APS360 Fundamentals of AI

Lisa Zhang

Lecture 11; June 13, 2019

Agenda

First hour:

- Midterm discussion
- Recurrent Neural Network implementation
- LSTM and GRU units

Second hour (tutorial, in MY330):

- Projects
- Group formation
- Project ideas

Midterm

Logistics

- Thursday, June 20
- Room: GB304
- Time: 6:10pm - 8:00pm

Rules

- Closed-book test
- No calculators
- Write in either pen or pencil
- You may not leave in the last 10 minutes of the test
- Bring your T-Card

Topics

- Lectures 1-5 materials, up to but not including RNNs
- Labs 1-4
- Detailed list of topics on the course website
- Themes:
  - Neural Network Training
  - Convolutional Arithmetics
  - Autoencoders
  - Interpreting Training Curves
  - Conceptual Questions

Coding

- You are expected to be able to reason about code, and write or choose small pieces of code.
- You might be asked to explain what some code or parameter does. (If the code isn't something we've seen frequently, I like to include the PyTorch API documentation.)

However...

- You are not graded on Python syntax
- You are not required to remember the order of the parameters
- When in doubt, try your best and write down your intentions and assumptions

Short-Answer Questions

- Using proper terminology is *very* important!
- If the TA cannot understand what you're saying, then we cannot award part marks, *even if you turned out to have the right idea*

We can't read minds, so please be clear with your answers.

Midterm Preparation

- Study questions
- Last year's midterm
- Lab questions

Do you want to build a repository of answers as a class?

Midterm Review

- Next class
- Bring a device (phone/laptop)
- Bring a pen or pencil

Recurrent Neural Network (cont'd)

