

# PETER BELL

## Curriculum Vitae

### Post-graduation career

- 2010 – Senior Research Associate, School of Informatics, University of Edinburgh  
70% FTE Sept 2012 – Apr 2018  
80% FTE Apr 2018 –
- 2012 – Senior Speech Scientist, Quorate Technology Ltd, 20% FTE

### Education

- 2005 – 2010 PhD in Automatic Speech Recognition, University of Edinburgh  
Thesis: *Full Covariance Modelling for Speech Recognition*
- 2004 – 2005 MPhil, Computer Speech, Text and Internet Technology, University of Cambridge  
(grade A)
- 1999 – 2002 BA, Mathematics, University of Cambridge (class 2:i).

### Research interests

Automatic speech recognition (ASR): domain adaptation and regularisation for neural networks; low resource systems; end-to-end training; multi-task learning; lightly supervised alignment; applications to broadcast media

### Research projects

- 2018 – SpeechWave, EPSRC responsive-mode grant; Co-investigator
- 2017 – SCRIPTS, IARPA MATERIAL programme; Co-investigator and manager of Edinburgh's ASR team
- 2016 – 2019 SUMMA, H2010; leader of the ASR team
- 2011 – 2016 Natural Speech Technology, EPSRC programme grant; responsible for structuring diverse data and wide domain coverage themes
- 2012 – 2015 EU-BRIDGE, FP17; responsible for ASR on lecture data, coordinating Edinburgh's entries to the IWSLT evaluation campaign

### Research grants awarded

- 2019 – Co-investigator, *Adapting end-to-end speech recognition systems*, Samsung, £137,400 (with Steve Renals)
- 2018 – 2021 Co-investigator, *SpeechWave*, EPSRC, £835,000 (with Steve Renals)
- 2017 – 2019 Co-investigator, *Systems for Cross-language Information Processing, Translation and Summarisation*, IARPA MATERIAL Programme, £614,900 (With Steve Renals and Kenneth Heafield)
- 2016 – 2018 PhD studentship, DataLab and Ericsson Broadcast and Media Services, £102,000 (with Steve Renals)
- 2015 – 2018 PhD studentship, Bloomberg Inc, £142,000 (with Steve Renals)
- 2014 Invited research visit to NICT, Japan, £5,000
- 2012 – 2013 Knowledge Transfer Partnership, UK Technology Strategy Board and France Telecom, £60,000

### **Supervision of research and taught postgraduate students**

- 2018 – Chao Luu, PhD co-supervisor
- 2018 Host of visiting PhD students Do Truong and Bertrand Higy
- 2017 – Shucong Zhang, PhD co-supervisor
- 2016 – Joanna Rownicka, PhD co-supervisor
- 2016 – Ondřej Klejch, PhD co-supervisor
- 2015 – Joachim Fainberg, PhD co-supervisor
- 2014 – 2018 Ahmed Ali, PhD co-supervisor
- 2011 – 2016 Pawel Swietojanski, PhD co-supervisor
- 2014 – 2018 Primary supervisor of ten MSc students
- 2010 – 2015 Co-supervisor of several MSc students

### **Teaching experience**

- 2013 – 2019 Guest lecturer, Speech Recognition (MSc)
- 2017 Guest lecturer, Machine Learning Practical (MSc)
- 2014 – 2016 Project supervisor (MInf)
- 2008 – 2009 Guest lecturer, demonstrator, lab practical and coursework designer, Computational Foundations of Cognitive Science (UG 1)
- 2009 Tutor and coursework designer, Informatics 2B (UG 2)
- 2008 – 2009 Tutor, Mathematics for Informatics (UG 1)

### **Knowledge exchange and impact**

- 2015 – 2019 Consultancy with Crossover Capital Ltd and Auris Technology Ltd, applying research outputs to child ASR, £50,000
- 2015 – 2017 Consultancy with CEDAR Audio Ltd, investigating ASR performance with microphone array processing, £3,000
- 2015 – Domain adaptation research extended through PhD studentship with Bloomberg Inc
- 2011 – 2013 Consultancy, research licensing and Knowledge Transfer Partnership with France Telecom, £40,000, plus further research funding
- 2012 Research tools licensed to Quorate Technology Ltd

