Computational Media Research

CMPM 202, W2019

Prof. Angus Forbes (instructor)
https://creativecoding.soe.ucsc.edu
angus@ucsc.edu

Course Website & Slack Channel

Website:

https://creativecoding.soe.ucsc.edu/courses/cmpm202

Slack:

https://ucsc-cmpm202-w2019.slack.com

About Me

- Assistant professor in UCSC's Computational Media Department, direct the Creative Coding lab
- I teach courses on data visualization, computer graphics, and, starting today, computational media research
- Research areas: Visualization & Visual Analytics, Computer Graphics, VR / AR, New Media Arts
- Recent focus on novel ways to represent and analyze dynamic networks
- Collaborate with scientists to design visualization tools to make it easier to reason about complex data

Research Communities

- ACM SIGGRAPH 2018 Arts Papers Chair
- IEEE VIS Program committee member for 2018-2021
- IEEE VIS Arts Program, General Chair, 2013-2017
- ACM/EG Expressive Program Chair 2019; Papers Chair 2016; Arts Chair 2015

- Have contributed to ISEA, VR, NIME, ICMC, CHI, BigData, CSCW, UIST, and others

Computational Media

- What is the discipline of Computational Media? How does it differ from other disciplines? How does it relate to Computer Science? Media Arts? Digital Humanities? Game Studies?

- What are some examples of Computational Media research outputs? Where can you see them? hear them? read them?

- How are successful examples of Computational Media research evaluated and judged?

- What are the most exciting research areas to you? Why is it exciting? Why are you drawn to these areas? What are you planning on working on over the next few years?

Computational Media Research

Topics in Computational Media:

- Generative Art & Design
- Digital Humanities / Humanities Computing
- AI & Ethics
- Visualization and Computational Photography
- Sound and Sonification

Applied Deep Learning:

- Convolutional Neural Networks
- Generative Adversarial Networks
- Reinforcement Learning

Course Structure

Week 1

Introduction

Weeks 2 – 7

Tuesdays:

Lecture and Discussion

Thursdays:

Lab

Weeks 8 – 11

Work on final project:

Labs, Discussions, Meetings, Presentations, Writing, …

Expectations & Deliverables

1. Writing portfolio

- Each week I will ask you to explore a series of topics, answer questions, take positions on and respond to articles and artifacts, and/or generate ideas

- By Monday at 11:59pm each week, you will submit your writing via Google Docs

- Before class on Tuesday, you will review your classmates responses to prepare for in class conversation

- You will include all of your writing responses as a packet in your final portfolio submission

Expectations & Deliverables

2. Smaller Projects (~1 week)

- You will develop a two smaller projects
- Goal is to practice understanding and implementing generative techniques (WFC, neural style transfer, etc.)
- You will create a GitHub repository with working code, clear instructions, and a short write-up with images and video documentation
- Can work alone or in small teams (3 people max)

Expectations & Deliverables

3. Final Project (~4 weeks)

- You will work in a team to develop a more innovative larger group project utilizing deep learning
- You will practice defining and sharing responsibilities
- You will create a GitHub repository with working code, clear instructions, clear delineation of responsibilities, and a public facing Readme with images and examples
- You will create a CS conference-style paper using a LaTeX ACM conference template

Goals of Course

- Practice writing Computer Science conference-style papers
- Introduction to deep learning concepts and applications
- Better understanding of different academic computational media-esque communities
- Continue to become more creative and rigorous thinkers, makers, scholars
- Practice working successfully in teams
- Prepare for First Year Exam (for PhD students)
- Increase proficiency in tools and frameworks (git, LaTeX, TensorFlow, etc)

CM Practical Learning Outcomes

- Synthesize fundamental CM knowledge to **describe**: the field's development, trade-offs between different project design and evaluation approaches, and the relations of these topics with each other
    - Contextualize topic areas effectively within the larger ***history*** of the field, draw in relevant sources beyond those directly covered in coursework, and demonstrate synthetic understanding of the material discussed

- Use knowledge of least one CM topic to **identify** a potentially *novel contribution* and relate it to prior work in the field
    - Delineate highly novel contributions that are both creative and original
    - Draw connections to contemporary areas of related work

- Use sufficient and appropriate technical knowledge to **develop** a *novel contribution*
    - Introduce new technical concepts and demonstrate successful technical innovation
    - Successfully extend an existing technical approach, or successfully use an approach new to the area

Video Traces

Dennis Hlynsky video processing

- Traces particular elements in a scene to emphasize patterns that would otherwise remain unseen, such as the motion of moths or gnats, or the flight of birds

- Temporal conflations that function as a form of in situ information visualization

