

Alexander H. Gruenstein

M.I.T. Computer Science and
Artificial Intelligence Laboratory
32 Vassar Street #32-G494
Cambridge, MA 02139, USA

alexgru@csail.mit.edu
<http://www.mit.edu/~alexgru/>

Education

Ph.D. in Computer Science
Supervisor: Dr. Stephanie Seneff
GPA: 5.0 / 5.0
Thesis Title: *Context-Sensitive Techniques for the Prevention, Detection, and Correction of Errors in Widely-Accessible Multimodal Dialogue Systems*

Massachusetts Institute of Technology
2004-Present

M.S. in Symbolic Systems
Concentration: *Natural Language Technologies*
Supervisors: Prof. Stanley Peters, Dr. Oliver Lemon
GPA: 4.07 / 4.00
Thesis Title: *Conversational Interfaces: A Domain-Independent Architecture for Task-Oriented Dialogues*

Stanford University
2001-2003

B.S. in Symbolic Systems
Concentration: *Artificial Intelligence*
Graduated with *Distinction*; member *Phi Beta Kappa*
Symbolic Systems is an interdisciplinary major incorporating Computer Science, Linguistics, Philosophy, Logic, and Cognitive Psychology.
GPA: 3.88 / 4.00

Stanford University
1998-2003

Experience

Spoken Language Systems Group, Computer Science and Artificial Intelligence Laboratory M.I.T.
Research Assistant Sep 2004 - Present

Primarily focused on building the next generation of multimodal dialogue systems. Developing an infrastructure for deploying multimodal dialogue systems over the web and in-car, and evaluating real user experiences with such systems – notably with a novel map-based system called *City Browser* which provides multimodal access to urban information. Constructing methods for integrating knowledge about discourse context throughout all dialogue system components, including the development of several new dynamic language modeling and confidence scoring techniques which have been shown to improve speech recognition and concept accuracy.

Experience (continued)

Google Mountain View, CA
Software Engineering Intern Jun 2008 - Aug 2008

Worked as a member of the Speech group to conceive and develop multimodal speech applications.

Vlingo Cambridge, MA
Speech Technology Intern Jun 2007 - Sep 2007

Developed statistical language modeling techniques to improve speech recognition accuracy in very large vocabulary applications developed for mobile phones. For instance, Vlingo provides Yahoo's mobile search application.

Computational Semantics Lab, Center for the Study of Language and Information Stanford
Research Software Engineer Jul 2003 - Aug 2004; Summer 2005
Undergraduate/Graduate Research Assistant Jun 2000 - Dec 2002

Worked both full time as a Research Software Engineer and as a graduate and undergraduate student on various natural language technology projects. Developed a generic multimodal dialogue-management architecture which has been deployed at the core of multiple multimodal interfaces developed in the lab, including: an interface for controlling autonomous robots, a meeting assistant, an in-car dialogue system, and a tutorial dialogue system. Created NOMOS, a flexible interface for annotating and transcribing recorded media from natural language interactions, especially multiparty meetings. Collected a corpus of annotated meetings to investigate automatic topic segmentation.

BeVocal, Inc Mountain View, CA
Software Engineer Jan 2003 - Jul 2003

Developed voice applications using VoiceXML in a J2EE framework. Primarily developed application suites for leading wireless carriers for information-seeking (*e.g.* movie information, sports scores) and customer service voice applications. Conceived, researched, and developed prototypes for the next generation of marketable applications.

Archimedes Project, Center for the Study of Language and Information Stanford
Undergraduate Research Assistant Sep 1999 - Jun 2000

Designed machine-vision software which analyzed graphical user interfaces to present distilled information haptically (through touch) to visually impaired individuals.