### **Technical skills**

Extensive experience of Python, C/C++, Linux shell scripting, Kaldi and HTK, in both commercial and academic environments

### **Other professional activities**

Keynote speaker at Baltic Human Language Technology Conference, 2018

Invited talks at National Institute of Information and Communication Technology, Kyoto, Japan; National Institute of Informatics, Tokyo, Japan; University of Birmingham

Reviewer for IEEE/ACM Transactions on Audio, Speech and Language Processing; Computer Speech and Language; Interspeech; ICASSP; SLT; ASRU; ICLR; NIPS

Session chair at Interspeech

## Publications

- [1] Joanna Rownicka, Peter Bell, and Steve Renals, “Analysing deep CNN-based utterance embeddings for acoustic model adaptation,” in *Proc. SLT*, Dec. 2018.
- [2] Ondrej Klejch, Joachim Fainberg, and Peter Bell, “Learning to adapt: a meta-learning approach for speaker adaptation,” in *Proc. Interspeech*, Sept. 2018.
- [3] Peter Bell, Pawel Swietojanski, and Steve Renals, “Multitask learning of context-dependent targets in deep neural network acoustic models,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 2, pp. 238–247, 2017.
- [4] Peter Bell, Joachim Fainberg, Catherine Lai, and Mark Sinclair, “A system for real-time collaborative transcription correction,” in *Proc. Interspeech*, Aug. 2017.
- [5] Emiru Tsunoo, Ondrej Klejch, Peter Bell, and Steve Renals, “Hierarchical recurrent neural network for story segmentation using fusion of lexical and acoustic features,” in *Proc. ASRU*, Okinawa, Japan, December 2017.
- [6] Joanna Rownicka, Steve Renals, and Peter Bell, “Simplifying very deep convolutional neural network architectures for robust speech recognition,” in *Proc. ASRU*, December 2017.
- [7] Emiru Tsunoo, Peter Bell, and Steve Renals, “Hierarchical recurrent neural network for story segmentation,” in *Proc. Interspeech*, Aug. 2017.
- [8] Ahmed Ali, Preslav Nakov, Peter Bell, and Steve Renals, “WERd: Using social text spelling variants for evaluating dialectal speech recognition,” in *Proc. ASRU*, Dec. 2017.
- [9] Joachim Fainberg, Steve Renals, and Peter Bell, “Factorised representations for neural network adaptation to diverse acoustic environments,” *Proc. Interspeech*, pp. 749–753, Aug. 2017.
- [10] Ondřej Klejch, Peter Bell, and Steve Renals, “Sequence-to-sequence models for punctuated transcription combining lexical and acoustic features,” in *Proc. ICASSP*, 2017.
- [11] Ondřej Klejch, Peter Bell, and Steve Renals, “Punctuated transcription of multi-genre broadcasts using acoustic and lexical approaches,” in *Proc. SLT*, 2016, pp. 433–440.
- [12] Ahmed Ali, Peter Bell, James Glass, Yacine Messaoui, Hamdy Mubarak, Steve Renals, and Yifan Zhang, “The MGB-2 Challenge: Arabic multi-dialect broadcast media recognition,” in *Proc. SLT*, 2016.
- [13] Siva Reddy Gangireddy, Pawel Swietojanski, Peter Bell, and Steve Renals, “Unsupervised adaptation of recurrent neural network language models,” in *Proc. Interspeech*, 2016, pp. 2333–2337.
- [14] Joachim Fainberg, Peter Bell, Mike Lincoln, and Steve Renals, “Improving children’s speech recognition through out-of-domain data augmentation,” in *Proc. Interspeech*, 2016, pp. 1598–1602.
- [15] Peter Bell, Mark Gales, Thomas Hain, Jonathan Kilgour, Pierre Lanchantin, Xunying Liu, Andrew McParland, Steve Renals, Oscar Saz, Mirjam Wester, and Phil Woodland, “The MGB challenge: Evaluating multi-genre broadcast media recognition,” in *Proc. ASRU*, 2015.
- [16] Peter Bell and Steve Renals, “A system for automatic alignment of broadcast media captions using weighted finite-state transducers,” in *Proc. ASRU*, 2015.