Simplified RNN Model

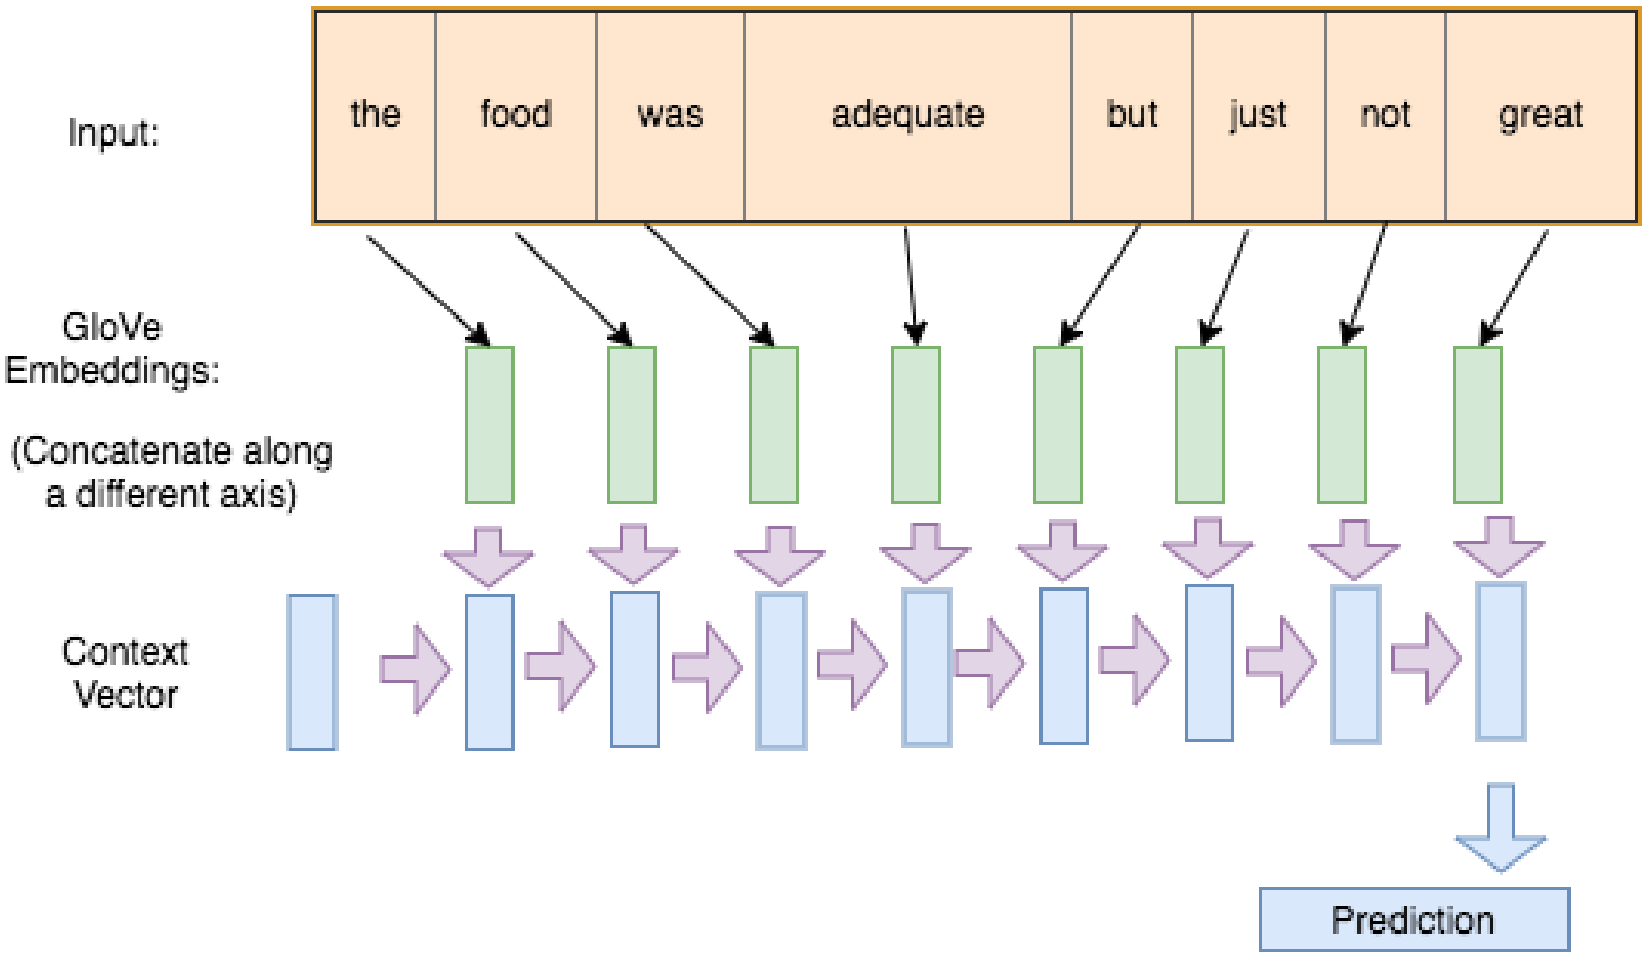

- `hidden = update_function(hidden, input)`
- `output = prediction_function(hidden)`

RNN Model (Colour change)

