

Joint CBS and COMP Research Seminar

Democratic Regression

Presented by

Tim Baldwin

Professor, School of Computing and Information Systems
Director, ARC Training Centre in Cognitive Computing for Medical Technologies
Associate Dean (Research Training), Melbourne School of Engineering
The University of Melbourne

Date: 22 November 2018 (Thursday)

Time: 11:00 - 12:30

Venue: PQ703

Abstract:

Two of the key tenets of modern-day democracy are the active participation of the people in politics and civic life, and political pledges to the people by political parties, e.g. in the form of manifestos. In this talk, I will present recent work on two political text regression problems: (1) prediction of the right-left leaning of a political party on the basis of their manifesto document; and (2) prediction of the popularity of an online petition, in terms of the number of votes it will receive. In both cases, we propose neural models with mixed objective functions incorporating auxiliary tasks, and in the case of manifesto analysis, further smooth the model output based on contextual and temporal constraints using probabilistic soft logic.

Biosketch:

Tim Baldwin is a Professor in the School of Computing and Information Systems, The University of Melbourne, and Associate Dean (Research Training) within the Melbourne School of Engineering. He is also Director of the recently-awarded ARC Centre for Cognitive Computing in Medical Technologies, in partnership with IBM Research. He has previously held visiting positions at Cambridge University, Tokyo Institute of Technology, University of Washington, University of Tokyo, Saarland University, NTT Communication Science Laboratories, and National Institute of Informatics. His primary research focus is on natural language processing (NLP), including social media analytics, computational lexical semantics, deep learning, and topic modelling.

Prof. Baldwin completed a BSc(CS/Maths) and BA(Linguistics/Japanese) at The University of Melbourne in 1995, and an MEng(CS) and PhD(CS) at the Tokyo Institute of Technology in 1998 and 2001, respectively. Prior to joining The University of Melbourne in 2004, he was a Senior Research Engineer at the Center for the Study of Language and Information, Stanford University (2001-2004). He is the author of over 350 peer-reviewed publications across diverse topics in natural language processing.