

A4834/6: Data Mining the City
DMC4 - Intelligent Design Machines

Instructor: Danil Nagy (dn2216@columbia.edu)

Meeting time: Wednesdays, 7:00pm-9:00pm

“Telling the future, when it comes right down to it, is not solely a human yearning. It is the fundamental nature of any organism, and perhaps any complex system. Telling the future is what organisms are for.”

- *Kevin Kelly, “Out of Control”*

Context:

The idea of an artificial intelligence fascinated technology thinkers and futurists even before our current era of computing had taken hold. Half a decade ago, early computer scientists theorised the possibility of creating computer systems that could mimic the organization and functions of the human brain, thereby lending the computer a form of intelligence similar to our own. However, it is only in the last decade or so that computer systems have become powerful enough to actually put these ideas into practice. In the last few years, even personal devices have reached a level sufficient to put these tools in the hands of average computer users.

This rapid technological development has led to an explosion of the field of Machine Learning in recent years, going from relative obscurity at the turn of the century, to a constant presence in our lives today. The tools of Machine Learning have revolutionised many industries including scientific research, banking, insurance, and medicine. Whether we realize it or not, they have also had a deep and profound impact on our everyday interaction with technology and with each other. Everything from online search, to social media, to new developments in digital assistants such as Apple’s Siri and Microsoft’s Cortana have been made possible by this technology. These are the algorithms that tell us what we want to buy on Amazon, what we want to watch on Netflix, and who we want to be friends with on Facebook.

While each of us interacts with Machine Learning algorithms on a daily basis, most of us have no idea how this technology actually works, or whether it can be useful to our own practice. While the low level implementation of these algorithms requires a knowledge of advanced mathematics and computer programming, the core concepts are actually fairly simple and intuitive, and there are now many tools available that allow even non-technical users to access the potentials of this technology. This seminar will introduce you to these concepts, and explore how they can be useful for our practices in urban planning, architecture, and design.

Expectations and Grading:

The class is divided into two sessions. Session A is devoted to an accelerated crash course in the tools and technologies of Machine Learning. Students will be expected to attend all lectures, which will introduce them the core concepts of this technology. They will also complete a series of weekly technical labs that will implement basic applications of Machine Learning algorithms for solving traditional problems in artificial intelligence.

Session B is devoted to targeted discussions and workshops focusing on how these technologies can be applied to urban planning, architecture, and design. At this time students will work either individually or in groups to develop a research project that applies Machine Learning to a given problem in their field.

The class welcomes students from all departments within GSAPP. All work will be done within the Python programming language. Although prior experience with coding and programming will be useful, it is not a prerequisite for taking the course. The material in Session A will cover all the technical skills you need to know, but be prepared that this can be a very challenging class, especially if you have never worked with code before. Session A can be taken on its own as an introduction to Machine Learning, but participation in Session A is required for taking Session B.

| | |
|---|-----|
| Class Attendance and Participation | 25% |
| Individual Labs (Session A) / Final Project (Session B) | 75% |

Schedule:

Session A – Introduction to Machine Learning

| Date | Lecture | Assignment Due |
|---------------------------|--|--|
| Sept 7 | Introduction to machine learning <ul style="list-style-type: none">• What is artificial intelligence?• How can a computer learn? | |
| Sept 14 | NO CLASS | |
| Sept 21 | Introduction to computer programming and the Machine Learning stack <ul style="list-style-type: none">• Software used in class• 5 basic elements of computer programming | Lab 1: Setting up environment |
| Sept 30 9:00am | Neural Networks part 1 - foundations <ul style="list-style-type: none">• Perceptron model• Activation function 1 - binary• Intuitive example• Linear algebra review• Network architecture• Softmax function• Cost function 1 - mean square error (MSE)• Forward propagation• Activation function 2 - sigmoid | Lab 2: Python basics (rescheduled) |
| Oct 5 | Neural Networks part 2 - training <ul style="list-style-type: none">• The gradient descent algorithm• Weight initialization• Cross validation (bias vs. variance tradeoff)• Regularization strategies (L1, L2, dropout)• Cost function 2 - cross entropy• Introducing non-linearities (ReLU function) | Lab 3: Classification |
| Oct 12 | Neural Networks part 3 - deep learning <ul style="list-style-type: none">• Convolutional Neural Networks (CNN)• Working with image data• Convolution layers• Pooling | Lab 4: Regression |
| Oct 19 | Neural Networks part 4 - learning from sequences <ul style="list-style-type: none">• Recurrent Neural Networks (RNN)• Working with time series data• RNN architecture• LSTM networks | Lab 5: Image prediction |

