

Unsupervised Dialogue Act Induction using Gaussian Mixtures

Tomáš Brychcín^{1,2} and Pavel Král^{1,2}

¹NTIS – New Technologies for the Information Society,
Faculty of Applied Sciences, University of West Bohemia, Czech Republic

²Department of Computer Science and Engineering,
Faculty of Applied Sciences, University of West Bohemia, Czech Republic

{brychcin, pkral}@kiv.zcu.cz

<http://nlp.kiv.zcu.cz>

Abstract

This paper introduces a new unsupervised approach for dialogue act induction. Given the sequence of dialogue utterances, the task is to assign them the labels representing their function in the dialogue.

Utterances are represented as real-valued vectors encoding their meaning. We model the dialogue as Hidden Markov model with emission probabilities estimated by Gaussian mixtures. We use Gibbs sampling for posterior inference.

We present the results on the standard Switchboard-DAMSL corpus. Our algorithm achieves promising results compared with strong supervised baselines and outperforms other unsupervised algorithms.

1 Introduction

Modeling the discourse structure is the important step toward understanding a dialogue. The description of the discourse structure is still an open issue. However, some low level characteristics have already been clearly identified, e.g. to determine the dialogue acts (DAs) (Jurafsky and Martin, 2009). DA represents the meaning of an utterance in the context of the full dialogue.

Automatic DA recognition is fundamental for many applications, starting with dialogue systems (Allen et al., 2007). The expansion of social media in the last years has led to many other interesting applications, e.g. thread discourse structure prediction (Wang et al., 2011), forum search (Seo et al., 2009), or interpersonal relationship identification (Diehl et al., 2007).

Supervised approaches to DA recognition have been successfully investigated by many authors

(Stolcke et al., 2000; Klüwer et al., 2010; Kalchbrenner and Blunsom, 2013). However, annotating training data is both slow and expensive process. The expenses are increased if we consider different languages and different methods of communication (e.g. telephone conversations, e-mails, chats, forums, Facebook, Twitter, etc.). As the social media and other communication channels grow it has become crucial to investigate unsupervised models. There are, however, only very few related works.

Crook et al. (2009) use Chinese restaurant process and Gibbs sampling to cluster the utterances into flexible number of groups representing DAs in a travel-planning domain. The model lacks structural information (dependencies between DAs) and works only on the surface level (it represents an utterance as a word frequency histogram).

Sequential behavior of DAs is examined in (Ritter et al., 2010), where block Hidden Markov model (HMM) is applied to model conversations on Twitter. Authors incorporate a topic model on the top of HMM to distinguish DAs from topical clusters. They do not directly compare the resulting DAs to gold data. Instead, they measure the prediction ability of the model to estimate the order of tweets in conversation. Joty et al. (2011) extend this work by enriching the emission distribution in HMM to also include the information about speaker and its relative position. A similar approach is investigated by Paul (2012). They use mixed-membership Markov model which includes the functionality of topic models and assigns a latent class to each individual token in the utterance. They evaluate on the thread reconstruction task and on DA induction task, outperforming the method of Ritter et al. (2010).

In this paper, we introduce a new approach to unsupervised DA induction. Similarly to previous works, it is based on HMMs to model the struc-




Figure 1: DA model based on Gaussian mixtures.

tural dependencies between utterances. The main novelty is the use of Multivariate Gaussian distribution for emissions (utterances) in HMM. Our approach allows to represent the utterances as real-valued vectors. It opens up opportunities to design various features encoding properties of each utterance without any modification of the proposed model. We evaluate our model together with several baselines (both with and without supervision) on the standard Switchboard-DAMSL corpus (Jurafsky et al., 1997) and directly compare them with the human annotations.

The rest of the paper is organized as follows. We start with the definition of our model (Sections 2, 3, and 4). We present experimental results in Section 5. We conclude in Section 6 and offer some directions for future work.

