Eric Fosler-Lussier

fosler@ieee.org
http://www.icsi.berkeley.edu/~fosler

Columbia University, Dept. of Electrical Engineering
Mail Code 4712
500 West 120th Street, Room 1312
New York NY 10027

3204 Astor Close
Hillsborough, NJ 08844
(908) 874-6632

Research Interests

My research expertise includes automatic speech recognition, spoken natural language processing, and computational linguistics, with specializations in phonological modeling within speech recognition systems and spoken user interface design. I am also interested in statistical investigations of linguistic phenomena, as well as combining multiple sources of information within natural language and speech systems.

Education

University of California, Berkeley Ph.D., Computer Science 12/99

Dissertation: *Dynamic Pronunciation Modeling for Automatic Speech Recognition*.

Examined how factors such as speaking rate and word predictability affect both word pronunciations and performance of automatic speech recognition (ASR) systems; employed pronunciation models that change dynamically according to linguistic context in an ASR system.

Committee: Nelson Morgan (adviser), Jerome Feldman, Steven Greenberg, John Ohala

University of Pennsylvania B.A., Linguistics, *summa cum laude* with distinction in the major 5/93

Adviser: Mark Liberman

University of Pennsylvania B.A.S., Cognitive Science, *summa cum laude* 5/93

Adviser: Mitchell Marcus

Honors

Invited lecturer, Eighth European Language and Speech Network (ELSNET) Summer School on Language and Speech Communication, Chios Island, Greece, 2000.

European Language Resources Association Prize for Best Student Paper of the 6th European Conference on Speech Communication and Technology, 1999.

Invited Papers, International Congress of Phonetic Sciences (ICPhS 1999), International Conference on Spoken Language Processing (ICSLP 1996).

Invited Researcher, Johns Hopkins Large Vocabulary Continuous Speech Recognition Summer Research Workshop, 1996.

National Defense Science Engineering Graduate Fellowship Honorable Mention, 1994.

National Science Foundation Graduate Fellowship Honorable Mention, 1993 and 1994.

Phi Beta Kappa, 1993.

University Scholar, University of Pennsylvania, 1989–1993.

Benjamin Franklin Scholar, University of Pennsylvania, 1989–1993.

Grants

"Incorporating Higher-Level Information into Dynamic Pronunciation Modeling for ASR," National Science Foundation SGER grant IRI-9713346, 1997–1998.

"Automatic Learning of a Model for Word Pronunciations," Johns Hopkins Center for Language and Speech Processing Large Vocabulary Continuous Speech Recognition Summer Research Workshop follow-on grant competition, 1996–1997.

Research Experience

Columbia University, Department of Electrical Engineering Visiting Scientist 1/03–present

Currently developing new algorithms for automatic speech recognition that incorporate multiple sources of information within the recognition process (with D. Ellis). Also collaborating with faculty and students in the Computer Science department to design novel techniques for the automatic identification of topic boundaries in recorded meetings.

Bell Labs, Lucent Technologies Member of Technical Staff, Research 8/00–12/02

As a member of the Dialogue Systems Research Department, provided technology innovations in spoken dialogue systems and related automatic speech recognition (ASR) and natural language (NL) areas. Achievements include:

- Invented a new framework for semantic and pragmatic interpretation in spoken dialogue systems (with E. Ammicht and A. Potamianos) that combines evidence of the user's intent over multiple dialogue turns, and has methods of resolving language ambiguities and errors made by the system in previous turns.

- Designed a novel set of algorithms for determining the lexical confusability inherent within a speech recognition system in order to promote better lexicon design. Led research team (including I. Amdal and J. Kuo) in experiments validating the approach.

- As part of a research team, developed a set of spoken dialogue systems to provide interfaces to internet-based services (e.g., email, travel reservations); personal contributions included spoken user interface design, providing internet service connections, and creating finite state machine representations of the application.

- Devised a new approach to natural language generation for spoken dialogue systems (i.e., the part of a system that determines what the computer agent says) that reuses information encoded in the part of the system that interprets what the user says; the system is quick to implement for a new domain, and can take advantage of human-to-human conversations as training data (with M. Galley and A. Potamianos).