Session B – Learning for Design

| Date | Topic | Assignment Due |
|-------------|---|-------------------------------------|
| Oct 26 | Introduction to research project Guest Lecture - Dr. Tat Lam Professor, Chinese University of Hong Kong CEO and founder, Shanzhai City | Lab 6: NLP text prediction |
| Nov 2 | Research project presentations | Group/individual research proposals |
| Nov 9 | Workshop 1 - Urban data - GIS | |
| Nov 16 | NO CLASS | |
| Nov 23 | Workshop 2 - 3d data - Rhino/grasshopper - visualization - Python | |
| Nov 30 | Individual meetings | Group interim proposals |

Dec 16 **Final Presentations**
6-9 pm
114 Avery Hall

Session B - research assignment

Using the tools developed in class, develop a research project that applies these tools to solve or gain insight on a particular problem in your field (Graphic/Industrial/Architectural/Urban design, Urban planning, Engineering, etc.) The research project can be developed individually or in small groups of up to 4 students. Each individual/group must complete the following tasks as part of the assignment:

1. Develop a unique research hypothesis about your chosen topic or issue
2. Develop a machine learning methodology to test the hypothesis using a version of the neural networks covered in class
3. Gather data as needed to support your hypothesis
4. Clean up and reformat data so that it can be processed with a neural network
5. Implement your neural network using Python, Tensorflow, and/or Keras
6. Train your neural network and analyze the results
7. Represent your results graphically and through text

Deliverables:

1. Nov 2nd - topic presentations
 - a. Create a series of slides to present your initial research topics to the class. The presentation should cover the following points [20% of grade]:
 - i. What is the topic/problem you are researching
 - ii. How will you apply machine learning to this problem
 - iii. What kind of machine learning application will you develop (supervised vs. unsupervised, classification vs. regression, predictive vs. generative)
 - iv. What kind of data do you need to train the model and how you plan to acquire it
 - v. What type of neural network model will you be using (basic multi-layer neural network, CNN, RNN, hybrid, etc.)
 - vi. Any challenges you foresee for the project
2. Dec 16 - final presentations
 - a. Presentation (as pdf or powerpoint) including the following sections as a guide but focusing on graphics and visualizations to communicate your research to the review audience [40% of grade]
 - i. Abstract
 - ii. Introduction and hypothesis
 - iii. Methodology

1. Data collection
2. Data processing
3. Model construction
4. Model training
5. Visualization

iv. Conclusions and future steps

3. Dec 19 - final deliverable

- a. By the end of the day, email me with the following documents [40% of grade]:
 - i. Your final pdf or powerpoint presentation
 - ii. All data and code used in your project
 - iii. Up to four (4) high-res (at least 3000 x 2000 pixel) images representing the project. These images will be submitted to Abstract (GSAPP's yearly publication) for publication and archiving by the school so you should focus both on their visual impact, as well as how well they communicate your research project and the machine learning process to a non-technical audience.

* Final assignment counts for 75% of your Session B grade

Appendix - Project development workflow

The following step-by-step process may be helpful to guide your development of the research project for Session B:

- Step 1. Choose a topic or area of interest
- Step 2. Catalog the data you already have access to about this topic
- Step 3. Catalog the insights you want to get from the data - what is the question you are trying to answer, or the hypothesis you are trying to prove
- Step 4. Match the insight you want in [Step 3] with specific machine learning models or applications based on the following categories:
 - a. Supervised vs. unsupervised - does your data contain labels which you'd like to predict, or do you want the computer to find structure in unlabeled data
 - b. Classification vs. regression - are you predicting a categorical or continuous variable (this will be informed partially by the data you have [Step 2])
 - c. Predictive vs. generative - will you use the model to make predictions or summarize the data, or to create new data?
- Step 5. Design framework for machine learning model (neural networks in our case)
 - a. Select the type of neural network you will use (basic multi-layer neural network, CNN, RNN, hybrid, etc.)
 - b. Develop model architecture and initial hyper-parameter settings
- Step 6. Determine what additional data is needed to train the model and get the additional data
- Step 7. Train model and get results
- Step 8. Analyze model and evaluate insights gained. If necessary go back to [Step 5]