2 Proposed Model

Assume we have a set of dialogues \mathcal{D} . Each dialogue $\mathbf{d}_j \in \mathcal{D}$ is a sequence of DA utterances $\mathbf{d}_j = \{d_{j,i}\}_{i=1}^{N_j}$, where N_j denote the length of the sequence \mathbf{d}_j . Let N denote the length of corpora $N = \sum_{\mathbf{d}_j \in \mathcal{D}} N_j$. We model dialogue by HMM with K discrete states representing DAs (see Figure 1). The observation on the states is a feature vector $\mathbf{v}_{j,i} \in \mathbb{R}^M$ representing DA utterance $d_{j,i}$ (feature representation is described in Section 4). HMMs thus define the following joint distribution over observations $\mathbf{v}_{j,i}$ and states $d_{j,i}$:

$$p(\mathcal{D}, \mathbf{V}) = \prod_{\mathbf{d}_j \in \mathcal{D}} \prod_{i=1}^{N_j} p(\mathbf{v}_{j,i} | d_{j,i}) p(d_{j,i} | d_{j,i-1}). \quad (1)$$

Analogously to \mathcal{D} , \mathbf{V} is a set of vector sequences $\mathbf{v}_j = \{\mathbf{v}_{j,i}\}_{i=1}^{N_j}$.

We can represent dependency between consecutive HMM states with a set of K multi-

nomial distributions $\boldsymbol{\theta}$ over K states, such that $P(d_{j,i} | d_{j,i-1}) = \theta_{d_{j,i-1}, d_{j,i}}$. We assume the probabilities $p(\mathbf{v}_{j,i} | d_{j,i})$ have the form of Multivariate Gaussian distribution with the mean $\boldsymbol{\mu}_{d_{j,i}}$ and covariance matrix $\boldsymbol{\Sigma}_{d_{j,i}}$. We place conjugate priors on parameters $\boldsymbol{\mu}_{d_{j,i}}$, $\boldsymbol{\Sigma}_{d_{j,i}}$, and $\boldsymbol{\theta}_{d_{j,i-1}}$: multivariate Gaussian centered at zero for the mean, an inverse-Wishart distribution for the covariance matrix, and symmetric Dirichlet prior for multinomials. We do not place any assumption on the length of the dialogue N_j . The full generative process can thus be summarized as follows:

1. For each DA $1 \leq k \leq K$ draw:
 - (a) covariance matrix $\boldsymbol{\Sigma}_k \sim \mathcal{W}^{-1}(\boldsymbol{\Psi}, \nu)$,
 - (b) mean vector $\boldsymbol{\mu}_k \sim \mathcal{N}(\boldsymbol{\mu}, \frac{1}{\kappa} \boldsymbol{\Sigma}_k)$,
 - (c) distribution over following DAs $\boldsymbol{\theta}_k \sim \text{Dir}(\boldsymbol{\alpha})$.
2. For each dialogue $\mathbf{d}_j \in \mathcal{D}$ and for each position $1 \leq i \leq N_j$ draw:
 - (a) DA $d_{j,i} \sim \text{Discrete}(\boldsymbol{\theta}_{d_{j,i-1}})$,
 - (b) feature vector $\mathbf{v}_{j,i} \sim \mathcal{N}(\boldsymbol{\mu}_{d_{j,i}}, \boldsymbol{\Sigma}_{d_{j,i}})$.

Note that κ and ν represents the strength of the prior for the mean and the covariance, respectively. $\boldsymbol{\Psi}$ is the scale matrix of inverse-Wishart distribution.

3 Posterior Inference

Our goal is to estimate the parameters of the model in a way that maximizes the joint probability in Equation 1. We apply Gibbs sampling and gradually resample DA assignments to individual DA utterances. For doing so, we need to determine the posterior predictive distribution.

The predictive distribution of Dirichlet-multinomial has the form of additive smoothing that is well known in the context of language modeling. The hyper-parameter of Dirichlet prior determine how much is the predictive distribution smoothed. Note that we use symmetrical Dirichlet prior so $\boldsymbol{\alpha}$ in the following equations is a scalar. The predictive distribution for transitions in HMM can be expressed as

$$P(d_{j,i} | d_{j,i-1}, \mathbf{d}_{\setminus j,i}) = \frac{n_{\setminus j,i}^{(d_{j,i} | d_{j,i-1})} + \alpha}{n_{\setminus j,i}^{(\bullet | d_{j,i-1})} + K\alpha}, \quad (2)$$

where $n_{\setminus j,i}^{(d_{j,i}|d_{j,i-1})}$ is the number of times DA $d_{j,i}$ followed DA $d_{j,i-1}$. The notation $\setminus j, i$ means to exclude the position i in the j -th dialogue. The symbol \bullet represents any DA so that $n_{\setminus j,i}^{(\bullet|d_{j,i-1})} = \sum_{1 \leq k \leq K} n_{\setminus j,i}^{(k|d_{j,i-1})}$.