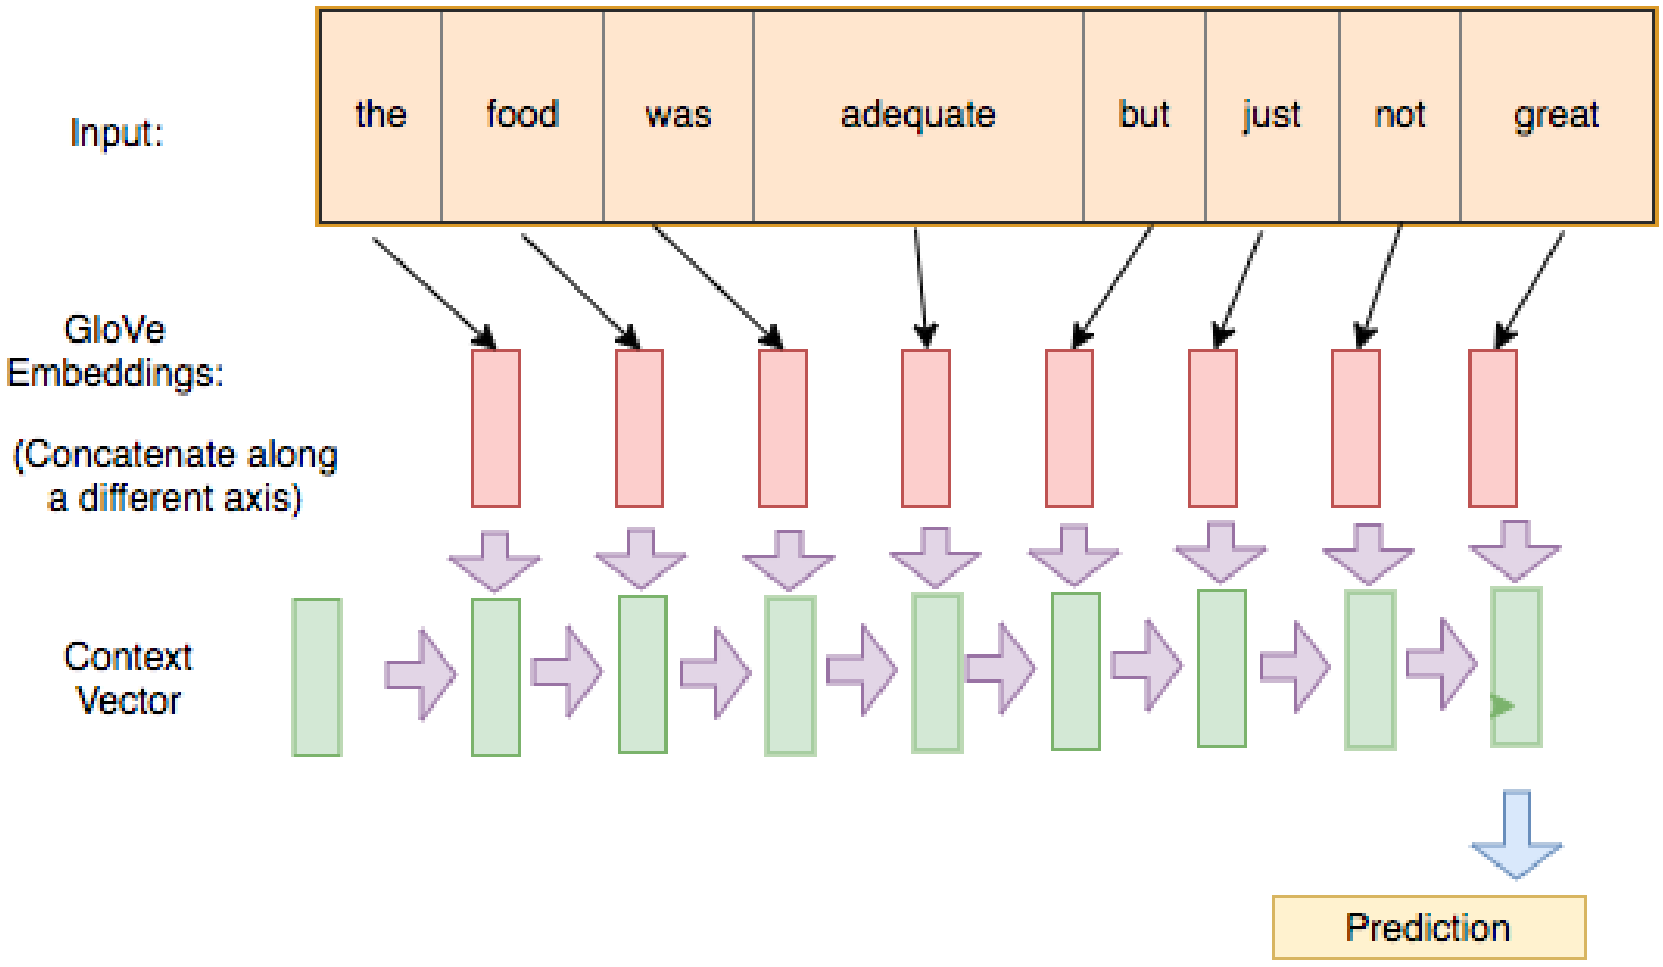

- `hidden = update_function(hidden, input)`
- `output = prediction_function(hidden)`

