



AFRL-RI-RS-TR-2018-129

**SCALABLE PROSODIC, ANOMALY AND RELATIONAL
KNOWLEDGE EXPLORATION OF LANGUAGE WITH ENHANCED
ROBUSTNESS (SPARKLER)**

RESEARCH FOUNDATION OF THE CITY UNIVERSITY OF NEW
YORK (CUNY)

MAY 2018

FINAL TECHNICAL REPORT

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

STINFO COPY

**AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE**

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with SAF/AQR memorandum dated 10 Dec 08 and AFRL/CA policy clarification memorandum dated 16 Jan 09. This report is available to the general public, including foreign nations. Copies may be obtained from the Defense Technical Information Center (DTIC) (<http://www.dtic.mil>).

AFRL-RI-RS-TR-2018-129 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE CHIEF ENGINEER:

/ S /
PETER ROCCI
Work Unit Manager

/ S /
JON S. JONES
Technical Advisor, Information Intelligence
Systems and Analysis Division
Information Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE**Form Approved
OMB No. 0704-0188**

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) MAY 2018			2. REPORT TYPE FINAL TECHNICAL REPORT		3. DATES COVERED (From - To) NOV 2012 – MAY 2017	
4. TITLE AND SUBTITLE SCALABLE PROSODIC, ANOMALY AND RELATIONAL KNOWLEDGE EXPLORATION OF LANGUAGE WITH ENHANCED ROBUSTNESS (SPARKLER)					5a. CONTRACT NUMBER FA9850-13-2-0041	
					5b. GRANT NUMBER N/A	
					5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Heng Ji, Andrew Rosenberg, Yang Liu					5d. PROJECT NUMBER DEFT	
					5e. TASK NUMBER 12	
					5f. WORK UNIT NUMBER 05	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Research Foundation Of The City University Of New York (CUNY) Queens College 65-30 Kissena Blvd Flushing, NY 11367-1575					8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory/RIED 525 Brooks Road Rome NY 13441-4505					10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/RI	
					11. SPONSOR/MONITOR'S REPORT NUMBER AFRL-RI-RS-TR-2018-129	
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited. This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with SAF/AQR memorandum dated 10 Dec 08 and AFRL/CA policy clarification memorandum dated 16 Jan 09						
13. SUPPLEMENTARY NOTES						
14. ABSTRACT We have developed trilingual (English, Chinese and Spanish) open-domain knowledge extraction techniques to extract, populate, and analyze unstructured data from heterogeneous sources in to acknowledge base, to relieve the reliance on manual knowledge base construction or in-person communication for knowledge management. This new frame work is able to discover schemes and extract facts from any input data in any domain, without any annotated training data, by incorporating distributional semantics and symbolic semantics. The resulting systems have won top performance at several NIST international research evaluations and been selected for DARPA and ARL demos and transitions. 4 PhD students have been supported, and 54 papers have been published at top conferences and journals.						
15. SUBJECT TERMS co-reference resolution, event extraction, relationship extraction						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON	
a. REPORT	b. ABSTRACT	c. THIS PAGE			Peter Rocci	
U	U	U	UU	28		

Table of Contents

1.0 Summary	1
2.0 Introduction	1
3.0 Methods, Assumptions, and Procedures	2
3.1 Relational Analysis – Overview	2
3.2 Relational Analysis - Liberal Extraction	3
3.3 Relational Analysis - Joint Extraction	3
3.4 Relational Analysis - Entity Discovery and Linking	4
3.5 Relational Analysis - Entity Morph Decoding	5
3.6 Relational Analysis - Event Extraction	5
3.7 Relational Analysis - Event Coreference Resolution	6
3.8 Relational Analysis - Open Relation Extraction and Grounding	6
3.9 Relational Analysis - Slot Filling and Slot Filling Validation	7
3.10 Relational Analysis - Event-Event Relation Extraction	7
3.11 Relational Analysis - Temporal Knowledge Base Population	8
3.12 Relational Analysis - Power Relation Classification	8
3.13 Relational Analysis: User Attribute Classification	9
3.14 Prosodic Analysis for Relational Analysis – Name Tagging	10
3.15 Relational Analysis - Chinese Event Extraction:	10
3.16 Relational Analysis - Influential User and Post Identification	11
3.17 Recognition of Emotion, Conflict and Laughter	12
3.18 Speech Segmentation	12
3.19 Anomaly Analysis - Emotion Recognition	12
3.20 Anomaly Analysis - Disfluency Detection	13
3.21 Language Processing Algorithms – Summarization	13
3.22 Language processing Algorithms - Text Normalization and Impact on other NLP Tasks:	13
4.0 Results and Discussion	15
4.1 NIST TAC-KBP Evaluation Results	16
4.2 PhDs Supported by DEFT	16
4.3 Post-docs supported by DEFT	17
4.4 Principal Investigator Honors	17
5.0 Conclusions	17
6.0 Publications	18

1.0 Summary

We have developed tri-lingual (English, Chinese and Spanish) open-domain knowledge extraction techniques to extract, populate, and analyze unstructured data from heterogeneous sources into a knowledge base, to relieve the reliance on manual knowledge base construction or in-person communication for knowledge management. This new framework is able to discover schemas and extract facts from any input data in any domain, without any annotated training data, by incorporating distributional semantics and symbolic semantics. The resulting systems have won top performance at several NIST international research evaluations and been selected for DARPA and ARL demos and transitions. 4 PhD students have been supported, and 54 papers have been published at top conferences and journals.