- Vimeo site: https://vimeo.com/87207954

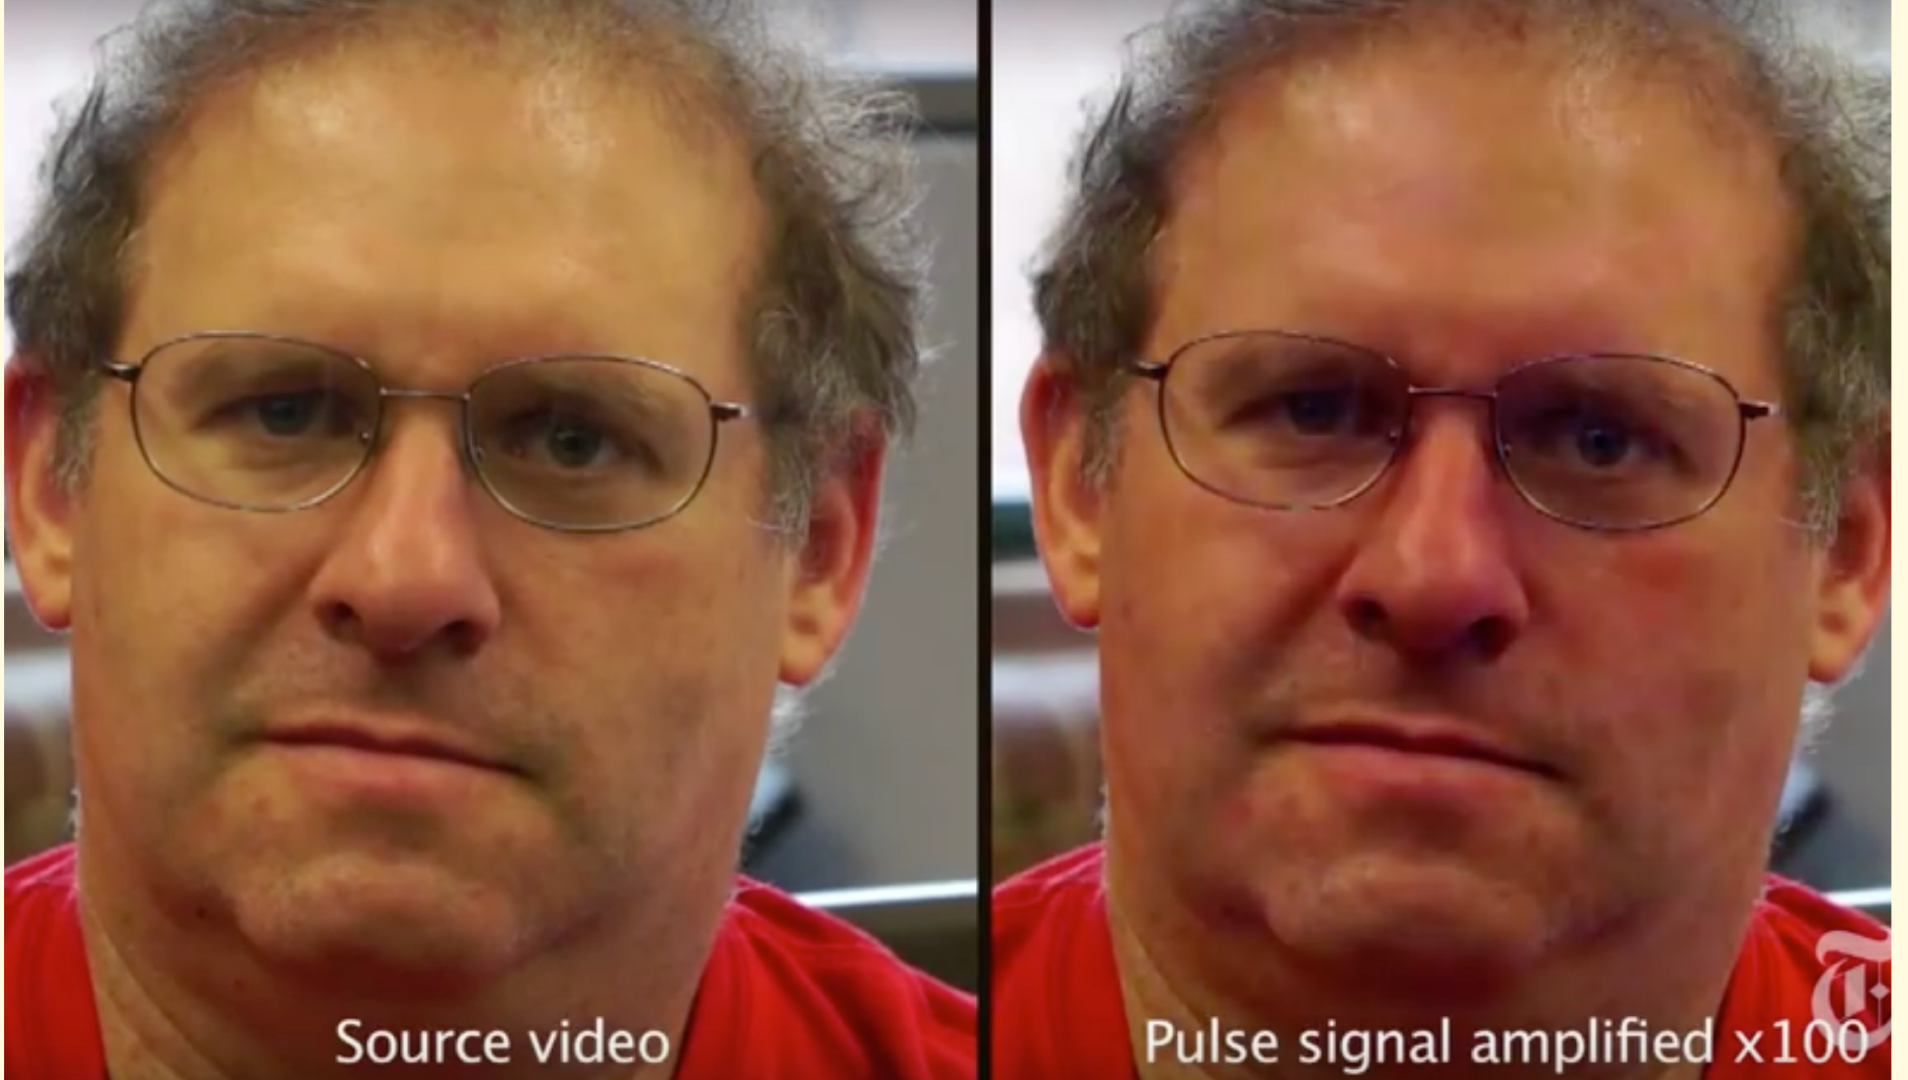
Source video
Pulse signal amplified x100