RNN Model - with Output

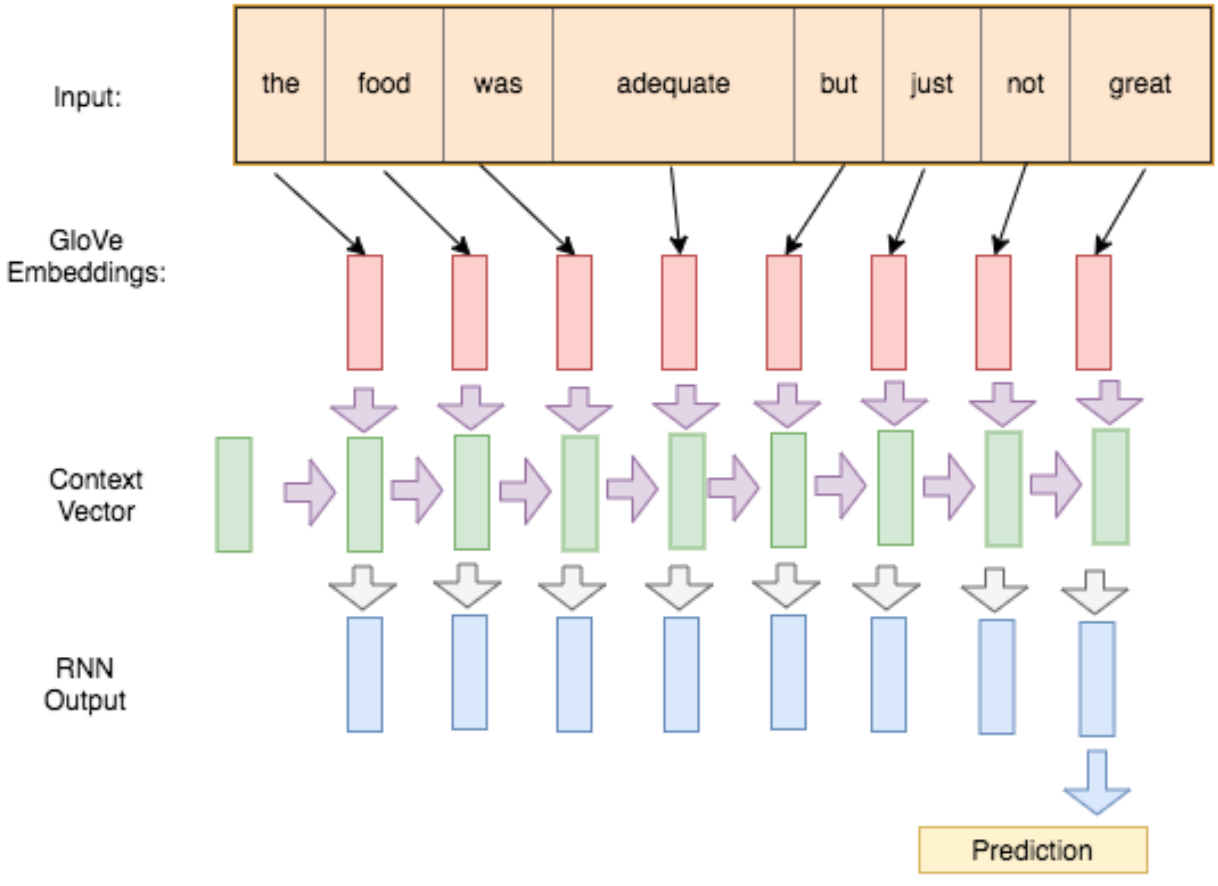

- `hidden = update_function(hidden, input)`
- `rnn_output = output_function(hidden)`
- `output = prediction_function(rnn_output)`

