Machine Learning for Artists and Designers (3 of 3): Parag Mital

This Document: https://goo.gl/Pf1fcT
The Other (live) Document: https://goo.gl/kGJ4z2

Title

“Creative Applications of Deep Learning w/ TensorFlow” Workshop, by Parag K. Mital

Time and Place

When: Saturday-Sunday, Mar. 24-25, 2017, 10am-6pm
Where: Frank-Ratchye STUDIO for Creative Inquiry
Who: Workshop by Parag Mital, for invited CMU students from
Sponsors: Carnegie Mellon’s School of Art, School of Architecture, School of Computer Science, Carnegie Institute of Technology, Human Computer Interaction IDEATE and Emerging Media Masters
Organizer: Ali Momeni

Links:

- This Document: https://goo.gl/Pf1fcT
- Discussion/Sharing: Creative Machine Learning Facebook Group
- GitHub Repo: https://github.com/pkmital/CADL
- Ali’s Notes, not facebook: Live Google Doc

Abstract

AI has found it way into nearly every aspect of our lives, from dictating our news, to fighting our wars, to monetizing our personal data through advertising. In this 2 day workshop, we'll try to disentangle the hype from reality and understand what is really possible with AI with a hands on approach. Through discussions and guided coding exercises, we'll start to explore our own datasets using a state-of-art technique to building AI known as Deep Learning. This workshop is geared towards creative practitioners with some coding background looking to develop an interdisciplinary practice with AI. Through this understanding, you will be better equipped to become more critical of AI, allowing you to potentially subvert and/or communicate issues with current AI practices, or explore creative and potentially inspiring directions for AI.

Workshop Scope

The workshop will start with a contextual and historical discussion about how AI and machine learning has developed. We'll then work on understanding common approaches to take when building datasets required for machine learning. We'll see what data actually looks like when presented to an AI, what is missing from it, how state of the art algorithms are used in production, what its strengths and weaknesses are, where we can witness it, how it shapes our world, how it differs from human perception, and speculate with a critical eye where it is all heading.

We'll then focus on building a technical understanding of Machine Learning and in particular, a state of the art approach for machine learning known as Deep Learning. This format will be a compact version of the Kadenze course, "Creative Applications of Deep Learning w/ TensorFlow". We'll look at the material from this course to understanding what's possible using TensorFlow, a Python and C++ library for building computational graphs. We'll then apply some of the algorithms we learn to our own data and start to ask some questions about what the algorithms are capable of learning and what is missing. The approach we take will be to understand the mechanics of the algorithm, to look inside the "black box", and use our understanding of what the algorithm is doing in order to help us understand what is and is not possible with AI.

The last day will be spent focusing on the creative potential such algorithms have and how they might integrate with your existing practice. We'll open the discussion to participants and structure the rest of the workshop around questions and potential directions suggested by participants.

Participant Prerequisites

- You should have some coding experience, ideally in Python though a background in Processing, Javascript, openFrameworks, or similar should suffice (i.e. at least one developed/completed application).
- You should have a clear understanding of basic coding primitives such as int versus float, understand program flow such as the use of if and for statements, and be somewhat comfortable with installing new software using the terminal/shell environment of your operating system, though guidance will be given along the way on these topics based on attendance.
- You should have completed a high school math curriculum, though we will have a friendly refresher of basic concepts such as equations and variables. This background will be used to cover fundamental though difficult topics required for building AI including basic statistics, matrix math, and convolution.
- You should also have a somewhat recent laptop (mine is 4 years old), capable of installing about 4 GB of material. Participants should then download the latest "Anaconda" distribution for their OS for Python 3.6.0: https://www.continuum.io/downloads (NOT Python 2.7!).
- Participants should also feel free to browse through the CADL repo here before coming to the workshop: https://github.com/pkmital/CADL

Agenda

Day 1

- 10am-11am:
    - Introduction
    - About You
    - About Me
    - Cultural Background
        - How do machines sense the world?
        - What can machines understand from such sensors?
        - How do machines learn?
        - How do humans learn?
        - What are some examples of machine intelligence?
        - What are some problems/fears with existing AI?
        - What are approaches we should take to resolve our fears?

        - What are examples of artists/media discussing potential fears/future s/realities of AI?
        - What are some initiative to understand/minimize risk in AI?
        - What are some potential problems machine intelligence may help solve?
    - History
- 11am-12pm
    - Data
        - What are datasets?
        - Labeled vs. unlabeled data
        - Common datasets used in ML/AI
        - What are common limitations amongst datasets?
        - Create your own dataset
        - Structuring data on the OS
    - Loading data
        - os.listdir
        - Loading an image
        - RGB Image Representation
        - Understanding data types and ranges (uint8, float32)
        - Visualizing your data as images
- 12pm-1pm
    - Image Manipulation
        - Cropping images
        - Resizing images
        - Cropping/Resizing Images
        - The Batch Dimension
        - Audio
            - scipy.io
            - wave
            - Bregman, Marsyas, Essentia
        - Text
            - open
            - Pandas
            - csv, json, ...
- 1pm-1:30pm: LUNCH
- 1:30pm-2:30pm (while eating)
    - Preprocessing datasets
        - Sound datasets
            - Magenta

                - PYMIDI
                - Bregman, Marsyas, Essentia
            - Text datasets
                - SpaCy
                - NLTK
                - GenSim
            - Image datasets
                - Cropping
                - Resizing
                - Colorspaces
                - Datatypes
                - OpenCV
- 2:30pm-3:00pm
    - Basic Statistics
        - Data structure
            - Image Shape
            - Batch Dimension
        - Mean and Bias
        - Standard Deviation and Variance
        - Factors / Invariances / Representations
- 3pm-3:30pm
    - Scope of Learning
        - What does an algorithm really learn?
        - Dataset creation
        - Amount of data
        - Transfer learning
        - Processing versus Understanding
        - "End-to-end" systems
- 4pm-5pm:
    - TensorFlow Basics
        - Variables
        - Tensors
        - Tensor Shapes
        - Graphs
        - Operations
        - Sessions
    - Convolution
        - Gaussian Kernel
        - Filtering an Image

        - Gabor Kernel
- 5pm-6pm:
    - Gradient Descent
        - Defining Cost
        - Minimizing Error
        - Backpropagation
        - Minima / Optima

Day 2

- 10am-11am:
    - Questions and Answers on Da 1
- 11am-12pm:
    - Neural Network
        - Defining Cost
        - Training Parameters
        - Stochastic Mini Batch Gradient Descent
        - Over vs. Underfitting
        - Nonlinearities
        - Going Deeper
- 12pm-1pm
    - Image Regression
        - Description
        - Building the Network
        - Training
- 1pm-1:30pm LUNCH
- 1:30pm-2:30pm
    - Autoencoder
        - Fully Connected
        - Convolutional Autoencoder
        - Variational Autoencoder
- 2:30pm-3:30pm
    - Pre-trained Networks
        - VGG
        - Inception
        - WordNet
        - Trump
- 3:30pm-6pm: LAB TIME

Pre-workshop Setup

Participants should do the following before coming to the workshop:

- Sign up with Ali and confirm the following:
    - Are they bringing a laptop
    - What OS is it running
- Install Anaconda Python 3.6: https://www.continuum.io/downloads
- Participants should also feel free to browse through the CADL repo here before coming to the workshop: https://github.com/pkmital/CADL