2.0 Introduction

The goal of this project is to extract structured facts from a wide spectrum of heterogeneous unstructured texts, and link them to knowledge bases. Traditional Information Extraction (IE) techniques are limited to a certain source X (X = a particular language, domain, limited number of pre-defined fact types, single data modality, ...). When moving from X to a new source Y , we need to start from scratch again by annotating a substantial amount of training data and developing Y -specific extraction capabilities. In this project we developed a new Universal IE paradigm to combine the merits of traditional IE (high quality and fine granularity) and Open IE (high scalability). This framework is able to discover schemas and extract facts from any input data in any domain, without any annotated training data, by incorporating

distributional semantics and symbolic semantics. We have conducted research in the following IE and Knowledge Base Population tasks:

- Entity Discovery and Linking
- Event Extraction and Event Coreference Resolution
- Relation Extraction and Slot Filling
- Event-Event Relation Extraction
- Native Language Detection
- Power Relation Detection
- Hierarchy Detection in Online Communications
- Influential User and Post Identification

3.0 Methods, Assumptions, and Procedures

3.1 Relational Analysis – Overview

We proposed and developed a "Liberal" Information Extraction (IE) paradigm, which can extract and translation information from multiple languages, genres and domains with low cost. The principle is to embrace symbolic semantics and distributional semantics, and perform IE in a bottom-up discovery instead of top-down classification fashion. This paradigm takes an arbitrary corpus as input, and simultaneously discover a customizable schema and extract knowledge elements in a "Cold-start" way, without human-defined schema or manual annotations. Thus it can be easily adapted to a new domain, genre or language. Using this new framework we are able to extend the extraction capability from 33 types of events to 3,000 types, 5 types of entities to 4,000 types, 41 types of relations to 2,000 types, 1-3 languages to 3,000 languages. The technical PI Ji has been coordinating the NIST TAC-KBP EDL track through the program [24, 35], and led the TinkerBell team

(Columbia, Cornell, JHU, RPI, Stanford, UIUC, UPenn) to build the first end-to-end tri-lingual cold-start++ KBP system [5].

3.2 Relational Analysis - Liberal Extraction

We developed a brand new Liberal Information Extraction paradigm which aims to discover schema and extract facts simultaneously from any input corpus in any domain. We incorporate symbolic (e.g., Abstract Meaning Representation) and distributional semantics to detect and represent event structures and adopt a joint typing framework to simultaneously extract event types and argument roles and discover an event schema. Experiments on general and specific domains demonstrate that this framework can construct high-quality schemas with many event and argument role types, covering a high proportion of event types and argument roles in manually defined schemas. We show that extraction performance using discovered schemas is comparable to supervised models trained from a large amount of data labeled according to predefined event types. The extraction quality of new event types is also promising. We applied this framework to build quality fine-grained entity extraction systems [1] and event extraction [22] systems with low cost.

3.3 Relational Analysis - Joint Extraction

We developed a single joint model to extract entity mentions, relations and event mentions simultaneously [45, 46, 53]. Relational Analysis approaches are typically organized as a pipeline architecture of processing stages (e.g. from speech recognition, to anomaly analysis, to name tagging, to coreference resolution, to relation extraction, to event trigger extraction, to event argument extraction, to cross-document entity linking and slot filling, etc.). Each of these stages has been studied separately and quite intensively over the past decade. Annotated corpora have been prepared for each stage, a wide range of models and machine learning methods have been applied, and separate official evaluations have been organized. There has clearly been a great deal of progress on some of these components. However, the

output of each stage is chosen locally and passed to the next step, and there is no feedback from later stages to earlier ones. Although this makes the systems comparatively easy to assemble, it comes at a high price: errors accumulate as information progresses through the pipeline, and an error once made cannot be corrected. We have developed a joint framework based on structured perceptron which extracts different types of facts (e.g. event triggers and arguments) together so that the local predictions can be mutually improved. In addition, we propose to incorporate global features which explicitly capture the dependencies of multiple facts. Experimental results on event extraction show that our joint approach with local features outperforms the sequential baseline, and adding global features further improves the performance significantly. Our approach advances state-of-the-art sentence-level event extraction, (7.2% higher trigger labeling F-measure and 16.2% higher argument labeling F-measure) and even outperforms previous argument labeling methods using external knowledge from other sentences and documents.

3.4 Relational Analysis - Entity Discovery and Linking

We demonstrate the effectiveness of Abstract Meaning Representation (AMR) to select high quality sets of entity “collaborators” to feed a simple similarity measure (Jaccard) to link entity mentions. Experimental results show that AMR captures contextual properties discriminative enough to make linking decisions, without the need for Entity Linking (EL) training data, and that system with AMR parsing output outperforms hand labeled traditional semantic roles as context representation for EL [34]. We further developed a deep semantic relatedness model (DSRM) based on deep neural networks (DNN) and semantic knowledge graphs (KGs) to measure semantic relatedness between entities. We learn latent semantic entity representations through knowledge graph embeddings by encoding heterogeneous types of knowledge from KGs into DNN. The non-linear transformations of hierarchical hidden layers of DNN and the rich and clean knowledge from KGs allow deep semantic understanding of entities. We then extend the unsupervised graph regularization

model to model topical coherence and perform collective entity disambiguation[44]. We extended this semi-supervised framework to EL problem for scarce KBs [2]. We developed an unsupervised algorithm named Quantified Collective Validation that avoids excessive linguistic analysis on the source documents and fully leverages the knowledge base structure for the entity linking task [30], which can be applied to any domain [41] and any language [3, 6].

3.5 Relational Analysis - Entity Morph Decoding

People create morphs, a special type of fake alternative names, to achieve certain communication goals such as expressing strong sentiment or evading censors. We developed an end-to-end context-aware entity morph decoding system that can automatically identify, disambiguate, verify morph mentions based on specific contexts, and resolve them to target entities. Our approach is based on an absolute “cold-start” - it does not require any candidate morph or target entity lists as input, nor any manually constructed morph-target pairs for training [32, 52].

3.6 Relational Analysis - Event Extraction

We apply Zero-shot Transfer Learning to improve the quality of Liberal Event Extraction. We leverage existing human constructed schema and manual annotations for a small set of pre-defined event types, and transfer the knowledge of these existing types to the extraction of new event types. We then further improved the performance of textual event extraction by incorporating external visual information. Experiments on both the existing ACE event schema and new event types demonstrate the effectiveness of our transferable neural architecture. Without using any manual annotations on new event types, our framework achieves comparable performance as state-of-the-art supervised model trained from more than three thousand sentences. We developed a hybrid neural network to capture both sequence and chunk information from specific contexts, and use them to train an event detector for multiple languages without any manually encoded features. Experiments

show that our approach can achieve robust, efficient and accurate results for multiple languages (English, Chinese and Spanish) [15]. Existing event extractors for social media suffer from two major problems: lack of context and informal nature. Instead of conducting event extraction solely on each social media message, we incorporate cross-genre knowledge to boost the event extractor performance. Experiment results demonstrate that without any additional annotations, our proposed approach is able to provide 15% absolute F-score improvement over the state-of-the-art [12, 25].