RNN Model Types

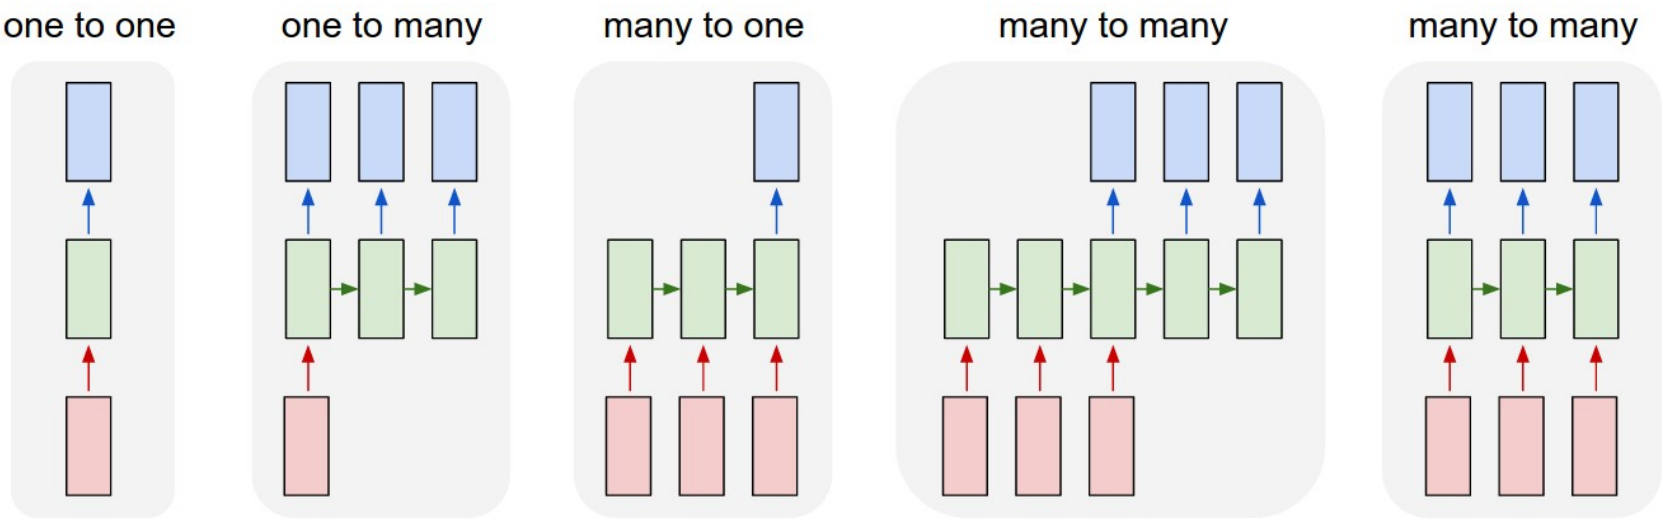

- Focus on many-to-one today (classification task)
- Other variations after reading week (generation task)
- Won't work with many-to-many (trasnlation task)

“Vanilla” RNN

- In an `nn.RNN` module, `rnn_output` is the same as the last `hidden` vector
- In more complex RNN models you’ll see later today, the two values may be different

Sentiment Analysis Model

Three (neural) components:

- GloVe embedding lookup
- Recurrent neural network
- Fully-connected prediction network

Let's take a look!

Batching

- When trainign an RNN, sequences in each batch must all have the same length
- So, sequences that are shorter needs to be **padded** with a special token
- In practice, try to batch similar-length training examples to minimize padding

Learning Long-Term Dependencies

- Historically, Recurrent Neural Networks were hard to train
- Better RNN units for learning long-term dependencies:
  - Long Short-Term Memory (LSTM): requires an extra **cell state**
  - Gated Recurrent Unit (GRU): only requires a hidden state