Eulearian Video Magnification

- Identify subtle patterns present in a video, and then amplify those patterns so that they are more easily perceivable by the viewer

- Exaggerates color or motion to show, for example, the breathing or blood pulse of an infant

- From William Freeman's lab at MIT: http://people.csail.mit.edu/mrub/vidmag/

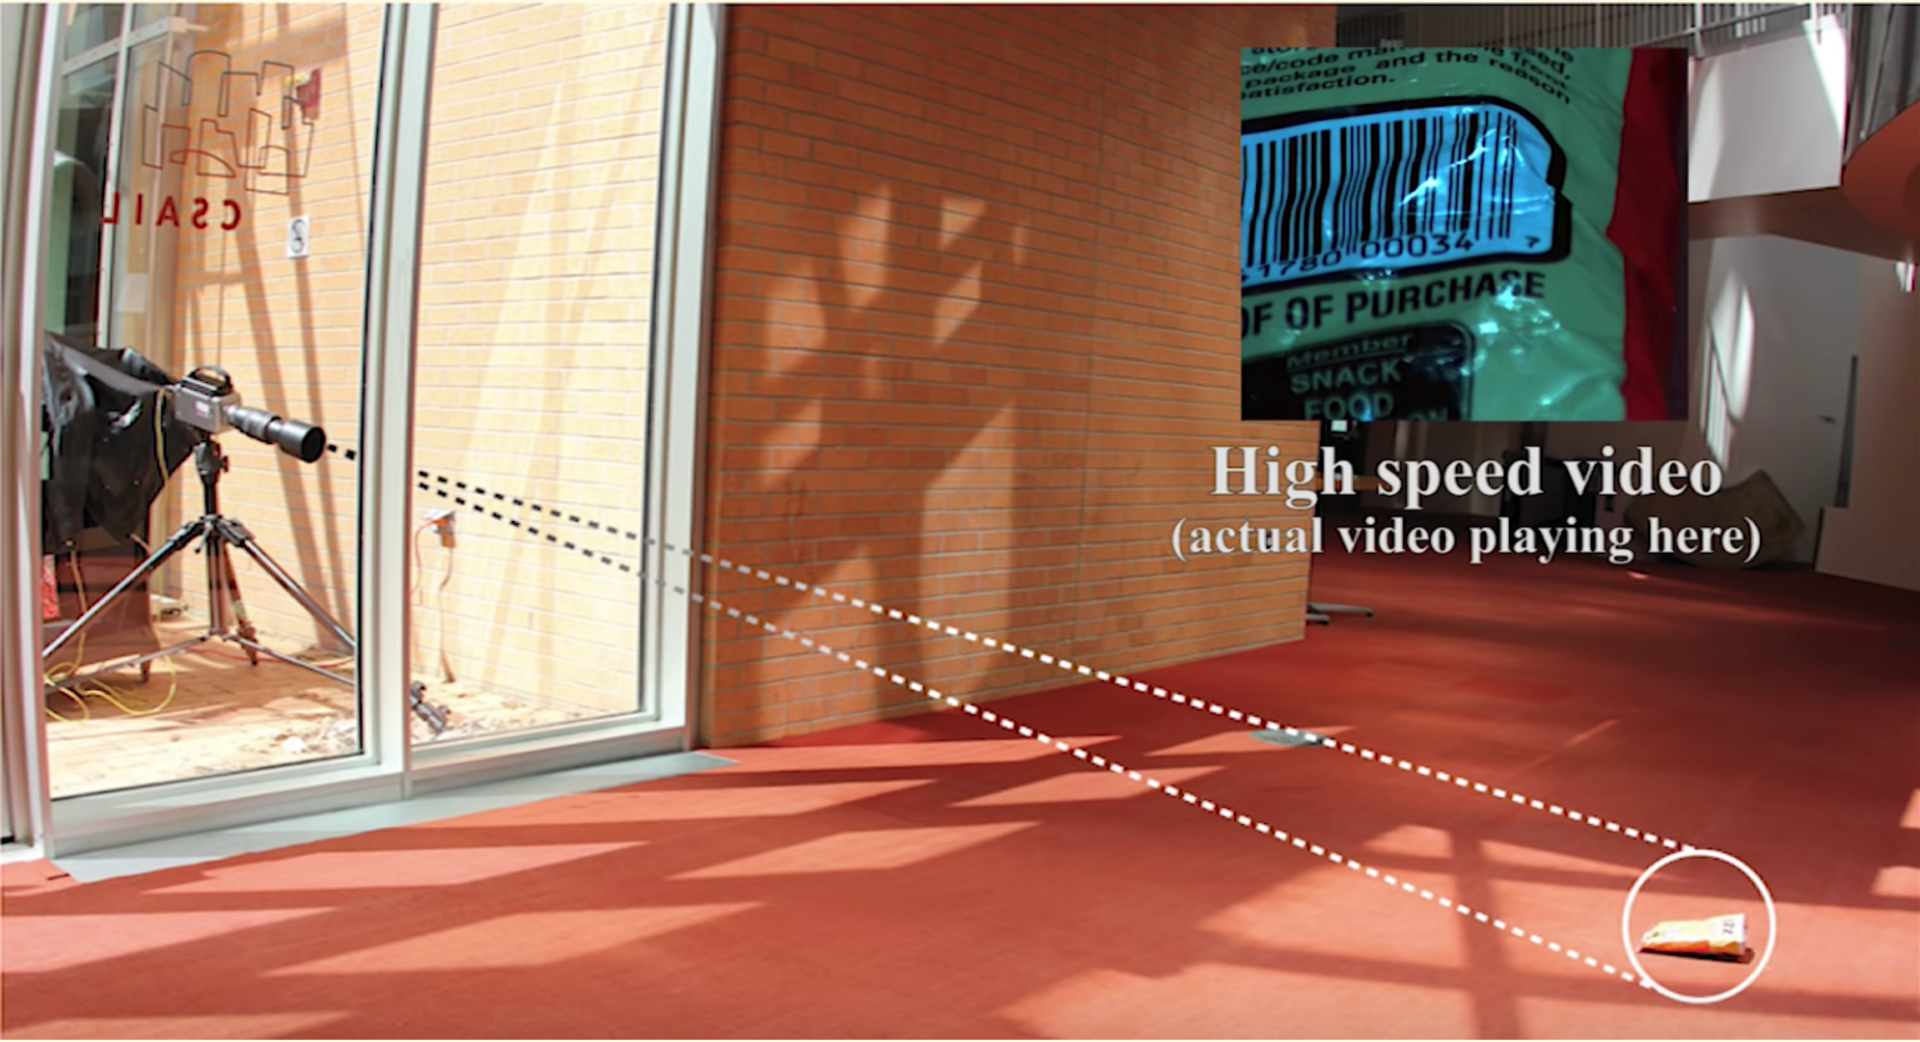
High speed video
(actual video playing here)

Visual Microphone

- Uses video as a “visual microphone” to recover sounds
- By decoding the subtle motions of plants, wrappers, or packaging, an algorithm can infer the sound playing in a room
- Website about project at: http://people.csail.mit.edu/mrub/VisualMic

