

Spatially Sparse Convolutional Neural networks

Benjamin Graham


Department of Statistics
University of Warwick

Tuesday, 31st May, 4 p.m.

D202 Seminar room, School of Engineering, 2nd Floor

Abstract: Convolutional neural networks (CNNs) perform well on 2D and 3D classification problems such as handwriting recognition, image classification, and 3D object recognition. If the input to a convolutional network is sparse, for example a 1D pen stroke on a 2D piece of paper, or a 2D surface in 3D space, then it makes sense to use that sparsity to reduce computational cost. This is just like matrix arithmetic: if you are dealing with sparse matrices, you can save memory and time by using an appropriate data structure. What is interesting is that this allows us to look at sparse spatial structures at higher resolution than would otherwise be practical.

More info: <http://www2.warwick.ac.uk/fac/sci/wcpm/seminars>



WCPM

Warwick Centre for Predictive Modelling