Abstract

Linear superiorization (LinSup) considers linear programming problems but instead of attempting to solve them with linear optimization methods it employs perturbation resilient feasibility-seeking algorithms and steers them toward reduced (not necessarily minimal) target function values. The two questions that we set out to explore experimentally are: (i) does LinSup provide a feasible point whose linear target function value is lower than that obtained by running the same feasibility-seeking algorithm without superiorization under identical conditions? (ii) How does LinSup fare in comparison with the Simplex method for solving linear programming problems? Based on our computational experiments presented here, the answers to these two questions are: 'yes' and 'very well', respectively.