- Also conducted research in automatic semantic class induction (finding groups of related words in text or speech data), prosodic analysis of speech, ASR language modeling, multi-modal (speech/keyboard/pen) system integration, design of VoiceXML based services, recursive finite state parsing, and natural language call routing.

International Computer Science Institute Postdoctoral researcher 9/99–8/00

Managed ICSI participation in SmartKom, a German consortial project directed towards building speech-enabled kiosks and personal data assistants; integrated the ICSI speech recognizer into the SmartKom system. Informally advised graduate students in both individual and group research projects.

University of California, Berkeley and the International Computer Science Institute Graduate Student Researcher 11/93–9/99

- Developed a method of dynamic pronunciation modeling for ASR systems that incorporated information about word predictability and speaking rate in determining the correct set of pronunciations to use.

- Conducted linguistic analyses of pronunciations in two large speech recognition corpora, showing correlations between word frequency, speaking rate, word pronunciations, and ASR errors.

- With Nelson Morgan and Nikki Mirghafori, developed methods of estimating speaking rate, both directly from the speech signal and from a first-pass recognition. Investigated ways to compensate for errors due to fast speaking rate.

- Also conducted research in probabilistic phonological rules for ASR, decoders that integrate multiple streams of acoustic information at syllable or word time scale, and a spoken dialogue system with knowledge of restaurants in Berkeley.

Johns Hopkins University
Large Vocabulary Continuous Speech Recognition Summer Workshop 7/96–8/96

As part of a team of researchers in a six-week workshop, developed automatic pronunciation-learning techniques using decision tree models.

Teaching Experience

Bell Labs, Lucent Technologies Member of Technical Staff, Research 8/00–12/02

Bell Labs hosts students for limited-term research projects that can extend from 10 weeks to 6 months. Served as project advisor or co-advisor for 2 undergraduate and 4 graduate students. One student, Michel Galley, won a best student researcher prize at the 2001 DARPA Communicator PI meeting for our natural language generation work.

ELSNET Summer School Invited Instructor 8/00

Developed and taught a one-week graduate-level course on pronunciation modeling for automatic speech recogition, with a focus on machine learning techniques and statistical descriptions of linguistic phenomena.

University of California, Berkeley Graduate Student Instructor 8/95–11/95

Teaching assistant with Prof. Jitendra Malik for Computer Science 188 (Introduction to Artificial Intelligence), an upper-division undergraduate class. Prepared and led weekly discussion sections; wrote and graded exam questions. Substituted for professor in lecture on two occasions (Topics: Planning, Speech Recognition). Prepared lecture notes on Markov models to augment class readings; these notes (available on the web) have been in use at least five universities.

Guest Lectures U.C. Berkeley

Electrical Engineering and Computer Science 225d, Audio Signal Processing in Humans and Machines, Spring 1997 (1 lecture), Spring 1999 (3 lectures). Topics: Introduction to linguistic categories for automatic speech recognition; Using statistical models of language in speech recognition and understanding systems.

Linguistics 158, Computational Linguistics, Spring 1997. Topic: Introduction to Speech Recognition

University of Pennsylvania 1/93–5/93

Volunteer teaching assistant for Prof. Sabine Iatridou. Held one-on-one tutorials for students in Linguistics 150, Introduction to Transformational Grammar. Substituted for professor in lecture on one occasion (Topic: Natural Language Processing).

Professional Activities

Publicity Chair, IEEE Workshop on Automatic Speech Recognition and Understanding, St. Thomas, U.S. Virgin Islands, November 2003.

Workshop organizer, International Speech Communication Association (ISCA) Tutorial and Research Workshop on Pronunciation Modeling and Lexicon Adaptation for Spoken Language Technology, Estes Park, Colorado, September 2002.

Referee for the journals *Speech Communication, Computer, Speech, and Language, IEEE Transactions on Speech and Audio Processing*, and *IEEE Systems, Man & Cybernetics*.

