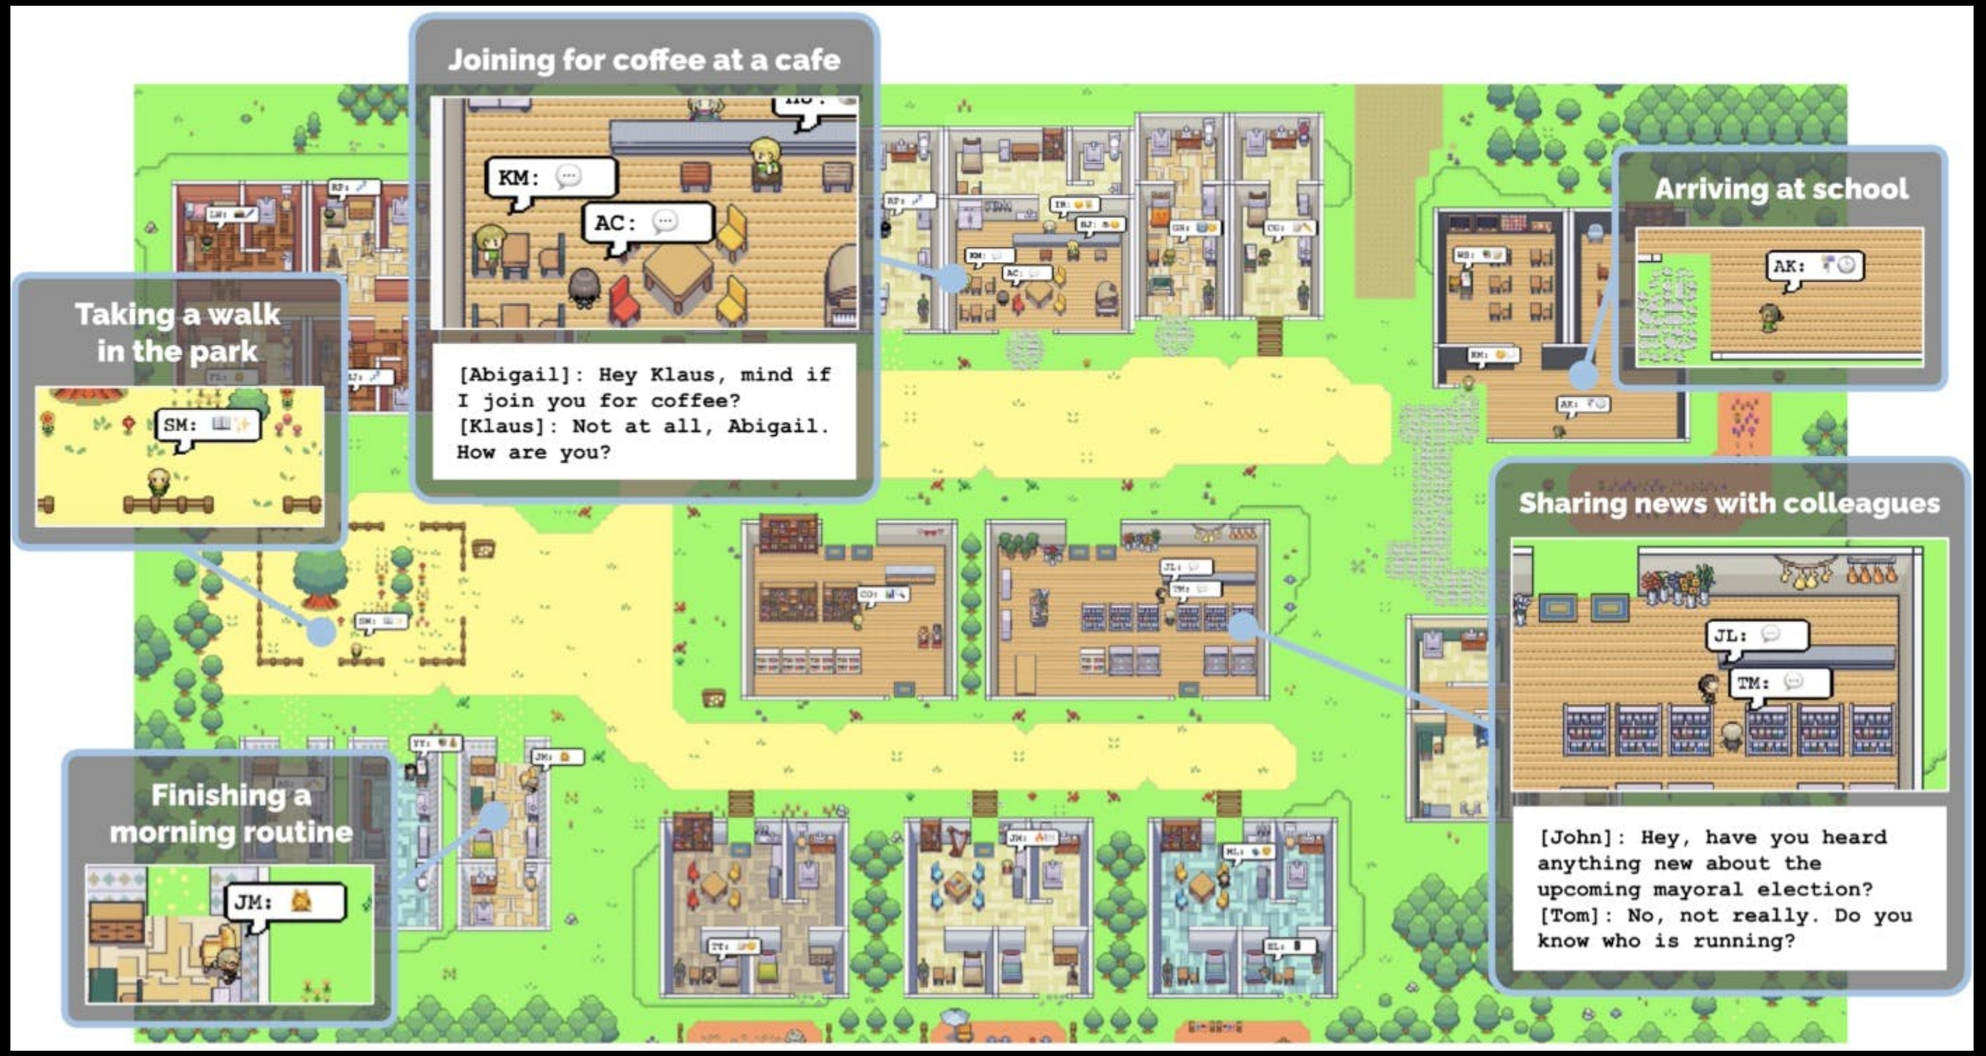

Illustration of "Generative Agents"
https://reverie.herokuapp.com/arXiv_Demo/

To Infinity and Beyond: SHOW-1 and Showrunner Agents in Multi-Agent Simulations

Philipp Maas Fable Studio
Frank Carey Fable Studio
Chris Wheeler Fable Studio
Edward Saatchi Fable Studio
Pete Billington Fable Studio
Jessica Yaffa Shamash Fable Studio

Abstract

In this work we present our approach to generating high-quality episodic content for IP's (Intellectual Property) using large language models (LLMs), custom state-of-