The predictive distribution of Normal-inverse-Wishart distribution has the form of multivariate student t -distribution $t_{\nu'}(\mathbf{v}|\boldsymbol{\mu}', \boldsymbol{\Sigma}')$ with ν' degrees of freedom, mean vector $\boldsymbol{\mu}'$, and covariance matrix $\boldsymbol{\Sigma}'$. According to (Murphy, 2012) the parameters for posterior predictive distribution can be estimated as

$$\begin{aligned} \kappa_k &= \kappa + n^{(k)}, & \nu_k &= \nu + n^{(k)}, \\ \boldsymbol{\Psi}_k &= \boldsymbol{\Psi} + \mathbf{S}_k + \frac{\kappa n^{(k)}}{\kappa_k} (\bar{\mathbf{V}}^{(k)} - \boldsymbol{\mu})(\bar{\mathbf{V}}^{(k)} - \boldsymbol{\mu})^\top, \\ \boldsymbol{\mu}_k &= \frac{\kappa \boldsymbol{\mu} + n^{(k)} \bar{\mathbf{V}}^{(k)}}{\kappa_k}, & \boldsymbol{\Sigma}_k &= \frac{\boldsymbol{\Psi}_k}{\nu_k - K + 1}, \end{aligned} \quad (3)$$

where $n^{(k)}$ is the number of times DA k occurred in the data, $\bar{\mathbf{V}}^{(k)}$ is the mean of vectors associated with DA k , and $\mathbf{S}_k = \sum_{d_{j,i}=k} (\mathbf{v}_{j,i} - \bar{\mathbf{V}}^{(k)})(\mathbf{v}_{j,i} - \bar{\mathbf{V}}^{(k)})^\top$ is scaled form of the covariance of these vectors. Note that κ , ν , $\boldsymbol{\mu}$, and $\boldsymbol{\Psi}$ are hyper-parameters which need to be set in advance.

Now we can construct the final posterior predictive distribution used for sampling DA assignments:

$$\begin{aligned} P(d_{j,i} = k | \mathbf{D}_{\setminus j,i}, \mathbf{V}_{\setminus j,i}) &\propto \\ &P(d_{j,i} | d_{j,i-1}, \mathbf{d}_{\setminus j,i}) \\ &\times P(d_{j,i+1} | d_{j,i}, \mathbf{d}_{\setminus j,i+1}) \\ &\times t_{\nu_k - K + 1}(\mathbf{v}_{j,i} | \boldsymbol{\mu}_k, \frac{\kappa_k + 1}{\kappa_k} \boldsymbol{\Sigma}_k). \end{aligned} \quad (4)$$

The product of the first two parts in the equation expresses the score proportional to the probability of DA at position i in the j -th dialogue given the surrounding HMM states. The third part expresses the probability of DA assignment given the current feature vector $\mathbf{v}_{j,i}$ and all other DA assignments.

We also present the simplified version of the model that is in fact the standard Gaussian mixture model (GMM). This model does not capture the dependencies between surrounding DAs in the dialogue. Posterior predictive distribution is as follows:

$$\begin{aligned} P(d_{j,i} = k | \mathbf{D}_{\setminus j,i}, \mathbf{V}_{\setminus j,i}) &\propto \frac{n_{\setminus j,i}^{(k)} + \alpha}{N - 1 + K\alpha} \\ &\times t_{\nu_k - K + 1}(\mathbf{v}_{j,i} | \boldsymbol{\mu}_k, \frac{\kappa_k + 1}{\kappa_k} \boldsymbol{\Sigma}_k). \end{aligned} \quad (5)$$

In Section 5 we provide comparison of both models to see the strengths of using DA context.

4 DA Feature Vector

The real-valued vectors $\mathbf{v}_{j,i}$ are expected to represent the meaning of $d_{j,i}$. We use semantic composition approach. It is based on *Frege's principle of compositionality* (Pelletier, 1994), which states that the meaning of a complex expression is determined as a composition of its parts, i.e. words.

We use linear combination of word vectors, where the weights are represented by the inverse-document-frequency (IDF) values of words. We use Global Vectors (GloVe) (Pennington et al., 2014) for word vector representation. We use pre-trained word vectors¹ on 6B tokens from Wikipedia 2014 and Gigaword 5. Brychcín and Svoboda (2016) showed that this approach leads to very good representation of short sentences.

For supervised approaches we also use bag-of-words (BoW) representation of an utterance, i.e. separate binary feature representing the occurrence of a word in the utterance.