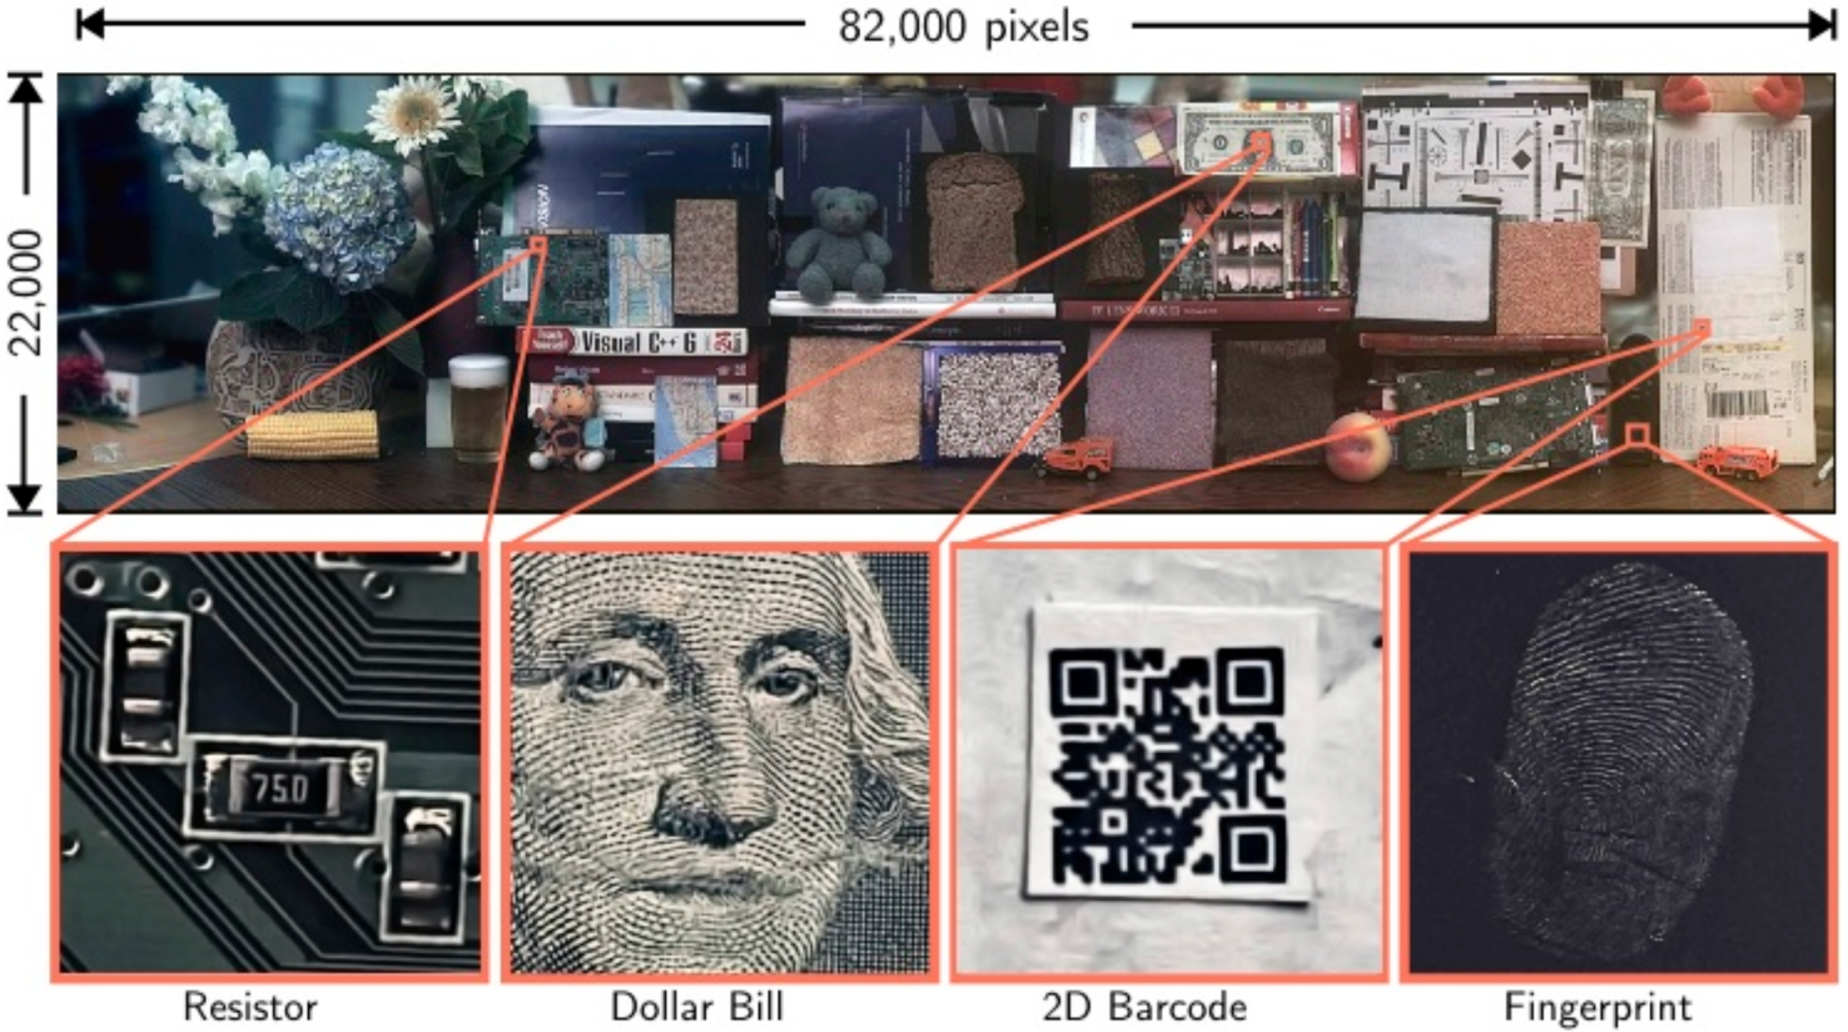

82,000 pixels
22,000
Visual C++ 6
Resistor
Dollar Bill
2D Barcode
Fingerprint

Hi-spatial resolution cameras

- Gigapixel Camera
- Camera arrays with multiple image sensors that capture light, and then are stitched together to create enormous images
- Shree Nayar's lab at Columbia: http://www.cs.columbia.edu/CAVE/projects/
- BigPixel: http://sh-meet.bigpixel.cn/


nature
Hidden
Mannequin
Femto-Camera
Wall
Door
Velten et al, Nature Communications 2012

Hi-spatial resolution cameras

- Femto-Photography Camera
- Can capture *one trillion* frames-per-second
- Able to detect/reconstruct a single photon as it moves through a space
- Can be used to see around corners, or read through a closed book
- Ramesh Raskar, MIT Camera Culture: http://web.media.mit.edu/~raskar/
- https://www.youtube.com/watch?v=JWDocXPy-iQ

Machine learning

- Basic idea: don't tell the computer what to do, give it lots of data and let it figure it out
- Rather than giving the computer instructions, you provide with an extensive training session.
- Initial breakthroughs were in classifying data
- Inverse of classification, can be used to generate new data with specified features


A
B
C
D
E
F

Style transfer

- Leon Gatys' style transfer: http://bethgelab.org/publications/leon+gatys/
- Style transfer is successful at learning (some) features that we can't explain clearly– non-ML algorithms are less successful at describing subtleties of image
- Has now been extended to videos, panoramic/VR, real-time
- Could an ML technique learn how to block a shot? How to light a scene? How to edit a film? How to write a script? Learn a director's style?

