A Bayesian Model of the Temporal Dynamics of Word Meaning

Lea Frermann
Institute for Language, Cognition and Computation
University of Edinburgh
l.frermann@ed.ac.uk

Abstract

Natural language is a dynamic system, constantly evolving and adapting to the needs of its users and their environment [1]. Statistical models, trained on large-scale collections of historic text data, can produce insights into why and how language changes, and provide meaning representations which can benefit historical linguistic research and applications like historical text mining.

In our work, we investigate how meaning of individual words changes over time. Meaning change is a smooth and gradual process which varies in degree from word to word. Furthermore word meaning changes in various ways: e.g., new meanings (or senses) emerge (e.g., the *pointing device* sense of the word `mouse`), word senses change in prevalence (the *pointing device* sense became increasingly dominant over the past 50 years), or internally (the mouse *device* was characterized by contextual words like {ball, cable} 30 years ago, but by {optical, wireless} today).

We formalize these intuitions in a Bayesian model which learns a representation of an individual target word c's meaning change over time [2]. Our model learns from textual data, covering more than 300 years, represented as time-tagged documents. Each document comprises one mention of c in local context, and our model is based on the assumption that local context is predictive of the word's sense in the document. We adopt a discrete notion of time as contiguous temporal intervals $t \in [1..T]$ of equal length. For each time interval a meaning representation is inferred as (a) a set of senses $s \in [1..S]$ (each as a probability distribution over words w in a vocabulary $p(w|s,t)$), and (b) a probability distribution over senses $p(s|t)$ indicating their prevalence at time t. We explicitly model the local smoothness of meaning change by coupling temporally adjacent model parameters (word- and sense probabilities). In particular we use first-order intrinsic GMRFs [3] which allow us to draw each parameter from a normal distribution centered around the values of its temporal neighbors. The extent of local change is learnt during inference as the precision (inverse variance) of these normal distributions. We use blocked Gibbs sampling for inference. Given a learnt model we can detect different kinds of word meaning change within and across senses by comparing posterior estimates of the time-specific parameters.

Contrary to previous work on modeling word meaning change over time [4], we propose one coherent model for both inferring meaning representations, and detecting their changes. Additionally, by explicitly modeling locally smooth dynamics we infer one global, consistent set of temporal representations, whereas previous work induced time-specific representations independent of each other. We empirically demonstrate the benefit of the explicit model of local coherence, as well as the practical utility of the inferred representations on the task of temporal text classification, where our general model performs highly competitively compared to task-specifically engineered systems.

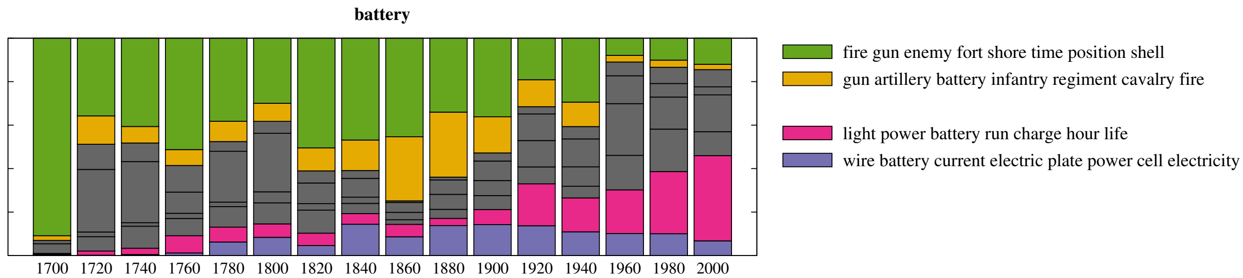

Figure 1: Dynamics of the meaning of the word `battery`. Sense prevalence for 20-year time intervals spanning the period of 1700-2010 is shown as stacked histograms (left); selected senses are highlighted in color (the remaining ones are shown in gray) and illustrated with their most highly associated words (right). The *military* sense of `battery` (green, yellow) decreases in prevalence over time, while its *electrical* sense increases. Two aspects of the *electrical* sense emerge: its chemical aspect (violet; with representative terms like {current, plate, cell}) which emerges from 1780; and its appliance-oriented sense ({light, run, life, charge}), which increases in prevalence from the 1920s.

We also qualitatively show that our model infers informative and interpretable representations. Figure 1 illustrates the diachronic representations inferred by our model for the word `battery`.

References

[1] April M.S. McMahon. *Understanding Language Change*. Cambridge University Press, 1994.

[2] Lea Frermann and Mirella Lapata. A Bayesian Model of Diachronic Meaning Change. *Transactions of the Association for Computational Linguistics (TACL)*, 2015. under submission.

[3] H. Rue and L. Held. *Gaussian Markov Random Fields: Theory and Applications*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. CRC Press, 2005.

[4] Sunny Mitra, Ritwik Mitra, Suman Kalyan Maity, Martin Riedl, Chris Biemann, Pawan Goyal, and Animesh Mukherjee. An automatic approach to identify word sense changes in text media across timescales. *Natural Language Engineering*, FirstView:1–26, 8 2015.