Proposal referee for the National Science Foundation and Netherlands Council of the Humanities.

Reviewer/scientific committee for the following conferences: Annual Meeting of the Association for Computational Linguistics (ACL) (1997, 2000, 2001), International Conference on Acoustics, Speech, and Signal Processing (ICASSP) (2001), IEEE Automatic Speech Recognition and Understanding Workshop (2001), Human Language Technologies Conference (2002, 2003).

U.C. Berkeley Computer Science Admissions Committee, Spring 1997, Spring 1998.

Member of IEEE, Association for Computational Linguistics, and International Speech Communication Association.

Member of University of Pennsylvania Secondary School Committee, 1998-2000.

Other Professional Experience

Linguistic Data Consortium, University of Pennsylvania Staff programmer 5/93–8/93
Created programs for indexing and extracting sound samples from the Switchboard corpus. Adapted scripts for aligning parallel multi-lingual text data derived from transcripts of United Nations meetings.

School of Engineering and Applied Science,
University of Pennsylvania Assistant systems administrator 5/91–8/93
Helped administer UNIX servers for school computer facility (Computer and Educational Technology Services). Wrote automatic system administration tool to equalize usage on multi-user system.

School of Engineering and Applied Science,
University of Pennsylvania Laboratory administrator 8/90–5/91
Installed and maintained computers and networks in a new Macintosh laboratory.

References

1. Prof. Nelson Morgan
EECS Department, U.C. Berkeley and
International Computer Science Institute
1947 Center Street, Suite 600
Berkeley, CA 94704-1198
(510) 666-2931
`morgan@icsi.berkeley.edu`

2. Dr. Joseph Olive
Bell Labs, Lucent Technologies
600 Mountain Ave., Room 2D-418
Murray Hill, NJ 07974
(908) 582-3110
`jpo@research.bell-labs.com`

3. Prof. Daniel Jurafsky
Department of Linguistics
University of Colorado
Boulder, CO 80309-0295
(303) 492-1300
`jurafsky@colorado.edu`

4. Prof. Mari Ostendorf
Box 352500, Dept. of Electrical Engineering
University of Washington
Seattle, WA 98195-2500
(206) 221-5748
`mo@ee.washington.edu`

Publications

Dissertation

J. E. Fosler-Lussier. "Dynamic Pronunciation Models for Automatic Speech Recognition," Ph.D. thesis, University of California, Berkeley, 1999. Reprinted as International Computer Science Institute technical report TR-99-015.

Patents

E. Fosler-Lussier, C. H. Lee, A. Pargellis, and A. Potamianos, "System and method for measuring domain independence of semantic classes," filed 2002 (in review).

E. Ammicht, E. Fosler-Lussier, and A. Potamianos, "System and method for representing and resolving ambiguity in spoken dialogue systems," filed 2002 (in review).

Book Chapters

E. Fosler-Lussier. "A Tutorial on Pronunciation Modeling for Large Vocabulary Speech Recognition," invited chapter for ELSNET book *Text and Speech Triggered Information Access*, Springer Verlag, to appear.

E. Fosler-Lussier. "Linguistic Categories for Speech Recognition," in N. Morgan and B. Gold, *Speech and Audio Signal Processing: Processing and Perception of Speech and Music*, John Wiley and Sons, Inc., 1999, Chapter 23.

Proceedings

W. Byrne, E. Fosler-Lussier, and D. Jurafsky, eds. *Proceedings of the ISCA Tutorial and Research Workshop on Pronunciation Modeling and Lexicon Adaptation in Spoken Language Technology*, Estes Park, Colorado, 2002.

Refereed Articles

A. Bell, D. Jurafsky, E. Fosler-Lussier, C. Girand, M. Gregory, and D. Gildea. "Form variation of English function words in conversation," *Journal of the Acoustical Society of America*, to appear, 2003.

A. J. Robinson, G. D. Cook, D. P. W. Ellis, E. Fosler-Lussier, S. J. Renals, and D. A. G. Williams. "Connectionist Speech Recognition of Broadcast News," *Speech Communication* 37:1-2, DARPA Broadcast News special issue, pp. 27–45, 2002.