https://fablestudio.github.io/showrunner-agents/

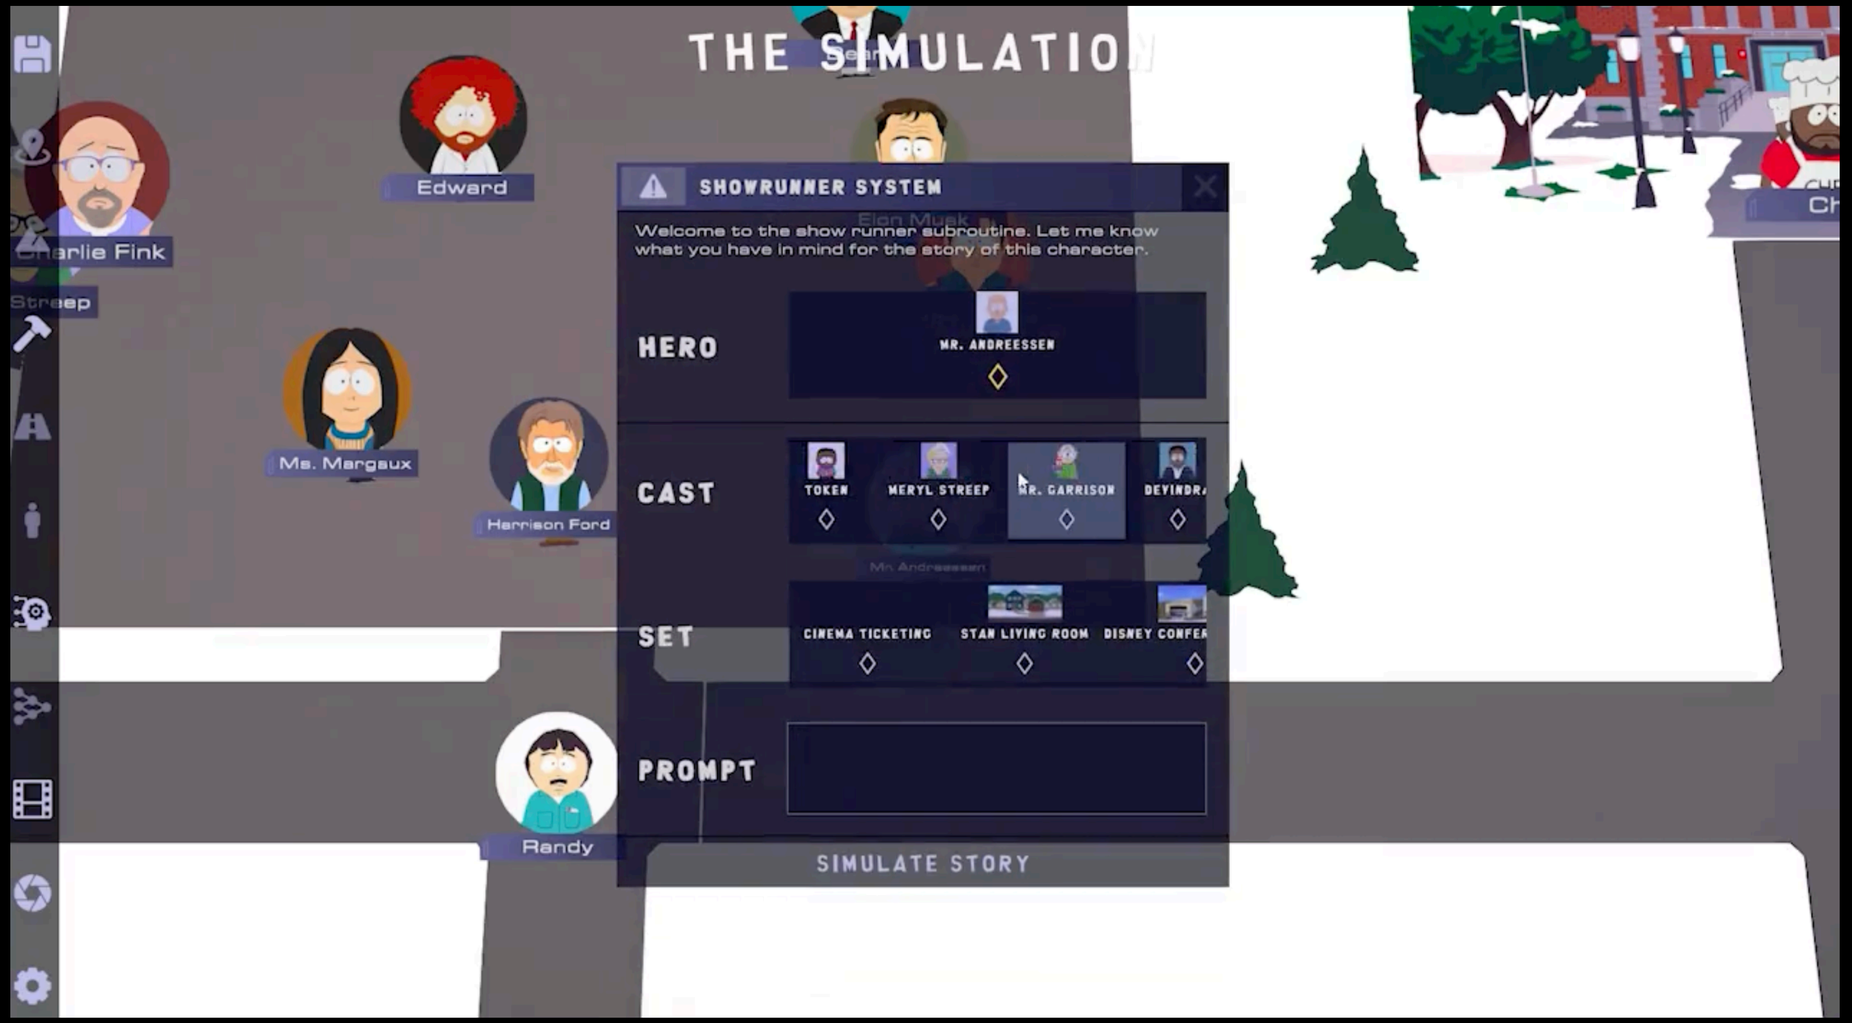

"The Simulation" - SouthPark - https://twitter.com/fablesimulation/status/1681352904152850437?lang=en

