

New Directions in Vector Space Models of Meaning

**Phil Blunsom, Edward Grefenstette
and Karl Moritz Hermann***
University of Oxford
first.last@cs.ox.ac.uk

Georgiana Dinu
Center for Mind/Brain Sciences
University of Trento
georgiana.dinu@unitn.it

1 Abstract

Symbolic approaches have dominated NLP as a means to model syntactic and semantic aspects of natural language. While powerful inferential tools exist for such models, they suffer from an inability to capture correlation between words and to provide a continuous model for word, phrase, and document similarity. Distributed representations are one mechanism to overcome these constraints.

This tutorial will supply NLP researchers with the mathematical and conceptual background to make use of vector-based models of meaning in their own research. We will begin by motivating the need for a transition from symbolic representations to distributed ones. We will briefly cover how collocational (distributional) vectors can be used and manipulated to model word meaning. We will discuss the progress from distributional to distributed representations, and how neural networks allow us to learn word vectors and condition them on metadata such as parallel texts, topic labels, or sentiment labels. Finally, we will present various forms of semantic vector composition, and discuss their relative strengths and weaknesses, and their application to problems such as language modelling, paraphrasing, machine translation and document classification.

This tutorial aims to bring researchers up to speed with recent developments in this fast-moving field. It aims to strike a balance between providing a general introduction to vector-based models of meaning, an analysis of diverging strands of research in the field, and also being a hands-on tutorial to equip NLP researchers with the necessary tools and background knowledge to start working on such models. Attendees should be comfortable with basic probability, linear algebra, and continuous mathematics. No substantial knowledge of machine learning is required.

2 Outline

1. Motivation: Meaning in space
2. Learning distributional models for words
3. Neural language modelling and distributed representations
 - (a) Neural language model fundamentals
 - (b) Recurrent neural language models
 - (c) Conditional neural language models
4. Semantic composition in vector spaces
 - (a) Algebraic and tensor-based composition
 - (b) The role of non-linearities
 - (c) Learning recursive neural models
 - (d) Convolutional maps and composition

3 Instructors

Phil Blunsom is an Associate Professor at the University of Oxford's Department of Computer Science. His research centres on the probabilistic modelling of natural languages, with a particular interest in automating the discovery of structure and meaning in text.

Georgiana Dinu is a postdoctoral researcher at the University of Trento. Her research revolves around distributional semantics with a focus on compositionality within the distributional paradigm.

Edward Grefenstette is a postdoctoral researcher at Oxford's Department of Computer Science. He works on the relation between vector representations of language meaning and structured logical reasoning. His work in this area was recently recognised by a best paper award at *SEM 2013.

Karl Moritz Hermann is a final-year DPhil student at the Department of Computer Science in Oxford. His research studies distributed and compositional semantics, with a particular emphasis on mechanisms to reduce task-specific and monolingual syntactic bias in such representations.

*Instructors listed in alphabetical order.