E. Fosler-Lussier and N. Morgan. "Effects of Speaking Rate and Word Predictability on Conversational Pronunciations," *Speech Communication* 29:2-4, pp. 137–158, 1999.

Work in progress

E. Fosler-Lussier, I. Amdal, and H.-K. J. Kuo. "A Framework for Predicting Speech Recognition Errors," invited paper for Speech Communication issue on Pronunciation Modeling and Lexicon Adaptation.

E. Ammicht, E. Fosler-Lussier, and A. Potamianos. "Pragmatic Understanding: Combining Semantics and Pragmatics in the Bell Labs Communicator System."

A. Pargellis, E. Fosler-Lussier, C.-H. Lee, A. Potamianos, and A. Tsai. "Automatic Induction of Semantic Classes for Spoken Dialogue Systems."

Refereed Conferences and Workshops

Speech Recognition: Pronunciation modeling

E. Fosler-Lussier, I. Amdal, and H.-K. J. Kuo. "On the road to improved lexical confusability metrics," ISCA Tutorial and Research Workshop on Pronunciation Modeling and Lexicon Adaptation (PMLA-2002), Estes Park, Colorado, 2002.

M. Wester and E. Fosler-Lussier. "A comparison of data-derived and knowledge-based modeling of pronunciation variation," International Conference on Speech and Language Processing (ICSLP-2000), Beijing, China, 2000.

E. Fosler-Lussier. "Contextual word and syllable pronunciation models," 1999 International Workshop on Automatic Speech Recognition and Understanding (ASRU '99), Keystone, Colorado, 1999.

E. Fosler-Lussier. "Multi-level decision trees for static and dynamic pronunciation models," Sixth European Conference on Speech Communication and Technology (Eurospeech '99), Budapest, Hungary, 1999.

E. Fosler-Lussier, S. Greenberg, and N. Morgan. "Incorporating Contextual Phonetics into Automatic Speech Recognition," Invited paper for the symposium "The Phonetics of Spontaneous Speech," International Congress of Phonetic Sciences (ICPhS '99), San Francisco, California, 1999.

E. Fosler-Lussier and G. Williams, "Not just what, but also when: Guided automatic pronunciation modeling for Broadcast News," DARPA Hub4E Evaluation Workshop, 1999.

E. Fosler, M. Weintraub, S. Wegmann, Y.-H. Kao, S. Khudanpur, C. Galles, M. Saraclar. "Automatic Learning of Word Pronunciations from Data," Invited paper, International Conference on Spoken Language Processing (ICSLP-96), Philadelphia, Pennsylvania, 1996.

G. Tajchman, E. Fosler, and D. Jurafsky. "Building Multiple Pronunciation Models for Novel Words using Exploratory Computational Phonology," Fourth European Conference on Speech Communication and Technology (Eurospeech '95), Madrid, Spain, 1995.

G. Tajchman, D. Jurafsky, and E. Fosler. "Learning Phonological Rule Probabilities from Speech Corpora with Exploratory Computational Phonology," Association for Computational Linguistics (ACL-95), Boston, Massachusetts, 1995.

Speech Recognition: Language modeling

H.-K. Kuo, E. Fosler-Lussier, H. Jiang, and C.-H. Lee. "Discriminative Training of Language Models for Speech Recognition," International Conference on Acoustics, Speech, and Signal Processing, Orlando, Florida, 2002.

R. Argiles Solsona, E. Fosler-Lussier, H.-K. Kuo, A. Potamianos and I. Zitouni. "Adaptive Language Models for Spoken Dialogue Systems," International Conference on Acoustics, Speech, and Signal Processing, Orlando, Florida, 2002.

E. Fosler-Lussier and H.-K. J. Kuo. "Using Semantic Class Information for Rapid Development of Language Models within ASR Dialogue Systems," International Conference on Acoustic, Speech, and Signal Processing (ICASSP-2001), Salt Lake City, Utah, 2001.

