Python4ML

An open-source course for everyone

Team: James Hopkins | Brendan Sherman | Zachery Smith | Eric Wynn
Client: Amirsina Torfi
Instructor: Dr. Edward Fox
CS 4624: Multimedia, Hypertext, and Information Access
Virginia Tech, Blacksburg VA 24061
5/12/2019

Outline

- Summary
- Deliverables
- Documentation
- Testing Plans
- Post-Semester Work
- Lessons Learned

Summary

Creating an open-source guide to Machine Learning using Python

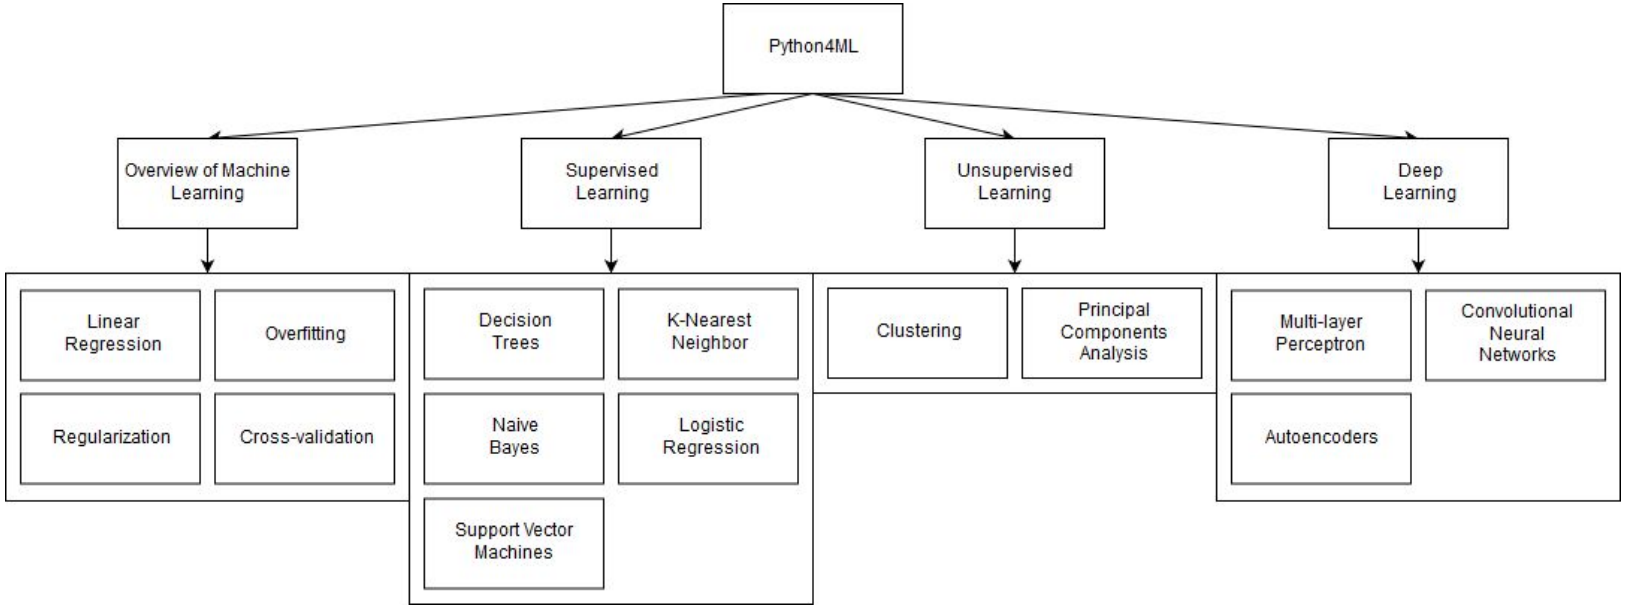

Deliverables

Course website

- Contains all module write-ups
- Entire site can be downloaded as PDF or other formats
- Links to code examples on GitHub

