

Structure of Quadraticizations of Pseudo-Boolean Functions

Aritanan Borges Garcia Gruber - UFABC

Abstract

Pseudo-Boolean functions are real-valued mappings over the set of binary vectors, and it is well known that such functions admit unique multilinear polynomial representations over their variables. The problem of minimizing pseudo-Boolean functions has received a lot of attention in the past five decades since it encompasses numerous optimization models, including the well-known MAX-SAT and MAX-CUT problems, and has applications in areas ranging from combinatorics to computational complexity to physics to computer vision. When the model or application leads to the minimization of a quadratic pseudo-Boolean function, the optimal solution can be efficiently obtained provided the function possesses the diminishing returns structure, also known as submodularity. When submodularity is absent, one of the most frequently used technique is based on roof-duality (Hamer, Hansen, and Simeone, 1984), which aims at finding in polynomial time a simpler form of the given quadratic minimization problem by fixing some of the variables at their provably optimum value (persistency) and decomposing the residual problem into variable disjoint smaller sub problems (Boros and Hammer, 1989). Evaluations on real and synthetic data have shown such method to be very effective, sometimes fixing up to 80-90% of the variables at their provably optimum value (Boros, Hammer, Sun, and Tavares, 2008). In the past decade, many applications of pseudo-Boolean optimization, especially those stemming from the machine learning and computer vision communities, require a higher-degree multilinear polynomial to capture the existing nuances, resulting in super-quadratic objective functions. In these cases, however, fewer effective techniques are available. In particular, there is no analog to the persistencies provided by roof-duality for the quadratic case. The increased interest led to the study of methods to reduce higher-degree minimization problems to quadratic ones. In the span of six years, many techniques were developed to achieve that goal. Despite the apparent variety, they all share the common theme of being based on local transformations, and it is not at all clear how good those transformations are and how large the resulting quadratic problems can be. Continuing the investigation started in (Boros and Gruber, 2012) and (Fix, Gruber, Boros, and Zabih, 2015), we follow a global approach, considering all possible quadraticizations of a given higher-degree pseudo-Boolean function. In fact, we consider several different types of quadraticizations and provide sharp lower and upper bounds on the number of extra variables required to ?quadraticize? the function (Anthony, Boros, Crama, and Gruber, 2016, 2017). Our results establish some interesting facts about the structure of pseudo-Boolean functions, show some nice connections to linear algebra, extremal combinatorics, and circuit complexity, and give rise to new quadraticization algorithms as well. (Joint work with M. Anthony, E. Boros, Y. Crama, and E. Rodriguez-Heck).