Structural or analytic derivation

The document Gram matrix and term Gram matrix have the same nonzero eigenvalues, because they arise from the same SVD. This is the basis of choosing the smaller dimension in latent semantic computations.

Comparative cost and preferred method

Compute a requested pair by intersecting its sparse term sets; materialize the full Gram matrix only when its complete entries are needed. For a linear operator action, $C^\top(Cv)$ avoids dense storage. No general time improvement is claimed for writing an output containing all document pairs.