

A Fast Algorithm for Matrix Balancing

Philip A. Knight (University of Strathclyde)
joint work with
Daniel Ruiz (ENSEEIHT)

As long as a square nonnegative matrix A contains sufficient nonzero elements, then the matrix can be balanced, that is we can find a diagonal scaling of A that is doubly stochastic. A number of algorithms have been proposed to achieve the balancing, the most well known of these being the Sinkhorn-Knopp algorithm. In this talk we introduce new algorithms based on inner-outer iteration schemes. We show that the Sinkhorn-Knopp algorithm belongs to this family, but other members can converge much more quickly. In particular, we show that while stationary iterative methods offer little or no improvement in many cases, a scheme using a preconditioned conjugate gradient method as the inner iteration can give quadratic convergence at low cost.