

XIAODAN ZHU

APT-1822, 30 Charles St. W.
Toronto, ON, M4Y 1R5
xzhu@cs.toronto.edu
+1 (647) 405-2266

Education

2003–2009 Ph.D. (**expected: Sep.**), Department of Computer Science, University of Toronto, Canada
1997–2000 M.S., Department of Computer Science, Tsinghua University, China
1993–1997 B.S.; Department of Computer Science, North China University of Technology, China

GPA

Ph.D. – 4.0/4.0 M.S. – Overall 3.5/4.0; Major 3.8/4.0 B.S. – Overall 3.8/4.0; Major 3.9/4.0

Research Experience

2003–
present **University of Toronto; Ph.D. student**

(1) **Salient utterance extraction from spoken documents** (Ph.D. thesis)

My thesis studied utterance extraction and redundancy removal in spoken documents. The effectiveness of a wide variety of textual features, speech-related features, and speech recognition errors was carefully examined [3][8]. Instead of treating spoken documents by themselves, we also utilized additional (e.g., semantic) knowledge available in relevant written text (e.g., textbooks for lecture recordings) [4][6]. The redundancy removal issue was addressed among multiple spoken documents that discuss the same topic [2]. All the research and related work will be assembled in a book we are composing [1].

(2) **Sentiment analysis and question-answering in medical text**

In a side project, we studied the polarity detection of clinic outcomes through combining linguistic and domain knowledge [10]. The results were used to boost the performance of question-answering [7] for medical text.

May –Sep.
2007 **IBM T.J. Watson Research Center (New York); Intern**
Machine translation

I worked in a project with the aim to improve a state-of-the-art speech-to-speech machine translation system. We proposed an improved formally syntax-based SMT model [5], which is enriched by linguistically syntactic knowledge obtained from statistical constituent parsers. A linguistically-motivated prior derivation model was

used to score hypothesis derivations on top of the baseline model during the translation decoding. Moreover, we devised a fast training algorithm to achieve such improved models based on tree kernel methods, Experiments on an English-to-Chinese task demonstrated that our proposed models outperformed the baseline formally syntax-based models, while both of them achieved significant improvements over a state-of-the-art phrase-based SMT system. For this project, my mentor and I also implemented a hierarchical phrase-based system from scratch.

Jun.–Sep.
2006

Google Inc. (New York); Intern
Entity and relation extraction

I designed and implemented an unsupervised method to detect entities and their relations from the Web. The algorithm worked by bootstrapping. Given a seed (any type of entity pairs or n-tuples such as <Bill Gates, Microsoft, Harvard>), the algorithm found all HTML-tag sequences that spanned over these words, e.g., HTML tables or other much looser structures. Reliable HTML-tag patterns were then learned through clustering. New patterns acquired were applied to extract more entities and their relation labels. URL domains were clustered to expand relations. The algorithm was paralleled and run on the whole Google Web repository within hours.

Jul.–Aug.
2005

Avaya Labs Research (Denver); Intern
Call-type classification

My duty was to conduct call-type classification for customer-service call-centers. By using the agent-side audio only and through speaker adaptive training, we found that the overall classification performance was comparable to that obtained by using both agents' and customers' audio, which suggests that there is no need to recognize customers' speech, given the inevitable higher word error rates on it.

Nov. 2001–
Sep. 2003

Microsoft Research (Beijing); Visiting Scholar
Information extraction

I applied log-linear reranking methods to improve the performance of a state-of-the-art source-channel model for Chinese named-entity identification. Transformation-based learning was used afterwards to adapt the system for different standards used in SIGHAN evaluation. I also studied the identification of single-character Chinese named entities [12]. I built a finite-state-transducer toolkit to extract factoids (named entities other than persons, locations, and organizations), which includes the functionalities of determinization (for sequentiable or subsequentially FSTs) and building bi-machines (for unambiguous FSTs).

