
Learning text representation

for Social Sciences

Salomé Do

February 2020

Introduction

This short course aims at introducing basic tools used to represent textual data in a linguistically powerful way. Representation is at the very heart of most current computational approaches to linguistics. Finding dense, numeric representations enables to use the transformed textual data in many other learning algorithms, and thus help solving many downstream tasks. For instance, finding numeric, robust representations for two texts enables to compute a distance, a similarity between them. Such representations are known to embed a lot of semantic information and thus are particularly relevant when trying to match semantically similar documents, rather than formally similar documents.

Yet, these representations have historically been designed and used to enhance natural language processing models, generally aimed at automatically answering questions, translating texts, extracting informations such as person's, places, organisations mentions, etc. These tasks are predictive rather than descriptive tasks, and might then be of limited interest for a social scientist. Finding ways to use the power of textual representations in a descriptive and exploratory setting is still a scientific challenge, which we will try to account for.

In the course, we chose to focus on the tools that were developed mainly between 2013 and 2019. These tools generally use deep learning, but the mathematical details of deep learning are not on the scope of this course. We will instead describe the core ideas and differences between the main deep learning architectures, and how they can achieve extracting high-level, and sometimes semantically abstract informations. A large part of the course is also dedicated to the recent exploratory and descriptive methods that can reveal to be meaningful for social scientists. Throughout the course, the most important concepts and algorithms will be put into practice, with a real-life corpus.

Contents

1	Learning to represent words	3
1.1	How do machines learn?	3
1.1.1	Supervised, unsupervised, self-supervised learning	3
1.1.2	Training objective and process	5
1.1.3	Evaluation	5
1.2	Representation learning	5
1.2.1	The experts' way : feature engineering	6
1.2.2	What is a good representation?	6
1.2.3	The advent of neural networks	6
1.3	Some word representations	6
1.3.1	A non-neural representation	6
1.3.2	Word2Vec	6
1.3.3	Other well-known embeddings	6
1.4	Application : Learning embeddings over a corpus	7
2	Exploring embedding spaces	8
2.1	Embeddings interpretability and exploration	8
2.1.1	What is interpretability?	8
2.1.2	Algebraic formulas	8
2.1.3	Using algebraic formulas in social sciences	8
2.2	Exploring social bias through embeddings	8
2.3	Comparing embeddings spaces	8
2.3.1	Comparison metrics	8
2.3.2	Embedding space alignments	8
2.4	Application : Exploring and comparing two corpuses	8
3	From words to sentences	9
3.1	Aggregating sequences	9
3.1.1	Main problem	9
3.1.2	Averaging	9
3.1.3	TF-IDF weights	9
3.2	Recurrent neural networks	9
3.3	Contextual embeddings and sequential representation : examples	9
3.4	Application : Sentiment analysis with Transfer Learning	10