

Neural Networks for Regular Matroid Polytopes

Linear optimization over a matroid base polytope and computing the basis generating polynomial of a matroid are basic computational questions. While both are ‘easy’ for regular matroids in the ‘usual’ model of computation, the story gets much more subtle when one restricts the allowed basic operations.

Based on a fine-tuned version of Seymour’s decomposition of regular matroids, we derive an efficient way to compute the basis generating polynomial of a regular matroid without subtractions. By ‘tropicalization’, this implies that there are small ReLU neural networks for weighted basis maximization of regular matroids. This again can be interpreted as a new representation of base polytopes of regular matroids as Minkowski difference of projections of polytopes with ‘few’ facets. More precisely, this is the first example where the best known extended formulation is bigger than the virtual extended formulation.

This is based on joint work with Christoph Hertrich and Stefan Kober.