Jul.2000–
Nov. 2001 **Intel Research (Beijing); Full-time Researcher**
Information retrieval

I studied the indexing and search of closed caption of Chinese broadcast news [14]. I also studied Chinese query processing to match users' natural-language (non-keyword) queries with the frequently-asked-question (FAQ) database. New queries were detected and tracked in order to update the FAQ database of Intel's corporate-wide technical call-center recordings.

Sep.1997–
Jul. 2000 **Tsinghua University; Master student**
Information extraction

My Master's thesis studied two typical information extraction tasks: named-entity identification and template filling, in a limited domain: financial news. I adapted general-purpose models for Chinese word segmentation and shallow parsing to this specific domain. The named-entity extraction was based on a source-channel model and the template filling also used the parsing output. I studied general Chinese temporal representation with the aim to understand the temporal relationship between financial events [13][15]. I also joined a side project that studied the combinatorial regulations of Chinese semantic classes [17].

Publications

- [1] Gerald Penn and Xiaodan Zhu. Speech Summarization. Book under composing. Publisher: Morgan & Claypool.
- [2] Xiaodan Zhu, Gerald Penn and Frank Rudzicz. Summarizing multiple spoken documents: finding evidence from untranscribed audio. (To appear) The 47th Annual Meeting of the Association for Computational Linguistics (ACL-2009), Singapore.
- [3] Gerald Penn and Xiaodan Zhu. A critical reassessment of evaluation baselines for speech summarization of spontaneous conversations. The 46th Annual Meeting of the Association for Computational Linguistics (ACL-2008), Columbus, USA.
- [4] Xiaodan Zhu, Xuming He, Cosmin Munteanu, and Gerald Penn. Using latent Dirichlet allocation to incorporate domain knowledge for topic transition detection. Proceedings of the International Conference on Spoken Language Processing (Interspeech-2008), Brisbane, Australia.
- [5] Bowen Zhou, Bing Xiang, Xiaodan Zhu, and Yuqing Gao. Prior derivation models for formally syntax-based translation using linguistically syntactic parsing and tree kernels. ACL-2008 Second Workshop on Syntax and Structure in Statistical Translation, Columbus, USA.

- [6] Xiaodan Zhu, Siavash Kazemian, and Gerald Penn. Identifying salient utterances from Web spoken documents using descriptive hypertext. IEEE Workshop on Spoken Language Technology (SLT-2008), Goa, India.
- [7] Yun Niu, Xiaodan Zhu, and Graeme Hirst. Question answering in the medical domain: the role of clinical outcome and polarity. Proceedings of the American Medical Informatics Association 2006 Annual Symposium (AMIA-2006), Washington, D.C., USA.
- [8] Xiaodan Zhu and Gerald Penn. Summarization of spontaneous conversations. Proceedings of the 9th International Conference on Spoken Language Processing (Interspeech-2006), Pittsburgh, Pennsylvania, USA.
- [9] Xiaodan Zhu and Gerald Penn. Comparing the roles of textual, acoustic and spoken-language features on spontaneous-conversation summarization. Proceedings of the 11th Human Language Technology Conference / 5th Meeting of the North American Chapter of the Association for Computational Linguistics (HLT/NAACL-06) (short), New York, USA.
- [10] Yun Niu, Xiaodan Zhu, Jianhua Li, and Graeme Hirst. Analysis of polarity information in medical text. Proceedings of the American Medical Informatics Association 2005 Annual Symposium (AMIA-2005), Washington, D.C.
- [11] Xiaodan Zhu and Gerald Penn. Evaluation of sentence selection for speech summarization. RANLP-2005 Crossing Barriers in Text Summarization Research Workshop, Borovets, Bulgaria.
- [12] Xiaodan Zhu, Mu Li, Jianfeng Gao, and Chang-Ning Huang. Single character Chinese named entity recognition, ACL-2003 Second SIGHAN Workshop on Chinese Language Processing, Sapporo, Japan.
- [13] Kam-Fai Wong, Wenjie Li, Chunfa Yuan, and Xiaodan Zhu. Temporal representation and classification in Chinese. International Journal of Computer Processing of Oriental Languages. 15(2): 211-230 (2002).
- [14] Xiaodan Zhu, Qian Diao, and Joe F. Zhou. A two-character hash function for Chinese word indexing. In Proceedings of the 6th Joint Conference of Computational Linguistics (in Chinese), 2001.
- [15] Xiaodan Zhu and Chunfa Yuan. An algorithm for situation classification of Chinese verbs, ACL-2000 Second Workshop on Chinese Language Processing, Hong Kong, China.
- [16] Xiaodan Zhu. Information extraction from financial news and the related temporal information analysis. Masters thesis (in Chinese), Tsinghua University, 2000.

