ANDREW ZAMMIT MANGION, **University of Wollongong**
Approximate Inference with Deep Compositional Spatial Models

Non-stationary, anisotropic spatial processes are often used when modelling, analysing and predicting complex environmental phenomena. One such class of processes considers a stationary, isotropic process on a warped spatial domain. The warping function is generally difficult to fit and not constrained to be bijective, often resulting in ’space-folding.’ In the first part of the talk I will propose modelling a bijective warping function through a composition of multiple elemental bijective functions in a deep-learning framework, which ensures that there is no space-folding by construction. In the second part I will discuss inference when the resulting deep spatial model is a non-stationary Gaussian process or a general non-Gaussian process composed from a set of processes that are conditionally Gaussian. In the latter case I show how stochastic variational inference can be used for parameter estimation and for approximating predictive distributions as Gaussian mixtures. Through experiments in one and two dimensions I show that the deep compositional spatial models are quick to fit using approximate inference algorithms, and are able to provide better predictions and uncertainty quantification than other deep stochastic models of similar complexity.