

1 Speaker Biography

Andrew Fitzgibbon is a principal researcher in the computer vision group at Microsoft Research Cambridge. He is best known for his work on 3D vision, having been a core contributor to the Emmy-award-winning 3D camera tracker “boujou” (www.boujou.com) and Kinect for Xbox 360, but his interests are broad, spanning computer vision, graphics, machine learning, and even a little neuroscience. He has published numerous highly-cited papers, and received many awards for his work, including nine “best paper” prizes, the Silver medal of the Royal Academy of Engineering, and the BCS Roger Needham award. He is a fellow of the Royal Academy of Engineering, the British Computer Society, and the International Association for Pattern Recognition. Before joining Microsoft in 2005, he was a Royal Society University Research Fellow at Oxford University, having previously studied at Edinburgh University, Heriot-Watt University, and University College, Cork.

or nuisance parameters. This tutorial will describe several tools and techniques for solving such optimization problems, covering a broad syllabus: vector and matrix calculus; nonlinear optimization algorithms such as gradient descent and its many variants, Quasi-Newton and Gauss-Newton derivatives such as LBFGS and Levenberg Marquardt; modelling options for curves and surfaces; how to deal with missing data; how to deal with outliers, including how to optimize functions with robust kernels. There will be maths, and lots of it, but I hope to make it clear enough that anyone with decent high-school linear algebra and calculus can benefit, while nevertheless showing mathematical experts a few useful tricks.



2 Abstract

In vision and machine learning, almost everything we do may be considered to be a form of model fitting. Whether estimating the parameters of a convolutional neural network, target tracking, computing low-dimensional representations of datasets, computing structure and motion from image collections, estimating parameters for an inference model such as Markov random fields, or extracting shape spaces such as active appearance models, it almost always boils down to minimizing an objective containing some parameters of interest as well as some latent