JOHNS HOPKINS
WHITING SCHOOL
of ENGINEERING

The JHU Summer Workshops: History and 2014 Goals

Sanjeev Khudanpur and Jack Godfrey
Center for Language and Speech Processing
The Johns Hopkins University
July 7, 2014

How it all began …

- The SCAMP tradition at IDA-CCR (Princeton)
  - Arthur House, Tom Crystal, et al
- 1993-1994: The first two Rutgers workshops
- CLSP was established at Johns Hopkins in 1994
  - Fred Jelinek was named its first Director
- 1995: The first "JHU" Summer Workshop
  - Best error rates on English SWBD … ca 55%
  - Spanish telephone speech … 80% word error rate

What was LM95 like?

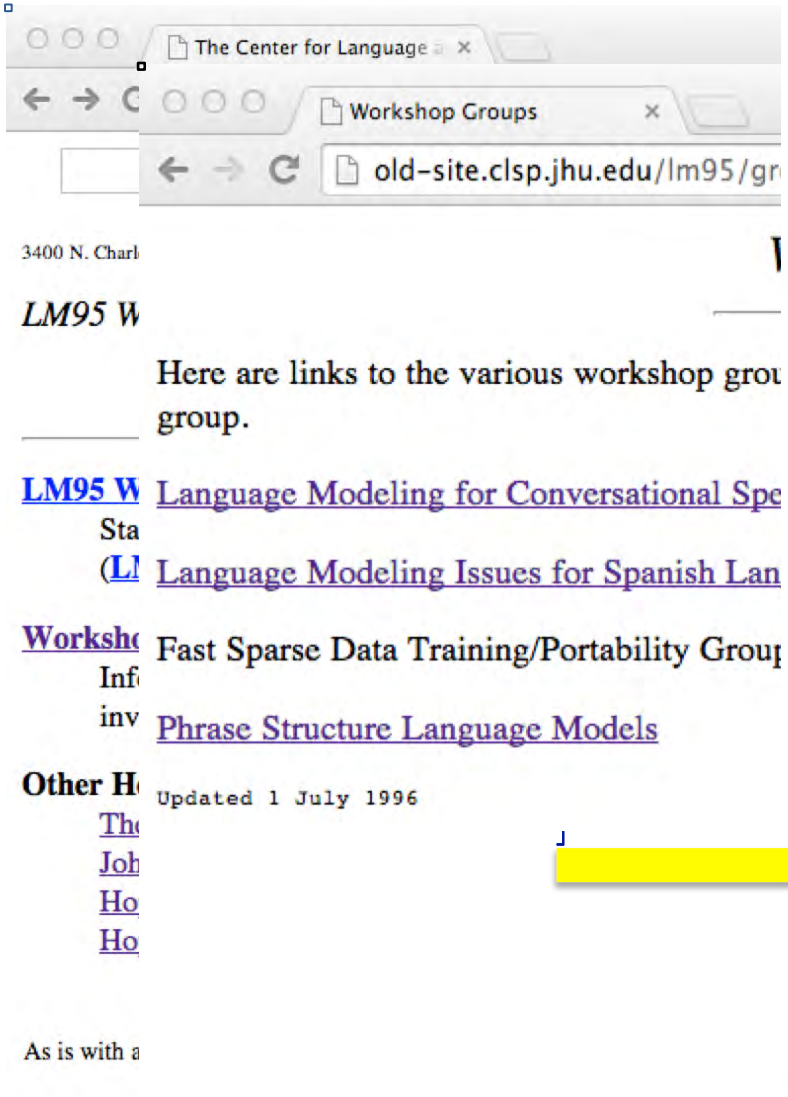

Questions

- What is the best WER attainable with the existing lattices
- Will we do better if we had more language training data?
- Can we work with Nbest lists? How much do we lose?
- Alternatives progress measures?
- Conversational speech has no clear linguistic boundaries. Is this a problem?
- Are some sentences more difficult than others? (long ones? dysfluent ones?) Will new ideas work better on some sentences than on others
- Can we predict when our output is correct? What are the errors most correlated with?
- Who is to blame for errors?
- Is trigram the best conventional model?

Two Decades of Summer Adventures

- **4 Teams in 1996**
  - Pronunciation Modeling (Mitch Weintraub)
  - Dependency Language Modeling (Andreas Stolcke)
  - Hidden Speaking Modes (Mari Ostendorf)
  - Speech Data Modeling (Jordan Cohen)
- **4 Teams in 1997**
  - Multi-scale Acoustic Processing (Andreas Andreou)
  - Syllable-based Modeling (George Doddington)
  - Data-driven Pronunciation Modeling (Michael Riley)
  - Discourse Modeling (Dan Jurafsky)
- **3 Teams in 1998**
  - Segmental Models of Speech (John Bridle and Li Deng)
  - Rapid Adaptation to New Speakers (Vassilis Digalakis)
  - Core NLP for Multiple Languages (Jan Hajic)

