

Two-level preconditioning for Ridge Regression

Karl Meerbergen¹ Joris Tavernier² Yves Moreau³ Jaak Simm⁴

Solving linear systems is often the computational bottleneck in real-life problems. Iterative solvers are the only option due to the complexity of direct algorithms or because the system matrix is not explicitly known. Here, we develop a two-level preconditioner for regularized least squares linear systems involving a feature or data matrix. Variants of this linear system may appear in machine learning applications, such as ridge regression, logistic regression, support vector machines and Bayesian regression. We use clustering algorithms to create a coarser level that preserves the principal components of the covariance or Gram matrix. This coarser level approximates the dominant eigenvectors and is used to build a subspace preconditioner accelerating the Conjugate Gradient method. We observed speed-ups for artificial and real-life data.

References:

[1] <https://onlinelibrary.wiley.com/doi/10.1002/nla.2371>

¹KU Leuven
Karl.Meerbergen@kuleuven.be

²OpenAnalytics
tavernierjoris@gmail.com

³KU Leuven
Yves.Moreau@kuleuven.be

⁴KU Leuven
jaak.simm@kuleuven.be