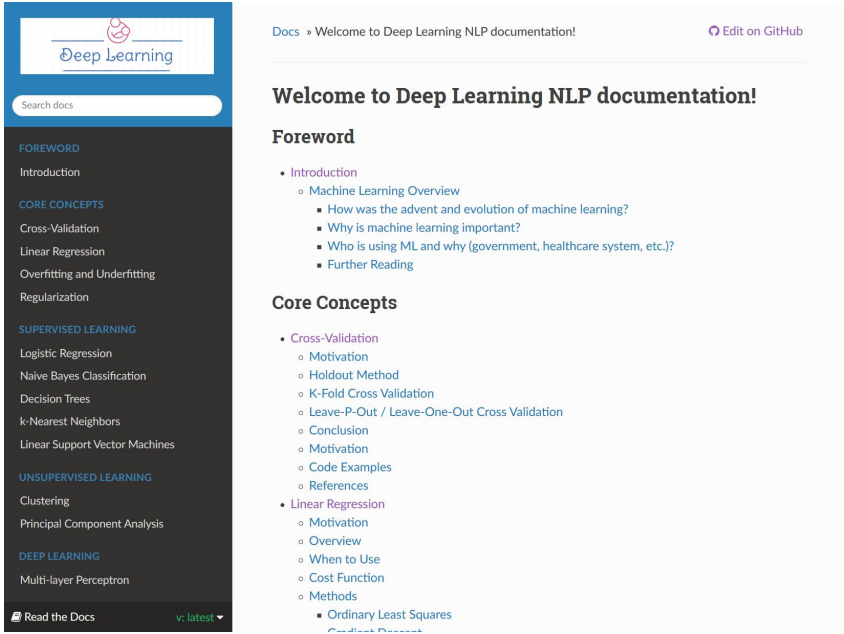

Deliverables

GitHub repository

- All course content is open-source
- Anyone can contribute and suggest changes
- Highly structured system for files

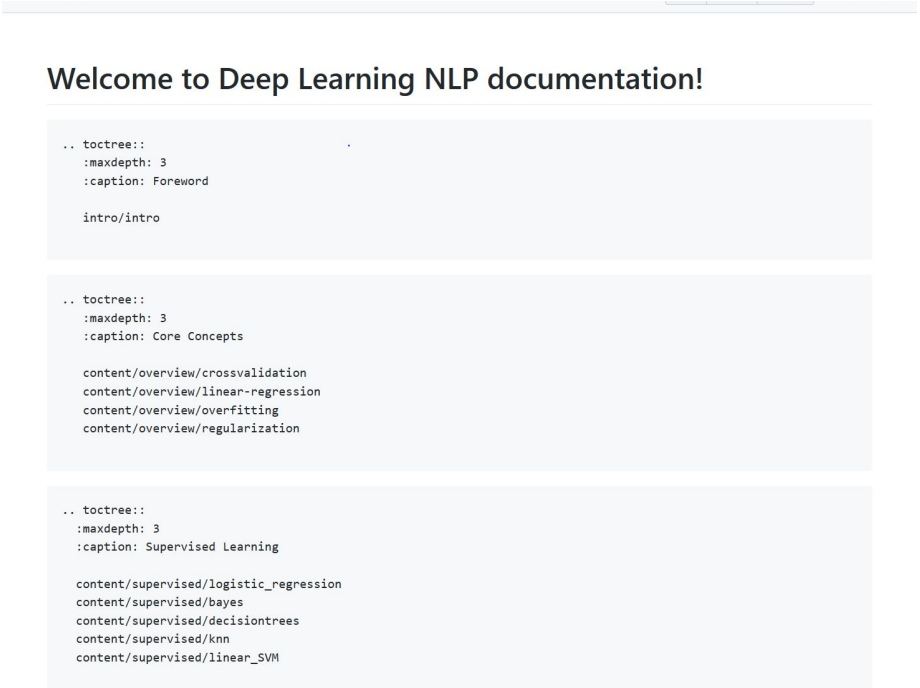

Welcome to Deep Learning NLP documentation!

```
.. toctree::
   :maxdepth: 3
   :caption: Foreword

   intro/intro
```

```
.. toctree::
   :maxdepth: 3
   :caption: Core Concepts

   content/overview/crossvalidation
   content/overview/linear-regression
   content/overview/overfitting
   content/overview/regularization
```

```
.. toctree::
  :maxdepth: 3
  :caption: Supervised Learning

  content/supervised/logistic_regression
  content/supervised/bayes
  content/supervised/decisiontrees
  content/supervised/knn
  content/supervised/linear_SVM
```

Documentation

Documentation written in reStructuredText (rST), a form of markup

Seamless integration with Sphinx

```
+---------------------+---------------------------+-------------------------------+
| Title               | Code                      | Documentation                 |
+=====================+===========================+===============================+
| Linear Regression   | `Python <linpython_>`_    | `Tutorial <lrtutorial_>`_     |
+---------------------+---------------------------+-------------------------------+
| overfitting         | `Python <overpython_>`_   | `Tutorial <overtutorial_>`_   |
+---------------------+---------------------------+-------------------------------+
| regularization      | `Python <regpython_>`_    | `Tutorial <regtutorial_>`_    |
+---------------------+---------------------------+-------------------------------+
| cross-validation    | `Python <crosspython_>`_  | `Tutorial <crosstutorial_>`_  |
+---------------------+---------------------------+-------------------------------+
```

Clean tables

```
.. figure:: _img/LR.png

   **Figure 1. A sample data set with a linear relationship** [`code`__]

   .. __: /code/overview/linear_regression/linear_regression.py
```

Annotated, linked figures

```
.. code-block:: python

   # Create a linear regression object
   regr = linear_model.LinearRegression()
```

Embeddable code

Demo

https://machine-learning-course.readthedocs.io/en/latest/

Project Stats

- 6500+ lines of content
- 100+ pages of course documentation (another 60+ in the final report)
- 70+ sources
- 25+ unique python examples

Testing Plans

We have been testing scripts as they are added, but the whole team will come together to re-test each one for the final deliverable.