RNN vs LSTM vs GRU

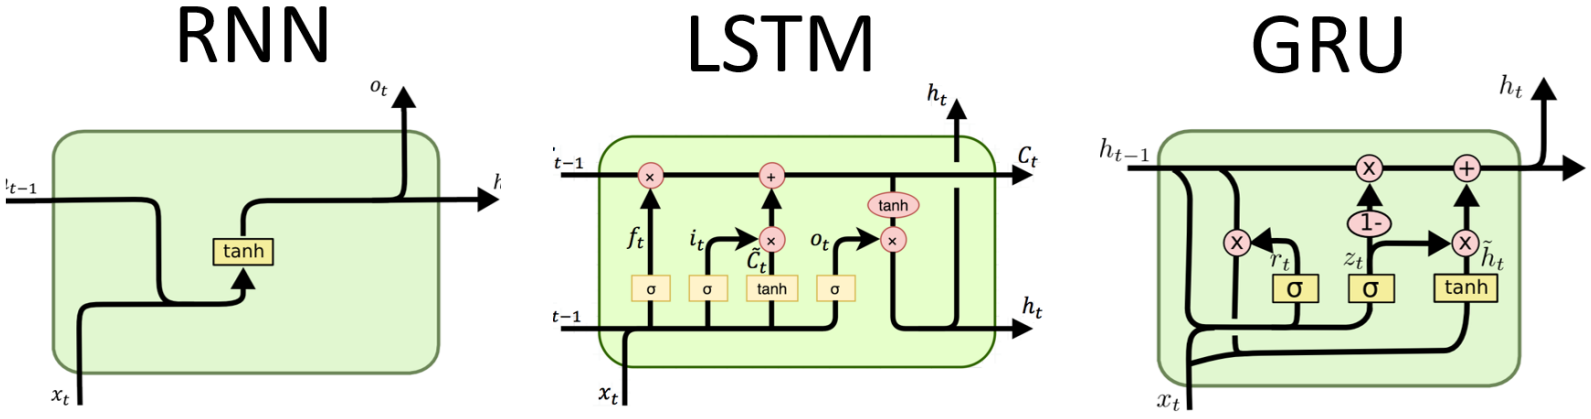

Documentations

https://pytorch.org/docs/stable/nn.html#rnn

APS360 Group Project

Project

Work in a group of 3 to implement a machine learning application in an area of your own choosing.

Why an open-ended project?

- Build practical implementation skills
- Become comfortable with the uncertainties of open-ended work
- Build oral and written communication skills

Groups & Topics

You can choose your own group and topic!

- Topic must make use of machine learning
- There must be some kind of a data collection or cleaning stage
- You must be able to articulate why machine learning is the appropriate tool for your task

Example: age estimation from head shot


tensor(67)
tensor(68.5408


tensor(9)
tensor(9.4429


tensor(25)
tensor(26.7493


tensor(31)
tensor(30.2026


tensor(37)
tensor(36.6638


tensor(69)
tensor(60.9164


tensor(46)
tensor(23.6994


tensor(23)
tensor(27.9990


tensor(59)
tensor(54.0926


tensor(59)
tensor(52.2167

Example: South Park Script Generation

Kyle: Hello?
Time Child: Yeah, Kyle. Maybe we can get like what happened in the sand?
Linda: We have to go and be the rosey to part of the school hallway monitor.
Chef: Phahahahahaaaa.
Mr. Brown: Please, come in.