D. Jurafsky, C. Wooters, J. Segal, A. Stolcke, E. Fosler, G. Tajchman, and N. Morgan. "Using A Stochastic Context-Free Grammar as a Language Model for Speech Recognition," International Conference on Acoustic, Speech, and Signal Processing (ICASSP-95), Detroit, Michigan, 1995.

Speech Recognition: Systems

G. Cook, J. Christie, D. Ellis, E. Fosler-Lussier, Y. Gotoh, B. Kingsbury, N. Morgan, S. Renals, T. Robinson, and G. Williams, "The SPRACH System for the Transcription of Broadcast News," DARPA Hub4E Evaluation Workshop, 1999.

N. Morgan, D. Ellis, E. Fosler-Lussier, A. Janin, and B. Kingsbury, "Reducing errors by increasing the error rate: MLP Acoustic Modeling for Broadcast News Transcription," DARPA Hub4E Evaluation Workshop, 1999.

D. Jurafsky, C. Wooters, G. Tajchman, J. Segal, A. Stolcke, E. Fosler, and N. Morgan. "The Berkeley Restaurant Project," International Conference on Spoken Language Processing (ICSLP-94), Yokohama, Japan, 1994.

Speaking rate estimation / Recognition of fast speech

N. Morgan and E. Fosler-Lussier. "Combining Multiple Estimators of Speaking Rate," International Conference on Acoustic, Speech, and Signal Processing (ICASSP-98), Seattle, Washington, 1998.

N. Morgan, E. Fosler, and N. Mirghafori. "Speech Recognition using On-line Estimation of Speaking Rate," Fifth European Conference on Speech Communication and Technology (Eurospeech '97), Rhodes, Greece, 1997.

N. Mirghafori, E. Fosler, and N. Morgan. "Towards Robustness to Fast Speech in ASR," International Conference on Acoustic, Speech, and Signal Processing (ICASSP-96), Atlanta, Georgia, 1996.

N. Mirghafori, E. Fosler, and N. Morgan. "Why Is ASR Harder For Fast Speech and What Can We Do About It?" Proceedings of the IEEE Signal Processing Society Snowbird95 Workshop, Snowbird, Utah, 1995.

N. Mirghafori, E. Fosler, and N. Morgan. "Fast Speakers in Large Vocabulary Continuous Speech Recognition: Analysis & Antidotes," Fourth European Conference on Speech Communication and Technology (Eurospeech '95), Madrid, Spain, 1995.

Spoken Dialogue Systems

S. Lee, E. Ammicht, E. Fosler-Lussier, H.-K. Kuo, and A. Potamianos, "Spoken Dialogue Evaluation for the Bell Labs Communicator System," Human Language Technologies Conference, San Diego, California, 2002.

M. Galley, E. Fosler-Lussier, and A. Potamianos. "Hybrid Natural Language Generation for Spoken Dialogue Systems," Seventh European Conference on Speech Communication and Technology (Eurospeech '01), Aalborg, Denmark, 2001.

Spoken Natural Language Understanding

A. Pargellis, E. Fosler-Lussier, and A. Tsai. "Using part-of-speech tags, context thresholding, and trigram contexts to improve the auto-induction of semantic classes," International Conference on Spoken Language Processing (ICSLP-2002), Denver, Colorado, 2002.

A. Pargellis, E. Fosler-Lussier, A. Potamianos, and C.-H. Lee. "A comparison of four metrics for auto-inducing semantic classes," IEEE Automatics Speech Recognition and Understanding Workshop (ASRU-2001), Trento, Italy, 2001.

E. Ammicht, A. Potamianos, and E. Fosler-Lussier. "Ambiguity Representation and Resolution in Spoken Dialogue Systems," Seventh European Conference on Speech Communication and Technology (Eurospeech '01), Aalborg, Denmark, 2001.

A. Pargellis, E. Fosler-Lussier, A. Potamianos, and C.-H. Lee. "Metrics for Measuring Domain Independence of Semantic Classes," Seventh European Conference on Speech Communication and Technology (Eurospeech '01), Aalborg, Denmark, 2001.

