

AI for Natural Language Processing (NLP), Images, and Forecasting

Course Summary

Description

Today, there is a great need for the introduction of AI into all aspects of software, making the enterprise software smart. The argument one often finds in articles describing the unsatisfactory state of business software is, "If smartphones can do it, why can't enterprise software?"

The course is intended for software architects and engineers. It gives them a practical level of experience, achieved through a combination of about 50% lecture, 50% demo work with student's participation.

Topics

- Text Preparation
- NLP and Deep Learning
- Unsupervised NLP (Optional)

Audience

This course is designed for Software Architects and Developers.

Prerequisites

The prerequisites for this course include:

- Familiarity with any programming language
- Be able to navigate Linux command line
- Basic knowledge of command line Linux editors (VI / nano)

Duration

Five Days

AI for Natural Language Processing (NLP), Images, and Forecasting

Course Outline

I. *Text Preparation*

- A. Bag-of-words
- B. Bag-of-n-Grams
- C. Filtering
- D. Stopwords
- E. Frequency-based
- F. Stemming
- G. Parsing and tokenization
- H. TF-IDF
- I. NLTK
- J. TextBlob
- K. SpaCy for semantic pipeline and named entity recognition
- L. AI overview
 - 1. A brief history of AI
 - 2. Types of AI systems
 - 3. Training machine learning models
 - 4. Applying models for prediction
 - 5. Demos and Labs
- M. AI with TensorFlow and Keras
 - 1. Google democratization of AI with TensorFlow
 - 2. Types of neural network (Perceptron, CNN) and their use
 - 3. Text Processing with TensorFlow
 - 4. Use cases and labs

II. *NLP and Deep Learning*

- A. Word2vec
- B. Learning word embedding
- C. The Skip-gram Model
- D. Building the graph
- E. Training the model
- F. Visualizing the embeddings
- G. Optimizing the implementation
- H. Text classification with TensorFlow
- I. Linear models and SVM
- J. Working with Unicode
- K. Automatic translation (seq2seq)
- L. Text generation with RNN
- M. Named entity extraction with RNNs (sequence modeling)
- N. Bidirectional LSTM with attention
- O. Transformer architecture
- P. Context-aware representations using pretrained language models (ELMo, BERT, ULMFiT) (Transfer learning probably)

- Q. Text processing pipelines

III. *Unsupervised NLP (Optional)*

- A. LDA (Latent Dirichlet Allocation)
- B. Topic modeling with gensim
- C. Applications of topic modeling
 - 1. Time series processing and forecasting elements
 - a. Traditional Time Series forecasting with ARIMA models
 - b. Defining Autocorrelation
 - c. Understanding the Dickey-Fuller Test
 - 2. Forecasting with TensorFlow and Keras
 - a. Google democratization of AI with TensorFlow
 - b. Types of neural network (Perceptron, CNN, LSTM) and their use
 - c. Forecasting with TensorFlow
 - d. Using RNN and LSTM in time series prediction.
 - e. Validation and metrics of Time Series Prediction models
 - f. Use cases and labs
 - 3. Image processing with Deep Neural Networks
 - a. Convolutional Neural Networks (CNN) as a new architecture designed for Images
 - b. Filtering, types of images filters
 - c. Labs showing how filters work
 - d. Image classification, digit recognition, clothing item recognition, etc.
 - e. Summary: what's next?