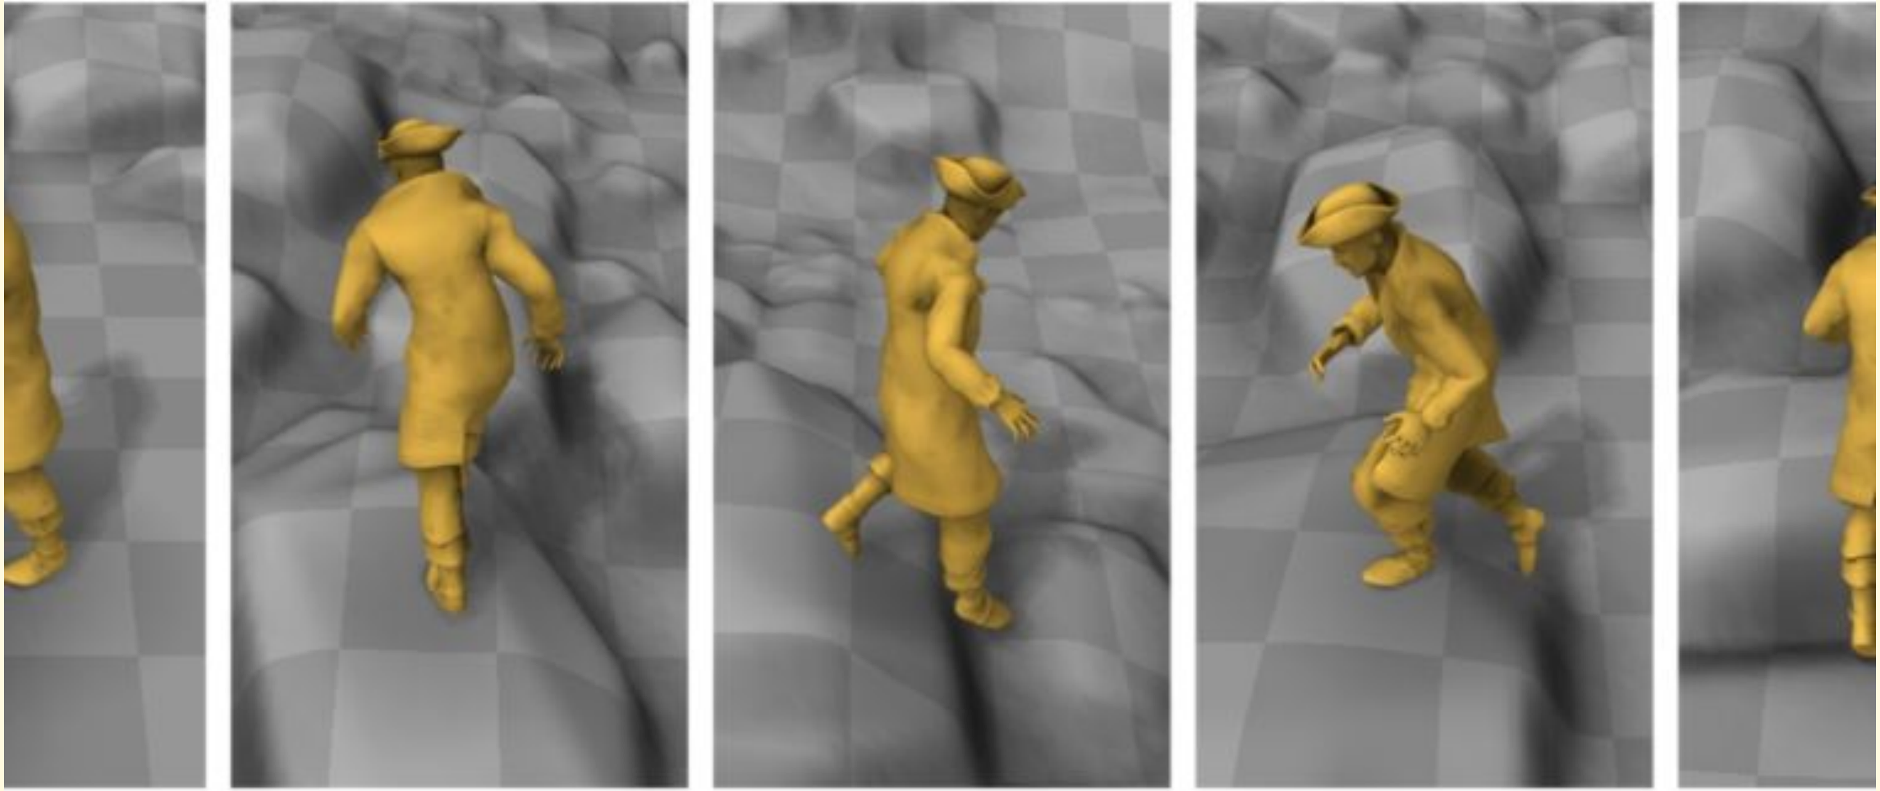

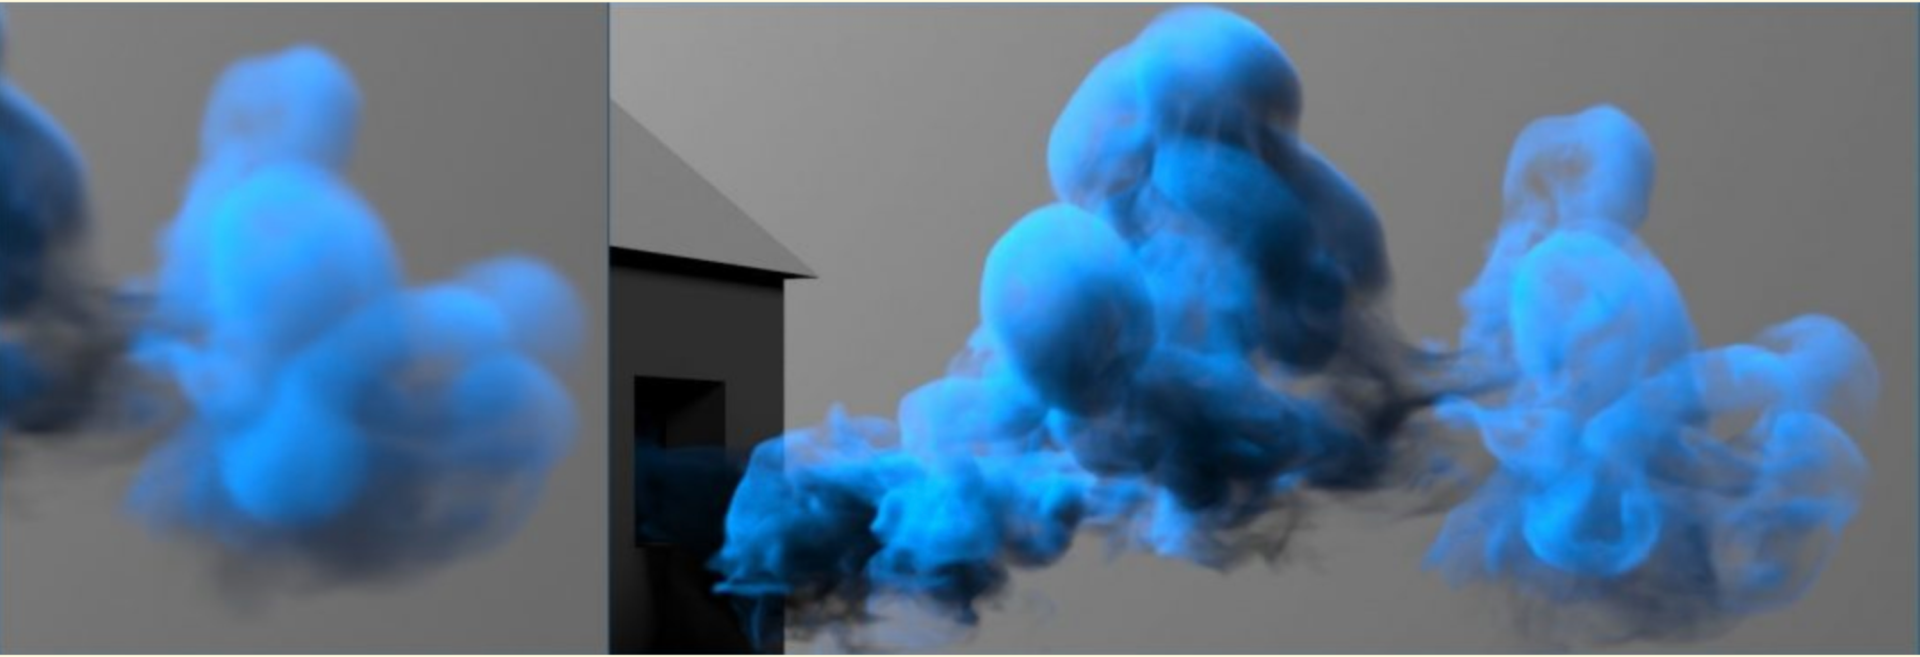

Generating motion

- Generating realistic motion in new environments
- Generating appropriate motion planning for specialized environments (e.g., rockclimbing)
- Generating crowd simulations
- Phase-Functioned Neural Networks for Character Control, https://dl.acm.org/citation.cfm?id=3073663
- TempoGAN: https://ge.in.tum.de/publications/tempogan/


Sogou搜狗
AI ANCHOR

Deep fakes

- Generating realistic video from audio/text
  - AI newscaster, Jordan Peele->Obama, Trump Puppet
  - https://grail.cs.washington.edu/projects/AudioToObama/

- Face2Face real-time face swapping
  - http://niessnerlab.org/projects/thies2016face.html

- Everybody Dance Now
  - https://carolineec.github.io/everybody_dance_now/

- Technology exacerbates our current inability to separate facts from opinions, to think critically, to identify propaganda, etc.

Cinematic machine learning?

- So far, ML is less successful at generating effective stories, plots, cinematic experiences. How come?

- Requires lots of elements simultaneously: lighting, editing, acting, multiple characters and plots, genre expectation, etc

- Assumes contextual information which we ourselves don’t know how to describe/discretize effectively and thoroughly

- However, these technologies are still in their infancy, and no doubt they will continue to increase in sophistication

- Rather than *replacing* an actor, a director, a screenwriter, or a composer, can we develop creative tools that utilize these technologies?

Thursday’s Homework

- Read the articles/Look at code about WaveFrontCollapse and Eulerian Video Magnification

- Project 1 will be assigned on Thursday 1/10

- Reading for Tuesday’s (1/15) discussion will be assigned on Thursday 1/10