

- Acoustic modeling expert
- Working with collaborators on open-source “Kaldi” speech recognition toolkit (release planned in March).
- Will support standard techniques but also SGMMs (see right pane).
- Tight integration with OpenFst (makes it easier to build hybrid word/phone models and the like).
- Freshly written, clean software with some original aspects (helpful for combination).

- Particularly well suited to small-data/multilingual.
- A factored form of GMMs (it expands to a big full-covariance GMM in each state).
- Most of the parameters are global (not tied to a state), so for multilingual you can share parameters without sharing phones across languages.
- We got impressive improvements in the “low-resource” case.
- Even more impressive improvements in “low-resource+multilingual” (e.g. 1 hour of in-language training data).
- Has also been used by IBM in GALE evaluation (big data); improvements were smaller but system still a bit better than baseline.
- Development continuing on this model (hope it will get decisively better than conventional models).

- Do not require funding
- Willing to help anyone who wants to use Kaldi/SGMMs
- Will consider “officially” joining a team.
- Regardless, for more hands-on help, talk to Brno University of Technology (collaborating with me on Kaldi development; can build the systems).

Daniel Povey
Researcher
Microsoft
dpovey@microsoft.com
(425) 706-2667
<http://research.microsoft.com/en-us/people/dpovey/>