

ISCA DL Travel Report

Prof. Jeff A. Bilmes
University of Washington, Seattle
<http://melodi.ee.washington.edu/~bilmes>

August 1st, 2013

From January 5 to January 22, 2013, I traveled to India as an ISCA Distinguished Lecturer. I visited the cities of Bangalore and Bombay and gave two lectures there. The cost for international as well as domestic flights will be covered by ISCA (I have, as of this writing, not yet turned in the receipts), and accommodation and food were covered by my local hosts.

India was chosen as a location for my set of ISCA talks for a number of reasons. First, as is obvious, there is an enormous amount of talent there, and encouraging students from the academic community in India to go into speech research could have real benefits to the community as a whole. Secondly, India is very populous, and its population is projected to exceed even China's population in the not-to-distant future <http://www.livemint.com/Politics/JD6HSg41JmbarVhcA03N2M/Indias-population-to-equal-Chinas-in-the-next-15-years.html>. It therefore seemed like a fantastic opportunity to establish further connection in order to foster further growth in the speech/language communities there.

1 ISCA DL in Bangalore, India



Figure 1: Photos of Prof. Jeffrey A Bilmes from the IEEE CONECC 2013 conference <http://conecct.iieebangalore.org/2013/> that was held in Bangalore India

My first talk was part of the IEEE CONECC 2013 conference <http://conecct.iieebangalore.org/2013/> which is held in Bangalore India, and is run by faculty at the Indian Institute of Science (IISc). The organizers for the conference, and also my hosts for the ISCA lectureship were Professor K.V.S. Hari <http://www.ece.iisc.ernet.in/~hari/> and Professor Thippur V. Sreenivas http://www.ece.iisc.ernet.in/~tvsree/T_V_Sreenivas.html who also runs a speech and audio group at the university there <https://sites.google.com/site/sagiisc/people/faculty>.

Several photos from the event are included in Figure1.

The talk that I gave was on a problem that we all face in the speech recognition community, namely the proliferation of training data. Many of the worlds' commercial speech labs have utilized such speech "big data" by somewhat indiscriminately adding it to the training set and then training on all of it. The hypothesis of the talk (and the set of methods that are being researched) is that some (if not much) of it is redundant,

and that by carefully selecting a subset of the training data, one can do just as well but at a significantly reduced cost. This problem is particularly acute these days with deep neural network models beginning to be so widely used. While deep models are found to be very accurate, they sometimes require weeks if not months of training time, and specialized implementations on specialized hardware (GPUs), something that often puts them out of reach of at least the academic environment. A goal of this research is to reduce data sizes without losing the benefit of big data, so that all research labs can partake in state-of-the-art speech research and test out advanced models that, otherwise, would be computationally prohibitive.

The abstract of the talk is given below.

Turning Big Data into Core Data

Jeffrey Bilmes

University of Washington, Seattle

<http://melodi.ee.washington.edu/~bilmes>

The wealth of available data has been a problem for human consumers of information for decades. At this point, the explosion of available data is a problem even for computer consumers of information. In this talk, we will discuss how submodular function optimization can address such issues. For example, the generic problem of joint active/semi-supervised learning, and in particular the problems of document summarization and data subset selection in speech recognition, can be successfully solved using submodular optimization. Given the computational expense of recent training methods for state-of-the-art speech recognition systems (such as deep models), training data subset selection is an even more important problem than in the past. We will also survey a number of recent problems in machine learning where submodular functions are appropriate. The talk will include sufficient background to make it accessible to everyone.

The attendance was around 250-300 people from academia (IISc in particular), and local industry in the Bangalore area (although there were attendees from all over India).

While in Bangalore, I also visited and spent some time at Professor Thippur V. Sreenivas's lab, and met many of his students. We talked about a number of interesting projects that were going on there, and I gave them some advice (and a sense of excitement about the speech community) as I could. I felt that some of the students there were quite good, and I would enjoy having them as my own students in my lab back in Seattle.

2 ISCA DL in Bombay, India

The IIT (Indian Institute of Technology) university system in India is known by many to be among the very best set of universities in India, and IIT-Bombay in particular is said by many to have become the very best of all of the IITs. It is for this reason that I chose to visit there with particular focus on the goals mentioned in the introduction of this document.

My next visit was to IIT-Bombay and my host at IIT was Ganesh Ramakrishnan <http://www.cse.iitb.ac.in/~ganesh/> of IIT-Bombay. He is himself in the computer science department there, and does not work on speech recognition per se, but does have an interest in the area.

The title/abstract (and goals) of my talk was the same as the IISc/IEEE CONECCT conference talk. The attendance was about 50-60 people, most of whom were students at IIT-Bombay.

I also met with a large number of faculty at IIT-Bombay. It should be noted that IIT-Bombay does not have a speech lab, although there are people there working on more general machine learning. It was a goal of mine to interest them in some speech applications, if anything as being seen as a fantastic application for some of their machine learning algorithms.

3 Discussion

Overall, my trip to India was very enjoyable and I think I had some level of success in my goals — many of the students I spoke to (in both locations) had never been out of the country, and they were quite eager to discuss with me both my research and life outside of the country.

It should also be noted that so far I've only given one series of talks (i.e., one trip) as part of the ISCA Lectureship — for a number of reasons (partly due to very full teaching schedule and partly due to having a relatively new baby), there were some scheduling conflicts so that I have yet not set up a second trip. I still hope to do so if it is agreeable to ISCA.

Lastly, I would like to thank my hosts in India (Profs. K.V.S. Hari, Thippur V. Sreenivas, and Ganesh Ramakrishnan) for being extraordinarily hospitable during my stays in Bangalore and Bombay. Also, I would like to thank Professor Satoshi Nakamura, and all of ISCA, for their fantastic work, and for establishing and maintaining this lectureship — I think the ISCA lectureship is an extremely beneficial thing to exist.