5 Experimental Results and Discussion

We use Switchboard-DAMSL corpus (Jurafsky et al., 1997) to evaluate the proposed methods. The corpus contains transcriptions of telephone conversations between multiple speakers that do not know each other and are given a topic for discussion. We adopt the same set of 42 DA labels and the same train/test data split as suggested in (Stolcke et al., 2000)².

In our experiments we set $\kappa = 0$, $\boldsymbol{\mu} = \mathbf{0}$, $\nu = K$, $\boldsymbol{\Psi} = \mathbf{1}$, and $\alpha = 50/K$. These parameters are recommended by (Griffiths and Steyvers, 2004; Murphy, 2012) and we also confirm them empirically. We always perform 1000 iterations of Gibbs sampling. The number of clusters (mixture size) is $K = 42$. The dimension of GloVe vectors ranges between $M = 50$ and $M = 300$.

DA induction task is in fact the clustering problem. We cluster DA utterances and we assign the same label to utterances within one cluster. Standard metrics for evaluating quality of clusters are *purity* (PU), *collocation* (CO), and their harmonic

¹Available at <http://nlp.stanford.edu/projects/glove/>.

²1115 dialogues (196,258 utterances) are used for training while 19 dialogues (4186 utterances) for testing. More information about the data split can be found at <http://web.stanford.edu/~jurafsky/ws97>.

	Model	AC	PU	CO	F1	HO	CM	V1
Extreme	Random labels	2.6%	31.5%	4.9%	8.5%	6.8%	4.1%	5.1%
	Distinct labels	0.0%	100.0%	0.9%	1.8%	100.0%	26.9%	42.4%
	Majority label	31.5%	31.5%	100.0%	47.9%	0.0%	100.0%	0.0%
Supervised	ME GloVe ($M = 50$)	63.2%	63.3%	77.8%	69.8%	41.0%	57.3%	47.8%
	ME GloVe ($M = 100$)	64.1%	64.4%	76.9%	70.1%	43.3%	57.3%	49.3%
	ME GloVe ($M = 200$)	64.8%	65.1%	77.2%	70.6%	43.5%	58.1%	49.7%
	ME GloVe ($M = 300$)	65.6%	65.8%	76.0%	70.6%	45.0%	57.7%	50.5%
	ME BoW	70.4%	70.7%	76.3%	73.4%	51.0%	62.9%	56.3%
	ME BoW + GloVe ($M = 300$)	71.5%	72.0%	76.0%	74.0%	53.2%	62.9%	57.7%
	ctx ME BoW + GloVe ($M = 300$)	72.9%	73.0%	76.1%	74.5%	53.9%	64.1%	58.6%
Unsupervised	BHMM (Ritter et al., 2010)	/	60.3%	31.2%	41.1%	43.1%	29.1%	34.7%
	M4 (Paul, 2012)	/	44.4%	45.9%	45.1%	19.4%	16.9%	18.0%
	K-means GloVe ($M = 50$)	/	57.1%	25.9%	35.6%	39.9%	27.5%	32.6%
	K-means GloVe ($M = 100$)	/	56.7%	29.5%	38.8%	39.9%	28.9%	33.5%
	K-means GloVe ($M = 200$)	/	56.9%	32.4%	41.3%	39.7%	31.2%	35.0%
	K-means GloVe ($M = 300$)	/	57.4%	31.2%	40.4%	40.2%	30.3%	34.6%
	GMM GloVe ($M = 50$)	/	54.4%	51.8%	53.1%	34.0%	37.7%	35.8%
	GMM GloVe ($M = 100$)	/	53.8%	58.1%	55.9%	33.7%	40.0%	36.5%
	GMM GloVe ($M = 200$)	/	52.1%	76.9%	62.1%	31.3%	43.6%	36.4%
	GMM GloVe ($M = 300$)	/	52.7%	79.8%	63.5%	30.1%	45.2%	36.1%
	ctx GMM GloVe ($M = 50$)	/	55.1%	60.0%	57.5%	36.4%	42.4%	39.1%
	ctx GMM GloVe ($M = 100$)	/	53.8%	81.7%	64.9%	32.3%	51.7%	39.8%
	ctx GMM GloVe ($M = 200$)	/	54.7%	81.4%	65.5%	32.1%	51.9%	39.7%
	ctx GMM GloVe ($M = 300$)	/	55.2%	81.0%	65.7%	34.4%	51.4%	41.2%

Table 1: Accuracy (AC), purity (PU), collocation (CO), f-measure (F1), homogeneity (HO), completeness (CM), and v-measure (V1) for proposed models expressed in percents.

mean (F1). In the last years, *v-measure* (V1) have also become popular. This entropy-based measure is defined as harmonic mean between *homogeneity* (HO – the precision analogue) and *completeness* (CM – the recall analogue). Rosenberg and Hirschberg (2007) presents definition and comparison of all these metrics. Note the same evaluation procedure is often used for different clustering tasks, e.g., unsupervised part-of-speech induction (Christodoulopoulos et al., 2010) or unsupervised semantic role labeling (Woodsend and Lapata, 2015).