3.7 Relational Analysis - Event Coreference Resolution

We developed a novel cross-document coreference resolution framework using cross-media features [33]. We model event coreference resolution as a spectral graph clustering problem. We then use event coreference resolution as feedback to enhance event nugget detection and realis detection, based on the hypothesis that if two event nuggets are coreferential, they should have consistent event types and realis labels.

3.8 Relational Analysis - Open Relation Extraction and Grounding

Previous open Relation Extraction (open RE) approaches mainly rely on linguistic patterns and constraints to extract important relational triples from large-scale corpora. However, they lack of abilities to cover diverse relation expressions or measure the relative importance of candidate triples within a sentence. It is also challenging to name the relation type of a relational triple merely based on context words, which could limit the usefulness of open RE in downstream applications. We propose a novel importance based open RE approach by exploiting the global structure of a dependency tree to extract salient triples. We design an unsupervised method to name relation types by grounding relational triples to a large-scale Knowledge Base (KB) schema, leveraging KB triples and weighted context words associated

with relational triples. Experiments on the English Slot Filling 2013 dataset demonstrate that our approach achieves 8.1% higher F-score over state-of-the-art open RE methods [4].

3.9 Relational Analysis - Slot Filling and Slot Filling Validation

We developed two new approaches based on attentive Deep Neural Networks for Slot Filling: (1). Take a regularized dependency graph as input to DNN, to compress the wide contexts between query and candidate filler; (2). Incorporate two types of attention mechanisms: local attention learned from query and candidate filler, and global attention learned from external knowledge bases, to guide the model to better select indicative contexts to determine slot type. Experiments show that this framework outperforms state-of-the-art by 16% absolute F-score gain [8]. We further captured more inter-dependency relations among different slot types, and incorporated them into the truth finding framework for slot filling validation. We also conducted large-scale graph pattern mining and matching to jointly extract inter-dependent slots. Experiments on the case study of Slot Filling Validation demonstrate that our approach can find truths accurately (9.4% higher F-score than supervised methods) and efficiently (finding 90% truths with only one half the cost of a baseline without credibility estimation) [42].

3.10 Relational Analysis - Event-Event Relation Extraction

We developed an overall framework including task definition, event-event relation taxonomy, data annotation and evaluation metrics. Our schema includes five main types of relations: Inheritance, Expansion, Contingency, Comparison and Temporality, along with 21 subtypes [14, 23]. This corpus has been shared with the entire

DEFT and KBP communities. We also designed a novel scenario projection approach method to predict relations between pairwise events, based on the hypothesis that events often co-occur in the same scenario tend to have similar relations.

3.11 Relational Analysis - Temporal Knowledge Base Population

We tackled temporal knowledge base population challenges from three perspectives: (1) Temporal information representation: we developed a new (4-tuple) representation framework for the task. (2) Annotation acquisition: The lack of substantial labeled training data for supervised learning is a limiting factor in the design of temporal slot filling systems. Our work examines the use of multi-class logistic regression methods to improve the labeling quality of training data obtained by distant supervision. (3) Temporal information classification: Another key challenge lies in capturing relations between salient text elements separated by a long context. We develop two approaches for temporal classification and combine them through cross- document aggregation: a flat approach that uses lexical context and shallow dependency features and a structured approach that captures long syntactic contexts by using a dependency path kernel tailored for this task. Experimental results demonstrated that our annotation enhancement approach dramatically increased the speed of the training procedure (by almost 100 times), and that the flat and structured classification approaches were complementary, together yielding a state-of-the-art temporal knowledge base population system [43].

3.12 Relational Analysis - Power Relation Classification

Understanding power relationships provides important context for better communication analysis. We investigate two different communities, business communication (ENRON corpus) and online forum (Reddit - currently collecting data), to develop robust prediction algorithms. For our experiments we took linguistically inspired approach. We use various word list and term vector approach to make predictions. The word lists

comprise well known LIWC and HierarchyPhrases developed by Eric Gilbert as well as new word lists developed by our group. Preliminary results can be found below. These experiments were conducted on ENRON corpus using ten-fold cross-validation. The task is to identify power hierarchy inside ENRON corporation given a single email. This is a three way classification task with labels UP_HIERARCHY (a person writes to somebody with higher position), DOWN_HIERARCHY (lower position), SAME_HIERARCHY (the same level). The baseline defined to be equal to chance (33.3%), our best performance is 56.33% accuracy – better than the current state-of-the-art.

3.13 Relational Analysis: User Attribute Classification

Wikipedia Talk data We focused on predicting two personal attributes from online text: (1) gender and (2) native language. The data used in our research was collected from Wikipedia Talk Page Posts using MediaWiki’s Wikipedia API. For native language identification (NLI), we collected 1300 English posts from users with 5 different native language backgrounds: Bulgarian, Czech, French, Spanish, and Russian. In addition to online text we also considered the International Corpus of Learner English, a collection of ESL student essays. We found we could predict gender 90% of time, native language in essays 84% of time, and native language in posts 56% of time. In addition, we explored cross-corpus native language classification, investigating how well these models generalize across domains.

Twitter We focused on demographic recognition across multiple languages in the Twitter domain. Specifically, we have evaluated how performance for predicting demographic attributes changes between English and Spanish Twitter users. Each user is represented by the aggregation of their tweets and the attributes in question include gender (male vs. female) and occupation (celebrity, politician, or journalist). In addition to exploring linguistic differences between languages, we also investigate previous claims about the communicative behavior of the demographic groups in question. For example, for men versus women, some have claimed that women tend to be more positive in their language use as well as use more inclusive in their speech. Whereas others have noted, men speak more negatively and tend to use first-person singular

pronouns. Also, some have highlighted that, in online forums, male topics of conversation tend to concentrate around prototypical ‘manly’ themes, such as sports. In English, we are able to detect occupation 93.62% and gender 72.66% of the time. In Spanish, we are able to detect occupation 75.16% of the time and gender 82% of the time. However, our best performing models in English for gender and occupation are not the same models in Spanish.