Figure 6. Example output from final model

Example: Skyline Generation using GAN

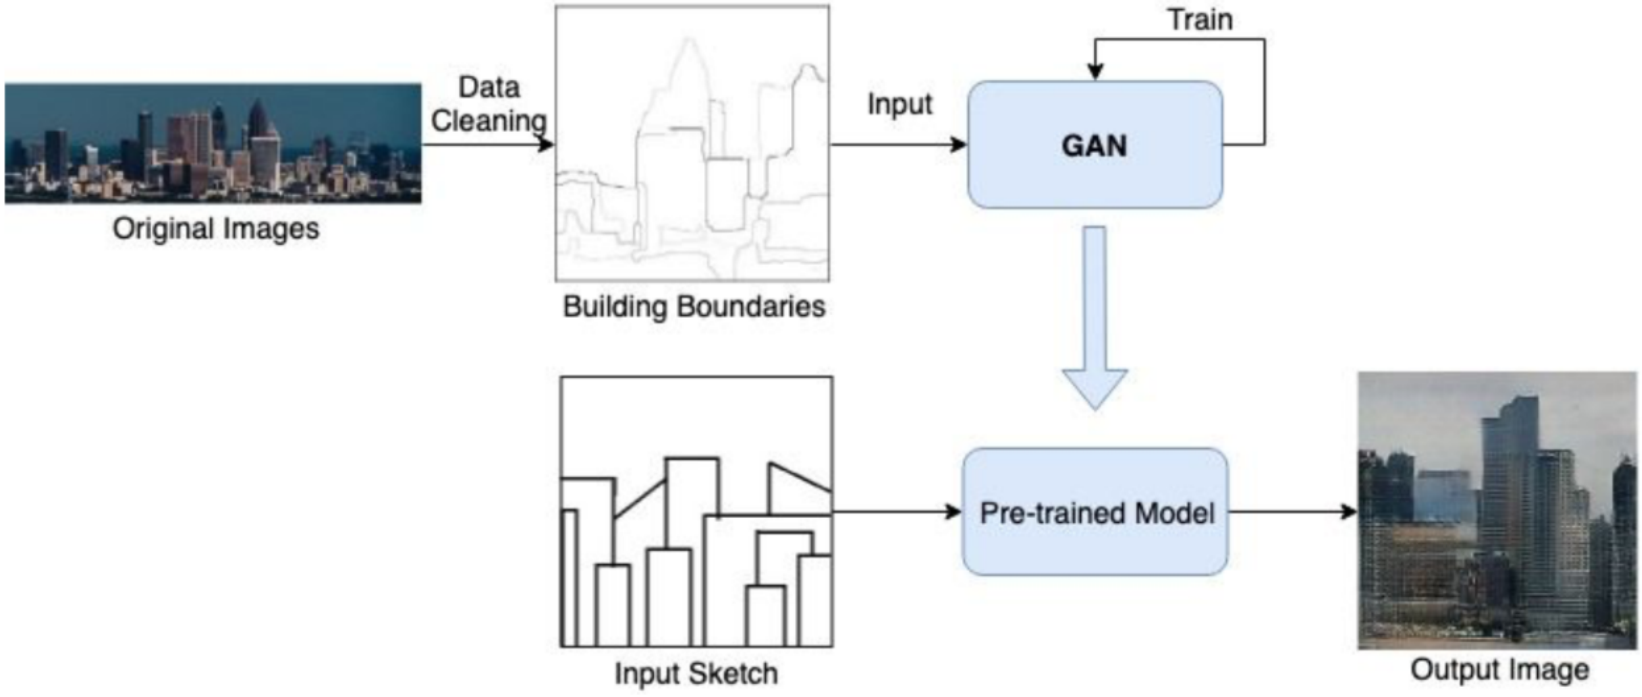

Example: Altering emotions of photographs

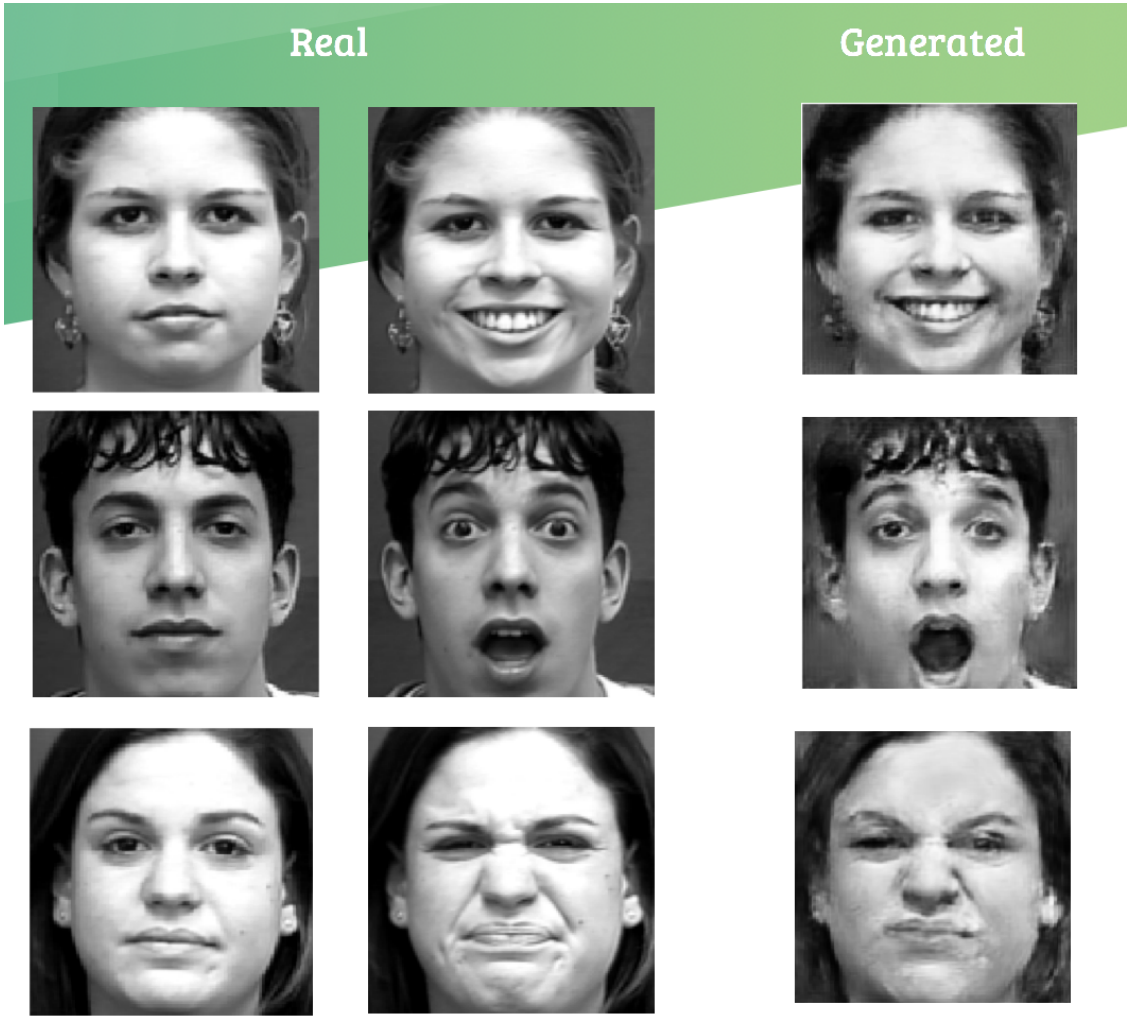

Project Topics

- TAs are here to help you decide on possible topics
- Every project in this class must be unique
- Every group should use a different dataset

Every group will be assigned a TA mentor, who will work with you and answer questions.

Uniqueness Approval

Send me an email with:

- the names, utorids, and email addresses of your group members
- the project idea that you wish to work on
- the data set(s) that you plan on using, and how you plan on collecting/cleaning/processing the data

I will either approve your idea, or let you know that another group has already taken the idea.

Your group must receive uniqueness approval by June 21, 2019. Otherwise there will be a 20% deduction applied to your Project Proposal grade.

Project Timeline

- Project Uniqueness: June 21
- Project Proposal: June 27
- Progress Meeting with your TA mentor
- Progress Report: July 24
- Presentation: last week of class
- Project Report: last day of class

Project Rules

- University of Toronto rules on plagiarism apply. All the code that you write must be your own.
- All team members must contribute to your project equally.
- You must complete your project **either** using Google Colab, **or** inside Github repository, so the TAs can track progress and contribution.
  - If you use Github, you have the option of writing your final report as a Github README. Speak to me if you would like to do this.
- Presentation slides must be created using Microsoft Powerpoint (depending on number of teams, this rule may be relaxed).

Project Proposal

- Due June 27, 9pm
- Goal is to demonstrate that you:
  - know what the goals of your project
  - know what dataset you will use
  - have a rough idea of what model you'll build
  - have a clear idea of how your team will work together

Rubric is posted on the course website.

Finding a team

- Piazza - “Search for Teammates!”
- Today!

What kind of project do you want to do?

- Type of data (image? text? biomedical?)
- Type of problem (classification? regression? generation?)
- Type of model (RNN? autoencoder? GAN?)
- Problem domain (biomedical? self-driving?)

Rest of Today

- Project ideas?
- Project team?
- Speed Dating?