

# CS 523: Multimedia Systems

Angus Forbes

[creativecommons.evl.uic.edu/courses/cs523](http://creativecommons.evl.uic.edu/courses/cs523)

# What is this class about?

## 1. Generative Systems / Algorithm Simulations using multimedia data

- Simple rules that create complex, emergent systems
- Cellular automata
- Flocking systems / Swarm behavior
- Genetic algorithms
- NPC behavior
- Simulating biological, ecological, sociological processes

# CS + New Media Arts

- Generative Systems

Simple rules that create complex, emergent systems

- Cellular automata
- Flocking systems / Swarm behavior
- Genetic Algorithms

- Generative Machine Learning Models

Models that can be used to create and simulate, rather than classify or categorize

- GANs, RNNs
- Deep Dream / Inceptionism
- Style Transfer

# CS + New Media Arts

Course is project-based, governed by the idea that you will spend some energy determining interesting, creative, meaningful projects, and then learn what you need in order to create those projects

Instead of (or alongside of) learning specific material and then choosing projects solely as examples to show that you understand particular concepts



# Creative Computational Intelligence

How can creative, multimedia applications amplify our intelligence?

Last week looked at projects by  
Dennis Hlynsky, Memo Atken, CSAIL  
projects that involved experiments with  
manipulating time and the juxtaposition  
of images

# Creative Computational Intelligence

We also talked about different notions of intelligence, and how that could (potentially) be encoded in computational systems.

# Homework from last week?

- Magnified Videos

# ML advances in:

- Realtime Speech Translation
- Identifying Location of a Photograph
- Self-Driving Cars
- Predictive Keyboards
- Gesture Recognition
- Lip Reading
- Product Recommendation
- Tumor Detection
- Speech Synthesis
- Image Processing

# ML projects

Overarching idea:

- Complex systems have too many rules, or rules that are difficult to quantify.
- Rather than try to come up with all of the rules, create a system that can learn the meaningful rules automatically, through examining lots of data where examples of the rules are expressed.
- Train your system on examples where you know the right answer, and then test to see if it works on new examples.

# Discriminative vs Generative

**Discriminative:** Detecting events, Finding patterns, Classifying objects, Recognizing elements

**Generative:** Synthesizing data, Inferring examples, Interacting with models

# Discriminative vs Generative

Inverses of each other:

If I have the knowledge of what features determine whether or not a specific sample belongs or doesn't belong to a particular category or class,

Then I can also use those features to create new samples that are examples of a particular category or class

# Online classifier for handdrawn objects (Google's A.I. Experiments)

[https://aiexperiments.withgoogle.com/  
quick-draw](https://aiexperiments.withgoogle.com/quick-draw)





# Deep Dream

“neural networks that were trained to discriminate between different kinds of images have quite a bit of the information needed to generate images too”

<https://youtu.be/DgPaCWJL7XI>

<https://research.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html>

# Neural Photo Editor

[https://www.youtube.com/watch?  
v=FDELBFSeqQs](https://www.youtube.com/watch?v=FDELBFSeqQs)



# Neural Style Transfer



# Generative Visual Manipulation on the Natural Image Manifold

<https://www.youtube.com/watch?v=9c4z6YsBGQ0>

# Neural Doodle




[https://www.youtube.com/watch?  
v=fu2fzx4w3ml](https://www.youtube.com/watch?v=fu2fzx4w3ml)







# PPGANs



redshank



ant



monastery




# PPGANs



volcano




# PPGANs






a red car parked on the side of a road

a blue car parked on the side of a road



a pizza on a plate at a restaurant

someone is just about to cut the pizza



oranges on a table next to a liquor bottle

a pile of oranges sitting in a wooden crate

# One-shot generalization

[https://www.youtube.com/watch?v=6S6Tx\\_OtvnA](https://www.youtube.com/watch?v=6S6Tx_OtvnA)

<https://www.youtube.com/watch?v=HkDxmnlWIM>

# DeepMind's Deep Q-learning

[https://www.youtube.com/watch?  
v=V1eYniJ0Rnk](https://www.youtube.com/watch?v=V1eYniJ0Rnk)

# Probabilistic Future Frame Synthesis

[https://www.youtube.com/watch?  
v=zidaYS85mCY](https://www.youtube.com/watch?v=zidaYS85mCY)

# Generative Design (Dreamcatcher)

[https://www.youtube.com/watch?  
v=CtYRfMzmWFU](https://www.youtube.com/watch?v=CtYRfMzmWFU)

a crumb of friction milks god

<https://vimeo.com/187931421>

# Face2Face

[http://www.graphics.stanford.edu/  
~niessner/thies2016face.html](http://www.graphics.stanford.edu/~niessner/thies2016face.html)

# Tensor Flow

Software library that makes it easier to conceive of numerical computing on big data using “data flow graphs,” which can represent different kinds neural network configurations, or other ML operations.

- New data types to work with data commonly used in Machine Learning (tensors, or multi-dimensional matrices)
- Workflow: 1. Configure a graph, 2. Create a session to run the graph, 3. Examine results



# Tensor Flow

- Large community of users, supported by Google, lots of tutorials, examples
- Runs on Linux, OSX, and Windows
- Supports CPU or GPU
- Incorporates Keras, a high-level NN library

# Homework for next week

A. Download and set up TensorFlow, go through MNIST tutorial.

B. Two (short) assignments – see PDF on website:

1. Implement the Eulerian Video Magnification code and create a video that accentuates a color or motion, be prepared to explain what new understanding of the scene you have after doing so.

2. Choose an interesting project from CreativeAI.net and present it in class.

C. Two readings (see website)