Service

Organizing Committee, Young Researchers' Roundtable on Spoken Dialog Systems 2006
Advisory Committee, Young Researchers' Roundtable on Spoken Dialog Systems 2005, 2007
Various conference and journal paper reviewing

Technical Skills

Natural Language Technologies:

Speech Recognition: SUMMIT, Nuance, Sphinx

Parsing: TINA, Gemini, Prolog, LKB

Synthesis: Festival, Dectalk

Language Modeling/Finite State Transducers: SLS Toolkit, Xerox Finite State Toolkit, SRILM

Inter-Agent Communication: GALAXY, Open Agent Architecture

Programming Languages:

Highly proficient in: Java, J2EE/JSP, C, Perl, VoiceXML

Proficient in: Javascript/DHTML, C++, Prolog, Lisp, Matlab

Customary Tools:

Eclipse, Emacs, Ant, CVS, Subversion, Tomcat, Apache, Latex, Office, TWiki;

Windows, Linux, Cygwin, Unix

Language

English, native speaker

Spanish, studied 6 years up to high school Advanced Placement level

Publications

*

Journal Papers

- [1] Oliver Lemon and Alexander Gruenstein. Multithreaded context for robust conversational interfaces: context-sensitive speech recognition and interpretation of corrective fragments. *ACM Transactions on Computer-Human Interaction*, 11(3):241–267, 2004.
- [2] Oliver Lemon, Alexander Gruenstein, and Stanley Peters. Collaborative activities and multi-tasking in dialogue systems. *Traitement automatique des langues*, 43(2):131–154, 2002. Special issue on dialogue.

*

Book Chapters

- [1] Alexander Gruenstein, John Niekrasz, and Matthew Purver. Meeting structure annotation. In Laila Dybkjaer and Wolfgang Minker, editors, *Recent Trends in Discourse and Dialogue*. Springer, 2008.
- [2] Brady Clark, Oliver Lemon, Alexander Gruenstein, Elizabeth Owen Bratt, John Fry, Stanley Peters, Heather Pon-Barry, Karl Schultz, Zack Thomsen-Gray, and Pucktada Treeratpituk. A general purpose architecture for intelligent tutoring systems. In Niels Ole Bernsen, Laila Dybkjaer, and Jan van Kuppevelt, editors, *Natural, Intelligent and Effective Interaction in Multimodal Dialogue Systems*. Kluwer, in press.
- [3] Oliver Lemon, Anne Bracy, Alexander Gruenstein, and Stanley Peters. An information state approach in a multi-modal dialogue system for human-robot conversation. In Hans Reiser, Peter Kühnlein, and Henk Zeevat, editors, *Perspectives on Dialogue in the new Millenium*, volume 114 of *Pragmatics and Beyond*, pages 229–242. John Benjamins, 2003.

Conference Papers

- [1] Alexander Gruenstein, Ian McGraw, and Ibrahim Badr. The WAMI toolkit for developing, deploying, and evaluating web-accessible multimodal interfaces. In *Proc. of 10th International Conference on Multimodal Interfaces*, 2008.
- [2] Alexander Gruenstein. Response-based confidence annotation for spoken dialogue systems. In *Proc. of the 9th SIGdial Workshop on Discourse and Dialogue*, 2008.
- [3] Alexander Gruenstein, Bo-June (Paul) Hsu, James Glass, Stephanie Seneff, Lee Hetherington, Scott Cyphers, Ibrahim Badr, Chao Wang, and Sean Liu. A multimodal home entertainment interface via a mobile device. In *Proc. of the ACL Workshop on Mobile Language Processing*, 2008.
- [4] Alexander Gruenstein and Stephanie Seneff. Releasing a multimodal dialogue system into the wild: User support mechanisms. In *Proc. of the 8th SIGdial Workshop on Discourse and Dialogue*, 2007.
- [5] Alexander Gruenstein and Stephanie Seneff. Context sensitive language modeling for large sets of proper nouns in multimodal dialogue systems. In *Proc. of IEEE/ACL Workshop on Spoken Language Technology*, 2006.
- [6] Alexander Gruenstein, Stephanie Seneff, and Chao Wang. Scalable and portable web-based multimodal dialogue interaction with geographical databases. In *Proc. of INTERSPEECH-ICSLP*, 2006.
- [7] John Niekrasz and Alexander Gruenstein. NOMOS: A semantic web software framework for annotation of multimodal corpora. In *Proc. of the 5th Conference on Language Resources and Evaluation (LREC 2006)*, 2006.
- [8] Alexander Gruenstein, John Niekraz, and Matthew Purver. Meeting structure annotation: Data and tools. In *Proc. of the 6th SIGdial Workshop on Discourse and Dialogue*, 2005.
- [9] Alexander Gruenstein, Chao Wang, and Stephanie Seneff. Context-sensitive statistical language modeling. In *Proc. of INTERSPEECH*, pages 17–20, 2005.
- [10] Alexander Gruenstein and Lawrence Cavedon. Using an activity model to address issues in task-oriented dialogue interaction over extended periods. In *Proceedings of AAAI Spring Symposium on Interaction Between Humans and Autonomous Systems over Extended Periods*, March 2004.
- [11] Alexander Gruenstein, Lawrence Cavedon, John Niekrasz, Dominic Widdows, and Stanley Peters. Managing uncertainty in dialogue information state for real time understanding of multi-human meeting dialogues. In *Proceedings of the 8th Workshop on Formal Semantics and Pragmatics of Dialogue (Catalog)*, 2004.
- [12] John Niekrasz, Alexander Gruenstein, and Lawrence Cavedon. Multi-human dialogue understanding for assisting artifact-producing meetings. In *Proceedings of the 20th International Conference on Computational Linguistics (COLING)*, 2004.
- [13] Ed Kaiser, David Demirdjian, Alexander Gruenstein, Xiaoguang Li, John Niekrasz, Matt Wesson, and Sanjeev Kumar. Demo: A multimodal learning interface for sketch, speak and point creation of a schedule chart. In *Proceedings of the Sixth International Conference on Multimodal Interfaces (ICMI 2004)*, 2004.
- [14] Anna Liess, Wendy Ellis, Janine Giese-Davis, Alexander Gruenstein, Mitch Golant, and David Spiegel. Emotional information available from videotapes vs transcripts. In *Proceedings of the 25th Annual Meeting of the Society of Behavioral Medicine*, March 2004.
- [15] Oliver Lemon, Alexander Gruenstein, Randolph Gullett, Alexis Battle, Laura Hiatt, and Stanley Peters. Generation of collaborative spoken dialogue contributions in dynamic task environment. In Reva Freedman and Charles Callaway, editors, *Working Papers of the 2003 AAAI Spring Symposium on Natural Language Generation in Spoken and Written Dialogue*, pages 85–90. AAAI Press, 2003.