- [17] Ahmed Ali, Walid Magdy, Peter Bell, and Steve Renals, “Multi-reference WER for evaluating ASR for languages with no orthographic rules,” in *Proc. ASRU*, 2015.
- [18] Adriana Stan, Yoshitaka Mamiya, Junichi Yamagishi, Peter Bell, Oliver Watts, Rob Clark, and Simon King, “ALISA: An automatic lightly supervised speech segmentation and alignment tool,” *Computer Speech and Language*, vol. 35, pp. 116–133, 2016.
- [19] Alessandra Cervone, Catherine Lai, Silvia Pareti, and Peter Bell, “Towards automatic detection of reported speech in dialogue using prosodic cues,” in *Proc. Interspeech*, Dresden, Germany, Sept. 2015.
- [20] Peter Bell, Catherine Lai, Clare Llewellyn, Alexandra Birch, and Mark Sinclair, “A system for automatic broadcast news summarisation, geolocation and translation,” in *Proc. Interspeech (demo session)*, Dresden, Germany, Sept. 2015. Winner of best demo award.
- [21] Peter Bell and Steve Renals, “Regularization of context-dependent deep neural networks with context-independent multi-task training,” in *Proc. ICASSP*, 2015.
- [22] Pawel Swietojanski, Peter Bell, and Steve Renals, “Structured output layer with auxiliary targets for context-dependent acoustic modelling,” in *Proc. Interspeech*, 2015.
- [23] Peter Bell and Steve Renals, “Complementary tasks for context-dependent deep neural network acoustic models,” in *Proc. Interspeech*, 2015.
- [24] Mark Sinclair, Peter Bell, Alexandra Birch, and Fergus McInnes, “A semi-markov model for speech segmentation with an utterance-break prior,” in *Proc. Interspeech*, September 2014.
- [25] Peter Bell, Joris Driesen, and Steve Renals, “Cross-lingual adaptation with multi-task adaptive networks,” in *Proc. Interspeech*, 2014.
- [26] A. Cervone, S. Pareti, P. Bell, I. Prodanof, and T. Caselli, “Detecting attribution relations in speech: a corpus study,” in *Proc. Italian Conference on Computational Linguistics*, 2014.
- [27] Peter Bell, Fergus McInnes, Siva Reddy Gangireddy, Mark Sinclair, Alexandra Birch, and Steve Renals, “The UEDIN english ASR system for the IWSLT 2013 evaluation,” in *Proc. International Workshop on Spoken Language Translation*, 2013.
- [28] Joris Driesen, Peter Bell, Mark Sinclair, and Steve Renals, “Description of the UEDIN system for German ASR,” in *Proc IWSLT*, December 2013.
- [29] Peter Bell, Hitoshi Yamamoto, Pawel Swietojanski, Youzheng Wu, Fergus McInnes, Chiori Hori, and Steve Renals, “A lecture transcription system combining neural network acoustic and language models,” in *Proc. Interspeech*, August 2013.
- [30] Yoshitaka Mamiya, Adriana Stan, Junichi Yamagishi, Peter Bell, Oliver Watts, Robert Clark, and Simon King, “Using adaptation to improve speech transcription alignment in noisy and reverberant environments,” in *8th ISCA Workshop on Speech Synthesis*, August 2013, pp. 61–66.
- [31] Heidi Christensen, Magda Aniol, Peter Bell, Phil Green, Thomas Hain, Simon King, and Pawel Swietojanski, “Combining in-domain and out-of-domain speech data for automatic recognition of disordered speech,” in *Proc. Interspeech*, August 2013.
- [32] Adriana Stan, Peter Bell, Junichi Yamagishi, and Simon King, “Lightly supervised discriminative training of grapheme models for improved sentence-level alignment of speech and text data,” in *Proc. Interspeech*, August 2013.