Every navigation link in the site needs to be tested.

We had sample users go through the modules and provide us with feedback.

End of Semester Plans

- Complete the last two topics on Neural Networks
- Build final site, test navigation, and test every script
- Get feedback from user testing

Python4ML

A Machine Learning Course for Everyone

Topics

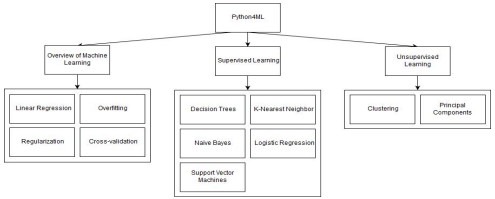

WHAT?
Python4ML is an open-source course for machine learning using the Python programming language.

HOW?
The course is made up of reStructuredText documents and example programs written in Python, using libraries such as scikit learn.

WHO?
This course is being developed with Virginia Tech's Open Source for Science organization, led by our client Amirsina Torfi.

WHERE?
The course is available on GitHub. The code and live course site can be found by scanning the QR code on this poster.

WHY?
The course is aimed at those with little knowledge of machine learning. We want to facilitate education in an open-source context, bringing important topics together in a high-level overview of ML.

Live Course Site

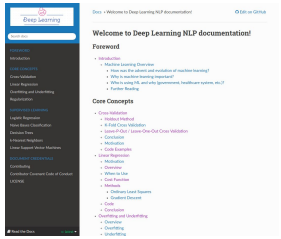

Overview

Linear Regression

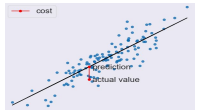

Overfitting / Underfitting

Regularization

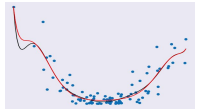

Cross Validation

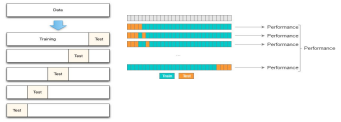

Technologies

matplotlib

Supervised Learning

Decision Trees

Logistic Regression

Support Vector Machines

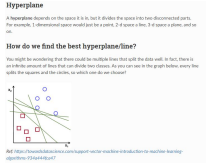

Unsupervised Learning

Clustering

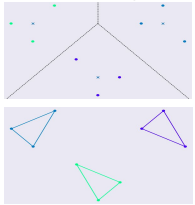

Principal Component Analysis

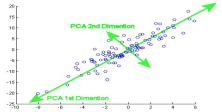

K-Nearest Neighbors

Naive Bayes

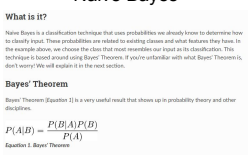

Learn More

Team: James Hopkins, Brendan Sherman, Zachery Smith, and Eric Wynn

Client: Amirsina Torfi,
Head of Open Source for Science @ VT

4/30/2019
Instructor: Edward A. Fox

Scan me

Post-Semester Work

This course can be expanded outside this capstone or in future semesters by:

- Writing more example scripts and content
- Integrating the project with Docker:
  - A containerized system, similar to virtual machines
  - Can pre-package dependencies together for each section and provide users with a single command to run scripts
  - Users won't need python **OR** dependencies installed

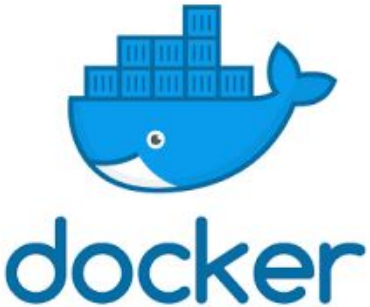

Lessons Learned / Takeaways

- Importance of high quality documentation
  - As important as or more important than actual code
  - Makes reviews faster and documents easier to understand
- Teaching about a subject requires deep understanding
  - Had to know why specific decisions were made at every step
  - Required lots of research on the topics

Lessons Learned / Takeaways

- Importance of finishing personal assignments in a timely manner
  - Several reviewing stages that required individual approval
  - Delayed when someone doesn't respond or approve changes

Acknowledgements

Client: Amirsina Torfi

- PhD student in Computer Science
- Focused on deep learning, neural networks

Scikit-Learn

- Open-source Python Machine Learning library