Table 1 presents the results of our experiments. We compare both supervised and unsupervised approaches. Models incorporating the information about surrounding DAs (context) are denoted by prefix *ctx*. We show the results of three unsupervised approaches: K-means clustering, GMM without context (Eq. 5), and context-dependent GMM (Eq. 4). We use Maximum Entropy (ME) classifier (Berger et al., 1996) for the supervised approach. For the context-dependent version we perform two-round classification: firstly, without

the context information and secondly, incorporating the output from the previous round.

In addition, Table 1 provides results for the three extreme cases: *random label*, *majority label*, and *distinct label* for each utterance (a single utterance per cluster). Note the last mentioned achieved v-measure of 42.4%. In this case, however, completeness approaches 0% with the rising size of the test data (so v-measure does too). So this number cannot be taken into account.

To the best of our knowledge, the best performing supervised system on Switchboard-DAMSL corpus is presented in (Kalchbrenner and Blunsom, 2013) and achieves accuracy of 73.9%. Our best supervised baseline is approximately 1% worse. In all experiments the context information proved to be very useful. The best result among unsupervised models is achieved with 300-dimensional GloVe (F1 score 65.7% and v-measure 41.2%). We outperform both the block HMM (BHMM) (Ritter et al., 2010) achieving F1 score 41.1% and v-measure 34.7% and mixed-membership HMM (M4) (Paul, 2012) achieving

F1 score 45.1% and v-measure 18.0%³. If we compare our method with the supervised version (F1 score 74.5% and v-measure 58.6%) we can state that HMM with GMMs is very promising direction for the unsupervised DA induction task.

6 Conclusion and Future Work

We introduced HMM based model for unsupervised DA induction. We represent each utterance as a real-valued vector encoding the meaning. Our model predicts these vectors in the context of DA utterances. We compared our model with several strong baselines and showed its strengths. Our Java implementation is available for research purposes at <https://github.com/brychcin/unsup-dial-act-induction>.

As the main direction for future work, we plan to experiment with more languages and more corpora. Also, more thorough study of feature vector representation should be done.

We plan to investigate the learning process much more deeply. It was beyond the scope of this paper to evaluate the time expenses of the algorithm. Moreover, there are several possibilities how to speed up the process of parameter estimation, e.g. by Cholesky decomposition of the covariance matrix as described in (Das et al., 2015). In our current implementation the number of DAs is set in advance. It could be very interesting to use non-parametric version of GMM, i.e. to change the sampling scheme to estimate the number of DAs by Chinese restaurant process.

Acknowledgments

This publication was supported by the project LO1506 of the Czech Ministry of Education, Youth and Sports under the program NPU I. Computational resources were provided by the CESNET LM2015042 and the CERIT Scientific Cloud LM2015085, provided under the programme "Projects of Large Research, Development, and Innovations Infrastructures". Lastly, we would like to thank the anonymous reviewers for their insightful feedback.

³Both implementations are available at <http://cmci.colorado.edu/~mpaul/downloads/mm.php>. We use recommended settings. Note the comparison with M4 is not completely fair, because it does not directly assign DAs to utterances (instead, it assigns DAs to each token). We always took the most frequent token DA in utterance as final DA.