3.14 Prosodic Analysis for Relational Analysis – Name Tagging

We investigated the use of prosodic information from speech to improve Relational Analysis on speech data, especially in where automatic speech recognition (ASR) performance is low. As a first test of this hypothesis we addressed the tasks of Name Tagging on ACE’05 broadcast news data. Using ASR transcription, and the Blender Lab’s CRF-based Name Tagging system, we augmented the feature vectors with two types of prosodic information. The first system extracts low-level acoustic/prosodic features using AuToBI and clusters them into 2-10 categories. The other approach to generating categorical prosodic features is to use AuToBI to generate predictions of ToBI labels. This generates 6 features corresponding to presence/absence of i) pitch accent, ii) intermediate phrase boundary, iii) intonational phrase boundary and iv) pitch accent type, v) phrase accent type and vi) phrase accent/boundary tone pairs. The baseline Name Tagger yielded an f-measure of 27.65% on this data. By extending the feature vector with clustered features, this is improved to 30.57%. Using the AuToBI features, it is further improved to 30.75%. This is a significant improvement, as the ASR on this data is particularly, and intentionally, poor.

3.15 Relational Analysis - Chinese Event Extraction:

We developed a Chinese event extraction system that uses word embedding vectors to represent language, and deep neural networks to learn the abstract feature representation in order to greatly reduce the effort of feature engineering. Our experiments on the ACE Chinese corpus showed that our proposed method performs better compared to the system using rich language features.

We developed a cross-lingual method for Chinese event detection (and sentence classification tasks in general, for low resource languages). It uses convolutional neural networks (CNN) for sentence modeling and classification, based on cross lingual word and phrase embeddings, and dependency grammar representations. The embeddings are trained in an unsupervised manner from data in two languages, with words randomly replaced by their corresponding translations in the other language. For the CNN model, instead of using a window of consecutive words, we use words along with their parent/child words in the dependency parses. The model is trained based on the annotated data in one language, and applied to the test language. The performance is comparable with and sometimes even better than the machine translation (MT) based methods, however, our method does not require the complex MT systems, nor the parallel corpora needed to train the MT systems.

3.16 Relational Analysis - Influential User and Post Identification

We participated in the pilot evaluation task, jointly organized by U. of Washington, Columbia, and our team. The task uses Reddit forum discussion data, and requires the identification of community endorsed person (based on the Karma scores a person obtains) in various sub-reddit forums, and the opinion influencer for ChangeMyView sub-reddit. Our system uses both the social interaction information (e.g., tree structure for the posts and comments) and the textual features derived from a post. Our results show that in general combining the two sources of information improves system performance.

For ChangeMyView we also propose a new task: ranking the posts based on their persuasiveness (Karma scores are used to approximate this). Again we combine social interaction and textual features. In particular, we incorporate some features based on argumentation (since for this domain a post's intention is to change a person's opinion). We find that at the beginning of the discussion, textual features are important, and as the number of comments increases, there is more benefit using social interaction information.

3.17 Recognition of Emotion, Conflict and Laughter

We developed systems to identify speaker state in terms of 1) emotion, 2) conflict, 3) laughter and 4) filled pauses. These will serve as social, communicative signals that can serve as the locus for anomaly. These systems were intended to serve as signals for Anomaly Analysis, but are applicable to First and Second Person IE as well. Our approach to conflict detection found that recognizing overlapping speech is a remarkably effective way to recognize conflict. This led to an unweighted average recall in the classification of low and high conflict of 83.1%.

3.18 Speech Segmentation

We integrated the AuToBI prosodic phrase segmentation into the DEFT codebase and API. While this effort was not directly supported by DEFT, improvements to phrasing have been checked into the repository and are accessible to the DEFT pipeline.

3.19 Anomaly Analysis - Emotion Recognition

We proposed to use the denoising autoencoder to generate robust feature representations for emotion recognition. In our method, the input of the denoising autoencoder (DAE) is the normalized static feature set (state-of-the-art features for emotion recognition). This input is mapped to two hidden representations: one is to capture the neutral information from the input, and the other one is used to extract emotional information. Model parameters are learned by minimizing the squared error between the original and the reconstructed input. After pre-training and fine-tuning, we use the hidden representation as features in the SVM model for emotion classification. Our experimental results show significant performance improvement compared to using the static features. Furthermore, we proposed to model gender information in the DAE framework. The emotional hidden representation is shared for the two genders in order to model more emotion specific characteristics, and is used as features in a back-end classifier for emotion recognition. We propose different optimization objectives in training the DAEs. To address the domain mismatch problem, we also propose a 'shared-hidden-layer autoencoder' (SHLA) approach

for learning common feature representations shared across the training and test set in order to reduce the discrepancy in them. Lastly we also combined the DAE framework with the i-vector representation. Our experimental results using all of these methods show better performance.

3.20 Anomaly Analysis - Disfluency Detection

We proposed a multi-step stacked learning model for disfluency detection. Our method incorporates refined n-gram features step by step from different word sequences and achieved state-of-the-art performance.

3.21 Language Processing Algorithms – Summarization

We investigated compressive summarization in both a pipeline and a joint modeling approach. First, we adopt the 'sentence compression + sentence selection' pipeline approach for compressive summarization, but propose to perform summary guided compression, rather than generic sentence-based compression. During summarization, we use multiple compressed sentences in the integer linear programming framework to select salient summary sentences. Our results on the TAC 2008 and 2011 summarization data sets show that by incorporating the guided sentence compression model, our summarization system can yield significant performance gain as compared to the state-of-the-art. In the joint modeling approach, we propose an efficient decoding algorithm for fast compressive summarization using graph cuts. Our approach first relaxes the length constraint using Lagrangian relaxation. Then we propose to bound the relaxed objective function by the supermodular binary quadratic programming problem, which can be solved efficiently using graph max-flow/min-cut. Since finding the tightest lower bound suffers from local optimality, we use convex relaxation for initialization. Experimental results on TAC2008 dataset demonstrate our method achieves competitive ROUGE score and has good readability, while is much faster than the integer linear programming (ILP) method.