- [33] Peter Bell, Pawel Swietojanski, and Steve Renals, “Multi-level adaptive networks in tandem and hybrid ASR systems,” in *Proc. ICASSP*, May 2013.
- [34] Ramya Rasipuram, Peter Bell, and Mathew Magimai.-Doss, “Grapheme and multilingual posterior features for under-resourced speech recognition: a study on Scottish Gaelic,” in *Proc. ICASSP*, May 2013.
- [35] Peter Bell, Mark Gales, Pierre Lanchantin, Xunying Liu, Yanhau Long, Steve Renals, Pawel Swietojanski, and Phil Woodland, “Transcription of multi-genre media archives using out-of-domain data,” in *Proc. IEEE Workshop on Spoken Language Technology*, December 2012, pp. 324–329.
- [36] Adriana Stan, Peter Bell, and Simon King, “A grapheme-based method for automatic alignment of speech and text data,” in *Proc. IEEE Workshop on Spoken Language Technology*, December 2012.
- [37] Eva Hasler, Peter Bell, Arnab Ghoshal, Barry Haddow, Philipp Koehn, Fergus McInnes, Steve Renals, and Pawel Swietojanski, “The UEDIN system for the IWSLT 2012 evaluation,” in *Proc. International Workshop on Spoken Language Translation*, 2012.
- [38] Myroslava Dzikovska, Amy Isard, Peter Bell, Johanna Moore, Natalie Steinhauser, and Gwendolyn Campbell, “Beetle II: an adaptable tutorial dialogue system,” in *Proceedings of the SIGDIAL 2011 Conference, demo session*. June 2011, pp. 338–340, Association for Computational Linguistics.
- [39] Peter Bell, Myroslava Dzikovska, and Amy Isard, “Designing a spoken language interface for a tutorial dialogue system,” in *Proc. Interspeech*, September 2012.
- [40] Myroslava O. Dzikovska, Peter Bell, Amy Isard, and Johanna D. Moore, “Evaluating language understanding accuracy with respect to objective outcomes in a dialogue system,” in *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*. April 2012, pp. 471–481, Association for Computational Linguistics.
- [41] Peter Bell, *Full covariance modelling for speech recognition*, Ph.D. thesis, University of Edinburgh, 2010.
- [42] Dong Wang, Simon King, Joe Frankel, and Peter Bell, “Stochastic pronunciation modelling and soft match for out-of-vocabulary spoken term detection,” in *Proc. ICASSP*, March 2010.
- [43] Dong Wang, Simon King, Joe Frankel, and Peter Bell, “Term-dependent confidence for out-of-vocabulary term detection,” in *Proc. Interspeech*, September 2009, pp. 2139–2142.
- [44] Peter Bell and Simon King, “Diagonal priors for full covariance speech recognition,” in *Proc. IEEE Workshop on Automatic Speech Recognition and Understanding*, December 2009.
- [45] Peter Bell and Simon King, “A shrinkage estimator for speech recognition with full covariance HMMs,” in *Proc. Interspeech*, September 2008, Shortlisted for best student paper award.
- [46] Peter Bell and Simon King, “Covariance updates for discriminative training by constrained line search,” in *Proc. Interspeech*, September 2008.
- [47] Peter Bell and Simon King, “Sparse gaussian graphical models for speech recognition,” in *Proc. Interspeech 2007*, August 2007.
- [48] Peter Bell, Tina Burrows, and Paul Taylor, “Adaptation of prosodic phrasing models,” in *Proc. Speech Prosody 2006*, May 2006.