Natural Language Call Routing / Topic Verification

H.-K. Kuo, C.-H. Lee, I. Zitouni, and E. Fosler-Lussier. "Minimum Verification Error Training for Topic Verification," International Conference on Acoustics, Speech, and Signal Processing (ICASSP-2003), Hong Kong, China, 2003.

H.-K. Kuo, I. Zitouni, E. Fosler-Lussier, E. Ammicht, and C.-H. Lee. "Discriminative Training for Call Classification and Routing," International Conference on Spoken Language Processing (ICSLP-2002), Denver, Colorado, 2002.

P. Durston, M. Farrell, D. Attwater, J. Allen, H.-K. J. Kuo, M. Afify, E. Fosler-Lussier, C.-H. Lee. "OASIS Natural Language Call Steering Trial," Seventh European Conference on Speech Communication and Technology (Eurospeech '01), Aalborg, Denmark, 2001.

Linguistic Studies

C. Shih, G. Kochanski, and E. Fosler-Lussier. "Implications of Prosody Modeling for Prosody Recognition," ICSA Research and Tutorial Workshop on Prosody in Automatic Speech Recognition (PROSODY-2001), Red Bank, New Jersey, 2001.

S. Greenberg and E. Fosler-Lussier. "The uninvited guest: Information's role in guiding the production of spontaneous speech," in the Proceedings of the Crest Workshop on Models of Speech Production: Motor Planning and Articulatory Modelling, Kloster Seeon, Germany, 2000.

A. Bell, D. Jurafsky, E. Fosler-Lussier, C. Girand and D. Gildea. "Reduction of English Function Words In Disfluent Contexts," International Congress of Phonetic Sciences (ICPhS '99), San Francisco, California, 1999.

M. L. Gregory, W. D. Raymond, A. Bell, E. Fosler-Lussier, and D. Jurafsky. "The effects of collocational strength and contextual predictability in lexical production," Proceedings of the Chicago Linguistics Society, v. 35, 1999.

D. Jurafsky, A. Bell, E. Fosler-Lussier, C. Girand, and W. Raymond. "Reduction of English function words in Switchboard," International Conference on Spoken Language Processing (ICSLP-98), Sydney, Australia, 1998.

E. Fosler-Lussier and N. Morgan. "Effects of Speaking Rate and Word Predictability on Conversational Pronunciations," ESCA Tutorial and Research Workshop on Modeling Pronunciation Variation for Automatic Speech Recognition, Kerkrade, Netherlands, 1998.

E. Fosler. "On Reversing the Generation Process in Optimality Theory," Proceedings of the Association for Computational Linguistics, Santa Cruz, California, 1996.

Technical Reports, Workshop Presentations, and Magazine Articles

I. Amdal and E. Fosler-Lussier, "Pronunciation Modeling in Speech Technology," *Telektronikk*, to appear.

E. Fosler-Lussier. "Markov Models and Hidden Markov Models: A Brief Tutorial," International Computer Science Institute Technical Report TR-98-041, 1998. CS 188 (Introduction to Artificial Intelligence), Fall 1995 Teaching Materials.

E. Fosler. "Automatic learning of a model for word pronunciations: Status report," Conversational Speech Recognition Workshop: DARPA Hub 5E Evaluation, Baltimore, Maryland, 1997.

M. Weintraub, E. Fosler, C. Galles, Y.-H. Kao, S. Khudanpur, M. Saraclar, and S. Wegmann, "WS96 Project Report: Automatic Learning of Word Pronunciation from Data," in F. Jelinek, editor, 1996 Large Vocabulary Continuous Speech Recognition Summer Research Workshop Technical Reports, Chapter 3. Center for Language and Speech Processing, Johns Hopkins University, Research Notes No. 24, 1997.

N. Mirghafori, E. Fosler, and N. Morgan. "Making Automatic Speech Recognition More Robust to Fast Speech," International Computer Science Institute Technical Report TR-95-067, Berkeley, California, 1995.