3.22 Language processing Algorithms - Text Normalization and Impact on other NLP Tasks:

Text normalization. We proposed two methods to improve normalization performance.

First is an unsupervised approach that automatically identifies pairs of a non-standard token and proper word from a large unlabeled corpus. We use semantic similarity based on continuous word vector representation, together with other surface similarity measurement. Second we propose a reranking strategy to combine the results from different systems. This allows us to incorporate information that is hard to model in individual systems as well as consider multiple systems to generate a final rank for a test case. Both word- and sentence-level optimization schemes are explored in this study. We evaluate our approach on data sets used in prior studies, and demonstrate that our proposed methods perform better than the state-of-the-art systems.

Joint text normalization and POS tagging. Text normalization and part-of-speech (POS) tagging for social media data have been investigated recently, however, prior work has treated them separately. We propose a joint Viterbi decoding process to determine each token's POS tag and non-standard tokens' correct form at the same time. In order to evaluate our approach, we create two new data sets with POS tag label and nonstandard tokens' correct form. This is the first time to have this data set in social media research community. The experiment results demonstrate the effect of non-standard words on POS tagging, and also show that our proposed methods perform better than the state-of-the-art system both in POS tagging and normalization task.

Named entity recognition. We investigated two methods using non-standard word (NSW) detection results for named entity recognition (NER) in social media data. One adopts a pipeline strategy, and the other uses a joint decoding fashion. We also create a new data set with newly added normalization annotation beyond the existing named entity labels. This is the first data set with such annotation. Our experiment results demonstrate the effectiveness of our NSW detection method and the benefit of NSW detection for NER. Our proposed methods perform better than the state-of-the-art NER system.

Language processing algorithms – parsing: We developed more efficient algorithms to improve the speed of incremental parsing with beam search.^[1] Beam search incremental parsers are accurate, but not as fast as they could be. We demonstrate that, contrary to

popular belief, most current implementations of beam parsers in fact run in $O(n^2)$ time, rather than linear time, because each state-transition is actually implemented as an $O(n)$ operation. We present an improved implementation, based on *Tree Structured Stack*, in which a transition is performed in $O(1)$, resulting in a real linear-time algorithm, which is verified empirically. We further improve parsing speed by sharing feature-extraction and dot-product across beam items. Practically, our methods combined offer a speedup of $\sim 2x$ in parsing over strong baselines.

Language learning algorithms - parallelized online learning: We developed minibatch parallelization algorithms to speedup online structured learning (for parsing and tagging). We developed a simple architecture based on “mini-batches”, which is trivially parallelizable. We show that, unlike previous methods, minibatch learning (in serial mode) actually *improves* the converged accuracy for both perceptron and MIRA learning, and when combined with simple parallelization, minibatch leads to very significant speedups (up to 9x on 12 processors) on state-of-the-art parsing and tagging systems in 2013.

4.0 Results and Discussion

We successfully delivered and fully integrated English joint IE (entity/relation/event) system, Tri-lingual entity discovery and linking system, Chinese slot filling system, Chinese event extraction system and hierarchy detection in online communities system into BBN ADEPT. These technologies have been successfully transited to government agencies including ARL, AFRL and DTRA. In 2017 our multi-lingual entity discovery and linking system and multi-media event recommendation system have been selected and presented on DARPA Demo Day. The TinkerBell Team led by PI Ji developed an end-to-end Cold-Start++ KB construction system. The KB constructed from KBP2017 corpus was used for JHU’s entity recommendation system, which was also presented at DARPA Demo Day. Moreover, BBN presented our entity discover and linking system as part of ADEPT at DARPA

Demo Day.

4.1 NIST TAC-KBP Evaluation Results

- 2012: Top 1 in Slot Filling Validation
- 2013: Top 1 in Slot Filling Validation, Top 2 in Temporal Slot Filling, and Top 3 in English Slot Filling
- 2014: Top 1 in Slot Filling Validation, Top 2 in English Slot Filling, Top 3 in Entity Discovery and Linking, and Top 3 in Event Argument
- 2015: Top 1 in Event Nugget Pilot, Top 1 in Chinese-English Slot Filling Pilot, Top 1 in Event Nugget Detection, Top 1 in Event Nugget Coreference Resolution, Top 1 in Tri-lingual Entity Linking, Top 1 in Slot Filling Validation Filtering
- 2016: Top 1 in Spanish Event Nugget Detection, Top 2 in Chinese Entity Discovery and Linking, and Top 3 in Spanish Entity Discovery and Linking
- 2017: Top 1 in ColdStart++ KBP, Top 1 in Tri-lingual Entity Discovery and Linking (first Window).

4.2 PhDs Supported by DEFT

- Dian Yu. 2017. "Unsupervised Graph-Based Relation Extraction and Validation for Knowledge Base Population"
- Hao Li. 2015. "Knowledge Enrichment for Event Extraction"
- Qi Li. 2015. "Joint Information Extraction"
- Hongzhao Huang. 2015. "Modeling Heterogeneous Networks for Information Ranking, Enrichment and Resolution on Microblogs"
- Taylor Cassidy. 2013. "Temporal Information Extraction and Knowledge Base Population"
- Suzanne Tamang. 2013. "Temporal Sequence Clustering"
- Zheng Chen. 2012. "Collaborative Ranking and Collaborative Clustering"
- Chen Li, 2016. Topic: summarization

- Rui Xia, 2015. Topic: emotion recognition

4.3 Post-docs supported by DEFT

- Yu Hong (2014-2016): Professor at Soochow University
- Jin Guang Zheng (2014-2015): Chief Scientist at Memect
- Haibo Li (2011-2013) : Research Scientist at Nuance
- Arkaitz Zubiaga (2011-2013): Senior Research Fellow at Center for Applied Data Analytics Research