- [17] Wei Xu, Chunfa Yuan, Changning Huang, and Xiaodan Zhu. A study on the combinatorial regulation of Chinese semantic classes, Communication of COLIPS.

Programming Experience (Non-research oriented)

- Jul.–Oct.
1998 **China Education Electronic Corp.;** Intern
For a multi-media classroom project, I implemented an audio communication component, which was based on the TCP/IP and Socket protocols and implemented with multiple-thread technology to ensure the smoothness of peer-to-peer speech and broadcasting among teachers and students. (VC++ and Delphi)
- Mar.–May
1998 Designed and implemented a hardware-controlling component of an expressway billing system for Guangdong Highway Management Bureau, China. (VC++)
- Jan.–Jun.
1997 **China National Software & Service Company Limited;** Intern
Built a real-time billing component for a GSM cellar-phone billing system, which fetched and processed data from exchangers in real-time. (VC++, Visual Basic)
- Apr.–Jul.
1995 Implemented a simplified preemptive multi-process demonstration system for the Computing Center of North China University of Technology. (8086/80286 Assembly language)

I also have extensive research-oriented development experience on both text and speech processing, which is closely associated with my research experience above and hence not listed here separately. I am skilled in the following programming languages:

C/C++, Perl, Python, Java, Shell Scripts, Matlab, Assembly, Delphi, Prolog.

Teaching Experience

2003–
present **Teaching assistant at the University of Toronto** (more than 20 appointments)

Graduate courses:

- CSC 2511 / 401 - Natural Language Computing
- CSC 2511 / 418 - Computer Graphics

Undergraduate courses:

- CSC 411 - Machine Learning and Data Mining
- CSC 309 - Web Programming
- CSC 190 - Data Structures and Algorithms
- CSC 180 - Introduction to Computer Programming

Awards and Prizes

- Wolfond Fellowship, University of Toronto, 2008-2009.
- IBM PhD Fellowship Finalist, 2008-2009.
- Ontario Graduate Scholarship, University of Toronto, 2008-2009.
- Ph.D. Fellowship, University of Toronto, 2003-2008.
- Wang Jiarong Scholarship, Tsinghua University, 1999-2000
- Guanghua Scholarship, Tsinghua University, 1998-1999
- First-Class Award, Creative Design Competition in Beijing Area, Tsinghua University, 1999
- Excellent Undergraduate Thesis Award, North China University of Technology, 1997
- First Class Scholarship, North China University of Technology, 1994 - 1997
- Excellent Student Award of Engineering College, North China University of Technology, 1994

Immigration Status

Permanent Resident of Canada; Citizen of People's Republic of China

References

Gerald Penn

Associate Professor
Department of Computer Science
University of Toronto
+1 (416) 978-7390
gpenn@cs.toronto.edu

Changning Huang

Senior Researcher
Microsoft Research, Beijing
+86 (10) 6261-7711
cnhuang@microsoft.com

Bowen Zhou

Research Staff Member
IBM T.J. Watson Research Center
+1 (914) 945-2852
zhou@us.ibm.com

Shubin Zhao

Software Engineer
Google Inc.
+1 (212) 589-8726
shubin@google.com