The problem of handling dense kernel matrices, arising as Green's operators in numerical PDE and as covariance operators in the theory of Gaussian processes, represents a major computational bottleneck in many fields of application.
Straightforward Cholesky decomposition of the kernel matrix results in cubic scaling in the number of data points, which is prohibitive for large data sets.
We introduce a novel compression method for kernel matrices, that allows for the rapid evaluation of the forward- and inverse operator, and the determinant, as well as rapid sampling from the underlying distribution.
As a byproduct, we also obtain a sparse approximate PCA of the kernel matrix.
For covariance kernels arising from Green's functions of elliptic boundary value problems and sampling with approximately uniform density, we prove that our method has near linear complexity in time and space, in the number of data points.
We also show numerically that we achieve good results for the Matérn family, even though it is not within the scope of our theory, due to the lack of boundary conditions.