4.4 Principal Investigator Honors

- Ji, World Economic Forum Global Future Council Member on the Future of Computing, 2016-Present
- Ji, "Young Scientist" selected by World Economic Forum, 2016-2017
- Ji, NAACL-HLT2018 Program Committee Chair, 2017-2018
- Ji, IBM Watson Faculty Award, 2012; 2014-2015
- Ji, Google Research Award, 2014-2015
- Ji, "Best of ICDM2013" paper by Knowledge and Information Systems, 2013
- Ji, "Best of SDM2013" paper by Statistical Analysis and Data Mining, 2013
- Ji, "AI's Top 10 to Watch" Award by IEEE Intelligent Systems, 2013
- Ji, PACLIC Best Paper Runner-up, 2012

5.0 Conclusions

This five-year project has been very fruitful. We proposed a new hypothesis that certain common semantic space exists across knowledge element types, languages, data modalities and tasks, therefore knowledge and resources can be effectively transferred across sources. For the first time, we model Information Extraction as a fundamentally different bottom-up discovery problem instead of top-down classification, combining symbolic and distributional semantics. As a result we were able to extend the extraction and knowledge base population capabilities from 33 event types to 3,000+ types, from 5 entity types to 16,000 types, and from 41 relation types to 2,000+ types. We achieved high quality, higher

than / comparable to supervised models, without requiring minimal manual annotations. This project also set a powerful platform for ongoing extensions to multi-lingual in the DARPA LORELEI program and multi-media in the DARPA AIDA program.

6.0 Publications

1. Lifu Huang, Jonathan May, Xiaoman Pan, Heng Ji, Xiang Ren, Jiawei Han, Lin Zhao and James Hendler. 2017. Liberal Entity Extraction: Rapid Construction of Fine-Grained Entity Typing Systems. *Big Data Journal*.
2. Ying Lin, Chin-Yew Lin and Heng Ji. 2017. List-only Entity Linking. *Proceedings of ACL*.
3. Xiaoman Pan, Boliang Zhang, Jonathan May, Joel Nothman, Kevin Knight and Heng Ji. 2017. Cross-lingual Name Tagging and Linking for 282 Languages. *Proceedings of ACL*.
4. Dian Yu, Lifu Huang and Heng Ji. 2017. Open Relation Extraction and Grounding. *Proc. the 8th International Joint Conference on Natural Language Processing (IJCNLP2017)*.
5. Mohamed Al-Badrashiny, Jason Bolton, Arun Tejavsi Chaganty, Kevin Clark, Craig Harman, Lifu Huang, Matthew Lamm, Jinhao Lei, Di Lu, Xiaoman Pan, Ashwin Paranjape, Ellie Pavlick, Haoruo Peng, Peng Qi, Pushpendre Rastogi, Abigail See, Kai Sun, Max Thomas, Chen-Tse Tsai, Hao Wu, Boliang Zhang, Chris Callison-Burch, Claire Cardie, Heng Ji, Christopher Manning, Smaranda Muresan, Owen C. Rambow, Dan Roth, Mark Sammons, Benjamin Van Durme. *TinkerBell: Cross-lingual Cold-Start Knowledge Base Construction. Proc. Text Analysis Conference (TAC2017)*.
6. Heng Ji, Xiaoman Pan, Boliang Zhang, Joel Nothman, James Mayfield, Paul McNamee and Cash Costello. 2017. Overview of TAC-KBP2017 13 Languages Entity Discovery and Linking. *Proc. Text Analysis Conference (TAC2017)*.
7. Tongtao Zhang, Spencer Whitehead, Hanwang Zhang, Hongzhi Li, Joseph Ellis, Lifu Huang, Wei Liu, Heng Ji and Shih-Fu Chang. 2017. Improving Event

- Extraction via Cross-Modal Integration. Proc. the 25th ACM International Conference on Multimedia (ACMMM2017)
8. Lifu Huang, Avirup Sil, Heng Ji and Radu Florian. 2017. Improving Slot Filling Performance with Attentive Neural Networks on Dependency Structures. Proc. Conference on Empirical Methods in Natural Language Processing (EMNLP 2017).
 9. Zhihao Zhou, Lifu Huang and Heng Ji. 2017. Learning Phrase Embeddings from Paraphrases with GRUs. Proceedings of the Workshop on Curation and Applications of Parallel and Comparable Corpora.
 10. Boliang Zhang, Xiaoman Pan, Ying Lin, Tongtao Zhang, Kevin Blissett, Samia Kazemi, Spencer Whitehead, Lifu Huang and Heng Ji. 2017. RPI BLENDER TAC-KBP2017 13 Languages EDL System. Proc. Text Analysis Conference (TAC2017).
 11. Yu Hong, Liang Yao, Mengyi Liu, Tongtao Zhang, Wenxuan Zhou, Jianmin Yao and Heng Ji. 2016. Image-Image Search for Comparable Corpora Construction. Proc. COLING2016 Workshop on Hybrid Approaches to Translation.
 12. Hao Li and Heng Ji. 2016. Cross-genre Event Extraction with Knowledge Enrichment. Proceedings of NAACL-HLT.
 13. Jesse Freitas and Heng Ji. 2016. Identifying News from Tweets. Proceedings of EMNLP.
 14. Yu Hong, Tongtao Zhang, Tim O'Gorman, Sharone Horowitz-Hendler, Heng Ji and Martha Palmer. 2016. Building a Cross-document Event-Event Relation Corpus. Proceedings of ACL2016 Workshop on Linguistic Annotation.
 15. Xiaocheng Feng, Lifu Huang, Duyu Tang, Bing Qin, Heng Ji and Ting Liu. 2016. A Language-Independent Neural Network for Event Detection. Proceedings of ACL.
 16. Xiang Ren, Wenqi He, Meng Qu, Lifu Huang, Heng Ji and Jiawei Han. 2016. AFET: Automatic Fine-Grained Entity Typing by Hierarchical Partial-Label Embedding. Proceedings of EMNLP.
 17. Di Lu, Clare R. Voss, Fangbo Tao, Xiang Ren, Rachel Guan, Rostyslav Korolov, Tongtao Zhang, Dongang Wang, Hongzhi Li, Taylor Cassidy, Heng Ji, Shih-Fu Chang, Jiawei Han, William Wallace, James Hendler, Mei Si and Lance Kaplan.