Two Decades of Summer Adventures

Two Decades of Summer Adventures

- **4 Teams in 1999**
  - Normalization of Non-standard Words (Richard Sproat)
  - Statistical Machine Translation (Kevin Knight)
  - Topic Based Novelty Detection (James Allen)
  - Towards Language Independent Acoustic Modeling (William Byrne)
- **4 Teams in 2000**
  - Reading Comprehension (Marc Light)
  - Cross-lingual Speech Retrieval (Helen Meng)
  - Pronunciation Modeling for Mandarin Speech (Pascale Fung)
  - Audio-Visual Speeech Recognition (Chalapathy Neti)
- **2 Teams in 2001**
  - Graphical Models for Speech Recognition (Jeff Bilmes and Geoff Zweig)
  - Automatic Summarization of Multiple Documents (Dragomir Radev)

Two Decades of Summer Adventures

Two Decades of Summer Adventures

- **4 Teams in 2002**
  - Generation for Machine Translation (Jan Hajic)
  - Weakly Supervised Learning of Parsers (Mark Steedman)
  - SuperSID: Speaker Recognition (Doug Reynolds)
  - Dialectal Arabic Speech Recognition (Katrin Kirchhoff)
- **3 Teams in 2003**
  - Syntax for Statistical Machine Translation (Franz J Och)
  - Semantic Analysis over Sparse Data (Louise Guthrie)
  - ~~Dialectal Chinese Speech Recognition (Thomas Zheng)~~
  - Confidence Estimation for NLP Applications (George Foster)
- **3 Teams in 2004**
  - Dialectal Chinese Speech Recognition (Richard Sproat)
  - Landmark-Based Speech Recognition (Mark Hasegawa-Johnson)
  - Joint Visual-Text Modeling for Video Retrieval (Giri Iyengar)

Two Decades of Summer Adventures

Two Decades of Summer Adventures

- **3 Teams in 2005**
  - Parsing Arabic Dialects (Owen Rambow)
  - Speech Parsing (Mary Harper)
  - Machine Translation by Parsing (Dan Melamed)
- **2 Teams in 2006**
  - Articulatory Features for Speech Recognition (Karen Livescu)
  - Open Source Toolkit for SMT (Philipp Koehn)
- **2 Teams in 2007**
  - Named Entity Disambiguation (Massimo Poesio)
  - Recovery from Model Inconsistency (Hynek Hermansky)
- **3 Teams in 2008**
  - Multilingual Spoken Term Detection (Richard Sproat)
  - Robust Speaker Recognition (Lukas Burget)
  - Vocal Aging Explained by Vocal Tract Modeling (Elmar Noeth)

Two Decades of Summer Adventures

- **3 Teams in 2009**
  - Scalable CCG Parsing (Stephen Clark)
  - Low Resource Speech Recognition (Dan Povey)
  - Unsupervised Acquisition of Lexical Knowledge (Dekang Lin)
- **3 Teams in 2010**
  - Speech Recognition with Segmental CRF (Geoff Zweig)
  - Inducing Synchronous Grammars for MT (Phil Blunsom)
  - Localizing Objects and Actions in Video (Jan Newmann)
- **3 Teams in 2011**
  - Synthesizing Emotional Speech (Alan Black)
  - Confusion-Based Language Models for MT and ASR (Brian Roark)
  - Learning from Visually Descriptive Text (Alex and Tamara Berg)
- **3 Teams in 2012**
  - Domain Adaptation for Machine Translation (Hal Daume III)
  - Detailed Understanding of Objects and Scenes (Andrea Vedaldi)
  - Complementary Evaluation Measures for ASR (Gerald Penn)

ummer Adventures

Figure 1. Our system automatically generates the following descriptive text for this example image: *"This picture shows one person, one grass, one chair, and one potted plant. The person is near the green grass, and in the chair. The green grass is by the chair, and near the potted plant."*

Two Decades of Summer Adventures

- **1 Team in 2013**
  - Speaker and Language Recognition (Self Organized)
- **3 Teams in 2014**
  - ASR Machines that Know When They Do Not Know (Hynek Hermansky)
  - Cross-Lingual Abstract Meaning Representations (Martha Palmer)
  - Probabilistic Representations of Linguistic Meaning (Jason Eisner and Ben Van Durme)

Workshop Sponsorship

- US National Science Foundation (NSF)
  - Office of International and Integrative Activities (IIA)
    - Partnerships for International Research and Education (PIRE)
- Fred Jelinek et al had received a PIRE Award in 2005
  - Collaborative Investigation of Meaning Representations in Language Understanding
    - Johns Hopkins, Brown, Charles and Saarland Universities
    - After Fred passed away, NSF generously allowed it to continue
    - Added Macquarie, VUT Brno and Colorado as new partners
- This PIRE award is the primary sponsor of WS 2104
  - Supplemental support from Johns Hopkins and in-kind support from several participants' home institutions

IIA-0530118: Project Goals

- This Partnership for International Research and Education links senior and junior researchers from Johns Hopkins University and Brown University with counterparts from Charles University in the Czech Republic and Saarland University in Germany.
- The international team will investigate formal representations of linguistic meaning for use in automatic speech recognition (ASR) and machine translation (MT) systems. Their goal is to augment current systems by applying a variety of formal models for deep syntactic/semantic representation.
- The projects complementary education component involves introducing participating U.S. graduate students to European-developed linguistic formalisms and training them to apply those formalisms to problems in natural language processing.