- [16] Beth Ann Hockey, Oliver Lemon, Ellen Campana, Laura Hiatt, Gregory Aist, James Hieronymus, Alexander Gruenstein, and John Dowding. Targeted help for spoken dialogue systems: Intelligent feedback improves naive user's performance. In *Proceedings of the 10th Conference of the European Chapter of the Association for Computational Linguistics (EACL)*, April 2003.
- [17] Oliver Lemon, Alexander Gruenstein, Lawrence Cavedon, and Stanley Peters. Collaborative dialogue for controlling autonomous systems. In *Proceedings of the AAAI Fall Symposium*, 2002.
- [18] Oliver Lemon, Alexander Gruenstein, Alexis Battle, and Stanley Peters. Multi-tasking and collaborative activities in dialogue systems. In *Proceedings of the 3rd SIGdial Workshop on Discourse and Dialogue*, pages 113–124, 2002.
- [19] Oliver Lemon, Alexander Gruenstein, Alexis Battle, Elizabeth Bratt, and Stanley Peters. Multi-tasking in practical multi-modal dialogue systems. In *Proceedings of the Association for Computational Linguistics (ACL)*, pages 106–107, July 2002. Demo.
- [20] Oliver Lemon, Anne Bracy, Alexander Gruenstein, and Stanley Peters. A multi-modal dialogue system for human-robot conversation. In *Proceedings of the Second Meeting of the North American Chapter of the Association for Computational Linguistics NAACL*, June 2001. Demo.
- [21] Oliver Lemon, Anne Bracy, Alexander Gruenstein, and Stanley Peters. The witas multi-modal dialogue system I. In *7th European Conference on Speech Communication and Technology (EuroSpeech)*, 2001.
- [22] Oliver Lemon, Anne Bracy, Alexander Gruenstein, and Stanley Peters. Multi-modal dialogue systems supporting human-robot communication. In *5th World Multi-Conference on Systemics, Cybernetics and Informatics (SCI)*, July 2001.
- [23] Oliver Lemon, Anne Bracy, Alexander Gruenstein, and Stanley Peters. Information states in a multi-modal dialogue system for human-robot conversation. In Peter Kühnlein, Hans Reiser, and Henk Zeevat, editors, *Proceedings of the 5th Workshop on Formal Semantics and Pragmatics of Dialogue (Bi-Dialog 2001)*, pages 57 – 67, 2001.