2016. Cross-media Event Extraction and Recommendation. Proceedings of NAACL-HLT.
18. Hongzhi Li, Joe Ellis, Heng Ji and Shih-Fu Chang. 2016. Event Specific Multimodal Pattern Mining for Knowledge Base Construction. Proceedings of ACMMM.
 19. Dian Yu, Xiaoman Pan, Boliang Zhang, Lifu Huang, Di Lu, Spencer Whitehead and Heng Ji. 2016. RPI_BLENDER TAC-KBP2016 System Description. Proceedings of TAC.
 20. Dian Yu and Heng Ji. 2016. Unsupervised Person Slot Filling based on Graph Mining. Proceedings of ACL.
 21. Di Lu, Xiaoman Pan, Nima Pourdamghani, Shih-Fu Chang, Heng Ji and Kevin Knight. 2016. A Multi-media Approach to Cross-lingual Entity Knowledge Transfer. Proceedings of ACL.
 22. Lifu Huang, Taylor Cassidy, Xiaocheng Feng, Heng Ji, Clare R. Voss, Jiawei Han and Avirup Sil. 2016. Liberal Event Extraction and Event Schema Induction. Proceedings of ACL.
 23. Ann Bies, Zhiyi Song, Jeremy Getman, Joe Ellis, Justin Mott, Stephanie Strassel, Martha Palmer, Teruko Mitamura, Marjorie Freedman, Heng Ji and Tim O'Gorman. 2016. A Comparison of Event Representations in DEFT. Proceedings of

- NAACL2016 Workshop on Events: Definition, Detection, Coreference and Representation.
24. Heng Ji, Joel Nothman and Hoa Trang Dang. 2016. Overview of TAC-KBP2016 Tri-lingual EDL and Its Impact on End-to-End KBP. Proceedings of TAC.
 25. Hao Li, Lin Zhao and Heng Ji. 2015. Social Event Extraction: Task, Challenges and Techniques. Proceedings of ASONAM.
 26. Dian Yu, Heng Ji, Sujian Li and Chin-Yew Lin. 2015. Why Read if You can Scan: Scoping Strategy for Biographical Fact Extraction. Proceedings of NAACL-HLT.
 27. Ofer Bronstein, Ido Dagan, Qi Li, Heng Ji and Anette Frank. 2015. Seed-based Event Trigger Labeling: How far can event descriptions get us? Proceedings of ACL.
 28. Yu Hong, Xiaobin Wang, Yadong Chen, Jian Wang, Tongtao Zhang and Heng Ji. 2015. Biography-Dependent Collaborative Entity Archiving for Slot Filling. Proceedings of EMNLP.
 29. Yu Hong, Di Lu, Dian Yu, Xiaoman Pan, Xiaobin Wang, Yadong Chen, Lifu Huang and Heng Ji. 2015. RPI_BLENDER TAC-KBP2015 System Description. Proceedings of TAC.
 30. Han Wang, Jin Guang Zheng, Xiaogang Ma, Peter Fox and Heng Ji. 2015. Language and Domain Independent Entity Linking with Quantified Collective Validation. Proceedings of EMNLP.
 31. Dian Yu, Yulia Tyshchuk, Heng Ji and William Wallace. 2015. Detecting Deceptive Groups Using Conversations and Network Analysis. Proceedings of ACL.
 32. Boliang Zhang, Hongzhao Huang, Xiaoman Pan, Sujian Li, Chin-Yew Lin, Heng Ji, Kevin Knight, Zhen Wen, Yizhou Sun, Jiawei Han and Bulent Yener. 2015. Context-aware Entity Morph Decoding. Proceedings of ACL.
 33. Tongtao Zhang, Hongzhi Li, Heng Ji and Shih-Fu Chang. 2015. Cross-document Event Coreference Resolution based on Cross-media Features. Proceedings of EMNLP.
 34. Xiaoman Pan, Taylor Cassidy, Ulf Hermjakob, Heng Ji and Kevin Knight. 2015. Unsupervised Entity Linking with Abstract Meaning Representation. Proceedings of NAACL-HLT.
 35. Heng Ji, Joel Nothman, Ben Hachey and Radu Florian. 2015. Overview of TAC-KBP2015 Tri-lingual Entity Discovery and Linking. Proceedings of TAC.
 36. Haibo Li, Masato Hagiwara, Qi Li and Heng Ji. 2014. Comparison of the Impact of Word Segmentation on Name Tagging for Chinese and Japanese. Proceedings of LREC.
 37. Fangbo Tao, Jiawei Han, Heng Ji, George Brova, Chi Wang, Ahmed El-Kishky, Jialu Liu, Xiang Ren, Yizhou Sun. 2014. NewsNetExplorer: Automatic Construction and Exploration of News Information Networks. Proceedings of SIGMOD.
 38. Tongtao Zhang, Haibo Li, Hongzhao Huang, Heng Ji, Min-Hsuan Tsai, Shen-Fu Tsai and Thomas Huang. 2014. Cross-media Cross-genre Information Ranking

