

We propose a new class of convex penalty functions, called variational Gram functions (VGFs), that can promote pairwise relations, such as orthogonality, among a set of vectors in a vector space. These functions can serve as regularizers in convex optimization problems arising from hierarchical classification, multitask learning, estimating vectors with disjoint supports, and other applications. We study necessary and sufficient conditions under which a VGF is convex, and give a characterization of its subdifferential. In addition, we show how to compute its proximal operator, and discuss efficient optimization algorithms for some structured loss-minimization problems using VGFs. Numerical experiments are presented to demonstrate the effectiveness of VGFs and the associated optimization algorithms.