Communicative Agents for Software Development

Chen Qian[♠] Xin Cong[♠] Wei Liu[♠] Cheng Yang[♣] Weize Chen[♠] Yusheng Su[♠]
Yufan Dang[♠] Jiahao Li[♦] Juyuan Xu[▲] Dahai Li[★] Zhiyuan Liu[♠✉] Maosong Sun[♠✉]
[♠]Tsinghua University [♣]Beijing University of Posts and Telecommunications
[♦]Dalian University of Technology [▲]Brown University [★]Modelbest Inc.
qianc62@gmail.com liuzy@tsinghua.edu.cn sms@tsinghua.edu.cn

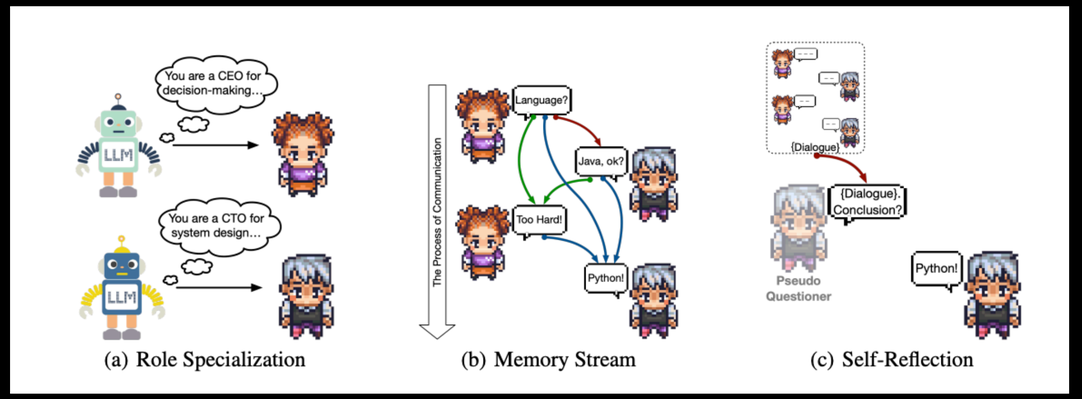

(a) Role Specialization (b) Memory Stream (c) Self-Reflection

(a) Naive Instruction in Coding (b) Thought Instruction in Coding

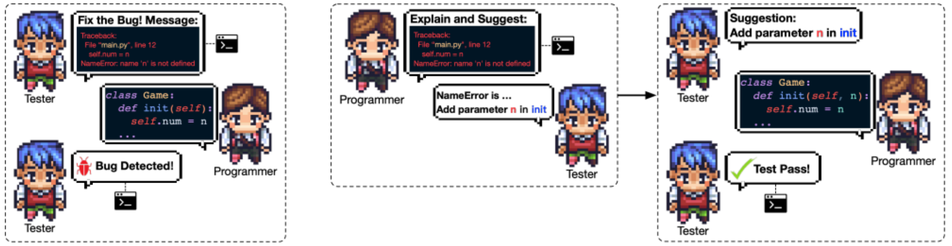

(c) Naive Instruction in Testing (d) Thought Instruction in Testing

https://arxiv.org/pdf/2307.07924v3.pdf

COURSE OUTLINE

Updated course schedule:

https://pkmital.com/home/teaching/ucla-cultural-automation-with-machine-learning/

COURSE GRADING

- 8 assignments, @ 7% each = 56%
- 8 feedback/ lab sessions, @ 3% each = 24%
- 1 final project = 30%

  Total = 110%

- > 90 = A
- 90 - 80 = B
- 80 - 70 = C
- 70 - 60 = D
- < 60 = 😟

HOMEWORK

Homework

No homework today, see you Wednesday :)

Upcoming

WENDESDAY

- Lecture: Risks of Machine Learning
- Homework assigned: Research presentations

MONDAY (LAB)

- Student research presentations
- Group review session with student feedback