- based on Multi-media Information Networks. Proceedings of COLING2014 Workshop on Vision and Language.
39. Yulia Tyshchuk, Hao Li, Heng Ji and William Wallace. 2014. Evolution of Communities on Twitter and the Role of their Leaders during Emergencies. Social Network Analysis - Community Detection and Evolution.
 40. Yu Hong, Xiaobin Wang, Yadong Chen, Jian Wang, Tongtao Zhang, Jin Zheng, Dian Yu, Qi Li, Boliang Zhang, Han Wang, Xiaoman Pan, Heng Ji. 2014. RPI-Soochow KBP2014 System Description. Proceedings of TAC.
 41. Jin Guang Zheng, Daniel Howsmon, Boliang Zhang, Juergen Hahn, Deborah McGuinness, James Hendler and Heng Ji. 2014. Entity Linking for Biomedical Literature. BMC Medical Informatics and Decision Making.
 42. Dian Yu, Hongzhao Huang, Taylor Cassidy, Heng Ji, Chi Wang, Shi Zhi, Jiawei Han, Clare Voss and Malik Magdon-Ismael. 2014. The Wisdom of Minority: Unsupervised Slot Filling Validation based on Multi-dimensional Truth-Finding. Proceedings of COLING.
 43. Taylor Cassidy and Heng Ji. 2014. Analysis and Refinement of Temporal Relation Aggregation. Proceedings of COLING.
 44. Hongzhao Huang, Yunbo Cao, Xiaojiang Huang, Heng Ji and Chin-Yew Lin. 2014. Collective Tweet Wikification based on Semi-supervised Graph Regularization. Proceedings of ACL.
 45. Qi Li and Heng Ji. 2014. Incremental Joint Extraction of Entity Mentions and Relations. Proceedings of ACL.
 46. Qi Li, Heng Ji, Yu Hong and Sujian Li. 2014. Constructing Information Networks Using One Single Model. Proceedings of EMNLP.
 47. Arkaitz Zubiaga, Heng Ji and Kevin Knight. 2013. Curating and Contextualizing Twitter Stories to Assist with Social Newsgathering. Proceedings of IUI.
 48. Arkaitz Zubiaga and Heng Ji. 2013. Harnessing Web Page Directories for Large-Scale Classification of Tweets. Proceedings of WWW.
 49. Arkaitz Zubiaga and Heng Ji. 2013. Tweet, but Verify: Epistemic Study of Information Verification on Twitter. Social Network Analysis and Mining.
 50. Weiwei Guo, Hao Li, Heng Ji and Mona Diab. 2013. Linking Tweets to News: A Framework to Enrich Online Short Text Data. Proceedings of ACL.
 51. Haibo Li, Jing Zheng, Heng Ji, Qi Li and Wen Wang. 2013. Name-aware Machine Translation. Proceedings of ACL.
 52. Hongzhao Huang, Zhen Wen, Dian Yu, Heng Ji, Yizhou Sun, Jiawei Han and He Li. 2013. Resolving Entity Morphs in Censored Data. Proceedings of ACL.
 53. Qi Li, Heng Ji and Liang Huang. 2013. Joint Event Extraction via Structured Prediction with Global Features. Proceedings of ACL.
 54. Zhongyu Wei and Yang Liu. Is This Post Persuasive? Ranking Argumentative Comments in Online Forum, ACL, 2016.
 55. Yandi Xia, Zhongyu Wei, and Yang Liu. An Efficient Cross-lingual Model for Sentence Classification Using Convolutional Neural Network. Student research workshop at ACL, 2016.
 56. Zachary Stallbohm, Yang Liu, and Zhongyu Wei. Negation Effect on Emotion Analysis in Twitter. Social NLP at IJCAI, 2016.

57. Rui Xia and Yang Liu. DBN-IVector Framework for Acoustic Emotion Recognition. Interspeech, 2016.
58. Zhongyu Wei, Yandi Xia, Chen Li, Yang Liu, Zachary Stallbohm, Yi Li, and Yang Jin. Disputation Behavior Analysis in the Online Debating Forum. 3rd Workshop on Argumentation Mining at ACL, 2016.
59. Chen Li and Yang Liu. Improving Named Entity Recognition in Tweets via Detecting Non-standard Words. ACL, 2015.
60. Xian Qian and Yang Liu. Feature Selection in Kernel Space: A Case Study on Dependency Parsing. ACL, 2015.
61. Chen Li, Yang Liu, Lin Zhao. Using External Resources and Joint Learning for Bigram Weighting in ILP-Based Multi-Document Summarization. NAACL, 2015.
62. Chen Li, Yang Liu, Lin Zhao. Improving Update Summarization via Supervised ILP and Sentence Reranking. NAACL, 2015.
63. Chen Li and Yang Liu. Joint POS Tagging and Text Normalization for Informal Text. IJCAI, 2015.
64. Xian Qian and Yang Liu. 2-Slave Dual Decomposition for Generalized High Order CRFs. Transactions of the Association for Computational Linguistics (TACL). 2014.
65. Chen Li, Yang Liu, Fei Liu, Lin Zhao, and FuliangWeng. Improving Multi-documents Summarization by Sentence Compression based on Expanded Constituent Parse Trees. EMNLP 2014.
66. Rui Xia, Jun Deng, Bjorn Schuller, and Yang Liu. Modeling Gender Information for Emotion Recognition Using Denoising Autoencoder. ICASSP 2014.
67. Jun Deng, Rui Xia, Zixing Zhang, Yang Liu, and Bjorn Schuller. Introducing Shared-hidden-layer Autoencoder for Transfer Learning and their Application in Acoustic Emotion Recognition. ICASSP, 2014.
68. Xian Qian and Yang Liu. Polynomial Time Joint Structural Inference for Sentence Compression. ACL, 2014.
69. Chen Li and Yang Liu. Improving Text Normalization via Unsupervised Model and Discriminative Reranking. ACL Student Research Workshop, 2014.
70. Rui Xia and Yang Liu. Using Denoising Autoencoder for Emotion Recognition. Proceedings of Interspeech, 2013.
71. Xian Qian and Yang Liu. Fast Joint Compression and Summarization via Graph Cuts. Proceedings of EMNLP, 2013.
72. Chen Li, Fei Liu, Fuliang Weng, and Yang Liu. Document Summarization via Guided Sentence Compression. Proceedings of EMNLP, 2013.
73. Guozhen An, David Guy Brizan, and Andrew Rosenberg. Detecting Laughter and Filled Pauses Using Syllable-based Features. Interspeech 2013. Keywords: laughter detection, filled pause detection, prosodic analysis
74. Felix Grezes, Justin Richards, and Andrew Rosenberg. Let me Finish: Automatic Conflict Detection Using Speaker Overlap. Interspeech 2013. Keywords: conflict detection, paralinguistics.
75. Denys Katerenchuk, Andrew Rosenberg. Improving Named Entity Recognition with Prosodic Features. Interspeech 2014.
76. Denys Katerenchuk, David-Guy Brizan, and Andrew Rosenberg. “Was that your

Mother on the Phone?" Interspeech 2014

77. Denys Katerenchuk, Andrew Rosenberg, "RankDCG: Rank-Ordering Evaluation Measure," LREC, 2016, Portoroz, Slovenia.
78. Denys Katerenchuk and Andrew Rosenberg. "Hierarchy Prediction in Online Communities." Student session, Association for the Advancement of Artificial Intelligence (AAAI) 2016, Phoenix, AZ, US.