References

- James Allen, Nathanael Chambers, George Ferguson, Lucian Galescu, Hyuckchul Jung, Mary Swift, and William Taysom. 2007. Plow: A collaborative task learning agent. In *Proceedings of the 22Nd National Conference on Artificial Intelligence - Volume 2*, AAAI'07, pages 1514–1519. AAAI Press.
- Adam L. Berger, Vincent J. D. Pietra, and Stephen A. D. Pietra. 1996. A maximum entropy approach to natural language processing. *Computational Linguistics*, 22:39–71, March.
- Tomáš Brychcín and Lukáš Svoboda. 2016. UWB at SemEval-2016 Task 1: Semantic Textual Similarity using Lexical, Syntactic, and Semantic Information. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 588–594, San Diego, California, June. Association for Computational Linguistics.
- Christos Christodoulopoulos, Sharon Goldwater, and Mark Steedman. 2010. Two decades of unsupervised POS induction: How far have we come? In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 575–584, Cambridge, MA, October. Association for Computational Linguistics.
- Nigel Crook, Ramon Granell, and Stephen Pulman. 2009. Unsupervised classification of dialogue acts using a Dirichlet process mixture model. In *Proceedings of the SIGDIAL 2009 Conference*, pages 341–348, London, UK, September. Association for Computational Linguistics.
- Rajarshi Das, Manzil Zaheer, and Chris Dyer. 2015. Gaussian lda for topic models with word embeddings. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 795–804, Beijing, China, July. Association for Computational Linguistics.
- Christopher P. Diehl, Galileo Namata, and Lise Getoor. 2007. Relationship identification for social network discovery. In *Proceedings of the 22nd National Conference on Artificial Intelligence*, AAAI'07, pages 546–552. AAAI Press.
- Thomas L. Griffiths and Mark Steyvers. 2004. Finding scientific topics. *Proceedings of the National Academy of Sciences of the United States of America*, 101(Suppl 1):5228–5235, April.
- Shafiq Joty, Giuseppe Carenini, and Chin-Yew Lin. 2011. Unsupervised modeling of dialog acts in asynchronous conversations. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, IJCAI'11, pages 1807–1813. AAAI Press.
- Daniel Jurafsky and James H. Martin. 2009. *Speech and Language Processing (2Nd Edition)*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA.

- Daniel Jurafsky, Elizabeth Shriberg, and Debra Bisca. 1997. Switchboard SWBD-DAMSL Shallow-Discourse-Function Annotation (Coders Manual, Draft 13). Technical Report 97-01, University of Colorado, Institute of Cognitive Science.
- Nal Kalchbrenner and Phil Blunsom. 2013. Recurrent convolutional neural networks for discourse compositionality. In *Proceedings of the Workshop on Continuous Vector Space Models and their Compositionality*, pages 119–126, Sofia, Bulgaria, August. Association for Computational Linguistics.
- Tina Klüwer, Hans Uszkoreit, and Feiyu Xu. 2010. Using syntactic and semantic based relations for dialogue act recognition. In *Coling 2010: Posters*, pages 570–578, Beijing, China, August. Coling 2010 Organizing Committee.
- Kevin P. Murphy. 2012. *Machine Learning: A Probabilistic Perspective*. The MIT Press.
- Michael J. Paul. 2012. Mixed membership markov models for unsupervised conversation modeling. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 94–104, Jeju Island, Korea, July. Association for Computational Linguistics.
- Francis Jeffrey Pelletier. 1994. The principle of semantic compositionality. *Topoi*, 13(1):11–24.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar, October. Association for Computational Linguistics.
- Alan Ritter, Colin Cherry, and Bill Dolan. 2010. Unsupervised modeling of twitter conversations. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 172–180, Los Angeles, California, June. Association for Computational Linguistics.
- Andrew Rosenberg and Julia Hirschberg. 2007. V-measure: A conditional entropy-based external cluster evaluation measure. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 410–420, Prague, Czech Republic, June. Association for Computational Linguistics.
- Jangwon Seo, W. Bruce Croft, and David A. Smith. 2009. Online community search using thread structure. In *Proceedings of the 18th ACM Conference on Information and Knowledge Management, CIKM '09*, pages 1907–1910, New York, NY, USA. ACM.
- Andreas Stolcke, Klaus Ries, Noah Coccaro, Elizabeth Shriberg, Rebecca Bates, Daniel Jurafsky, Paul Taylor, Rachel Martin, Carol Van Ess-Dykema, and Marie Meteer. 2000. Dialog act modeling for automatic tagging and recognition of conversational speech. In *Computational Linguistics*, volume 26, pages 339–373.
- Li Wang, Marco Lui, Su Nam Kim, Joakim Nivre, and Timothy Baldwin. 2011. Predicting thread discourse structure over technical web forums. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 13–25, Edinburgh, Scotland, UK., July. Association for Computational Linguistics.
- Kristian Woodsend and Mirella Lapata. 2015. Distributed representations for unsupervised semantic role labeling. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2482–2491, Lisbon, Portugal, September. Association for Computational Linguistics.