The Planning Began in March 2013

- Streamlined the "International" research plan of each PhD student
  - Speech recognition research at VUT Brno (Harish and Vjiay)
  - Machine Translation research at Charles University (Wei-Te)
  - Semantics research at Saarland University (Rachel)
  - Meaning representation research at Macquarie (Halley and Nick)
- Developed an integrative education effort to foster an esprit de corps
  - Weekly team meetings in the Fall 2013 semester
- Developed plans for a culminating integrative event for the project
  - A month-long residential workshop in Prague
- Reached out to 28 additional US researchers to join the workshop
  - 15 senior researchers, 13 PhD students
- Invited 10+ additional non-US researchers to the workshop
- Each team held a 2-day planning meeting (Feb 20-21 & April 19-20)

CLAMR
Cross-Lingual Abstract Meaning Representations

Martha Palmer	**University of Colorado**		
Ondrej Bojar	Charles University	Wei-Te Chen	University of Colorado
David Chiang	USC/ISI	Ondrej Dusek	Charles University
Frank Drewes	Umea University	Jeffrey Flanigan	Carnegie Mellon University
Daniel Gildea	University of Rochester	Tim O'Gorman	University of Colorado
Jan Hajic	Charles University	Xiaochang Peng	University of Rochester
Adam Lopez	Johns Hopkins University	Martin Popel	Charles University
Giorgio Satta	University of Padua	Adiyta Renduchintala	Johns Hopkins University
Zdenka Uresova	Charles University	Naomi Saphra	Johns Hopkins University
Nianwen Xue	Brandeis University	Chuan Wang	Brandeis University
		Yuchen Zhang	Brandeis University

PRELIM
Probabilistic Representations of Linguistic Meaning

Jason Eisner and Benjamin Van Durme	**Johns Hopkins University**

Oren Etzioni	Allen Institute for AI	Nicholas Andrews	Johns Hopkins University
Craig Harman	Johns Hopkins University	Frank Ferraro	Johns Hopkins University
Shalom Lappin	King's College London	Drew Reisinger	Johns Hopkins University
Staffan Larsson	University of Gothenburg	Halley Orshan	Johns Hopkins University
Dan Lassiter	Stanford University	Rachel Rudinger	Johns Hopkins University
Percy Liang	Stanford University		
David McAllester	Toyota Technological Institute		
James Pustejovsky	Brandeis University		
Kyle Rawlins	Johns Hopkins University		

ASR
Machines That Know When They Do Not Know

Hynek Hermansky	**Johns Hopkins University**		
Lukas Burget	VUT Brno	Matthew Maciejewski	Carnegie Mellon University
Jordan Cohen	Spelamode Consulting	Sri Harish Mallidi	Johns Hopkins University
Naomi Feldman	University of Maryland	Anjali Menon	Carnegie Mellon University
Tetsuji Ogawa	Waseda University	Vijayaditya Peddinti	Johns Hopkins University
Richard Rose	McGill University	Matthew Wiesner	McGill University
Richard Stern	Carnegie Mellon University	Karel Vesely	VUT Brno

Workshop Agenda (Week 1)

	Mon Jul 7	Tue Jul 8	Wed Jul 9	Thu Jul 10	Fri Jul 11
09:00	Workshop Opening Day Presentations				
11:00		PRLM Lecture	PRLM Lecture	PRLM Lecture	PRLM Lecture
12:30	Lunch Break				
14:00	Parallel Team Meetings				
15:00	Town Hall	Town Hall	Town Hall	Town Hall	Town Hall
15:30	PRLM Lecture	PRLM Lecture	PRLM Lecture	PRLM Lecture	PRLM Lecture
17:00	Steering Committee				
18:00	Welcome Dinner			Group Dinner	

Workshop Agenda (Week 2 & 3)

	Monday	Tuesday	Wednesday	Thursday	Friday
10:00	Lecture	Lecture	Lecture	Lecture	Lecture
14:00	Optional Lecture	Optional Lecture	Optional Lecture	Optional Lecture	Optional Lecture
15:00	Town Hall	Town Hall	Town Hall	Town Hall	Town Hall
17:00	Steering Committee				
18:00				Group Dinner	

Workshop Agenda (Week 4)

	Mon Jul 28	Tue Jul 29	Wed Jul 30	Thursday	Fri Aug 1
09:00					Workshop Closing Day Presentations
10:00	Lecture	Lecture	Lecture		
14:00	Optional Lecture	Optional Lecture	Optional Lecture		Closing Day Presentations (Cont'd)
15:00	Town Hall	Town Hall	Town Hall	Town Hall	
17:00	Steering Committee				
18:00			Group Dinner		Farewell Dinner

Fred would have loved to be among us today!