Pitfalls to avoid while using multiobjective optimization for machine learning

Junaid Akhter[1]

[1] University of Paderborn
junaid.akhter@uni-paderborn.de

Recently, there has been an increasing interest in exploring the application of multiobjective optimization (MOO) in machine learning (ML), the reason being the numerous situations in real-life applications where multiple objectives need to be optimized simultaneously. A key aspect of MOO is the existence of a Pareto set rather than a single optimal solution, illustrating the inherent trade-offs between objectives. Despite its potential, there is a noticeable lack of satisfactory literature that could serve as an entry-level guide for ML practitioners who want to use MOO. Hence, our goal in this paper is to produce such a resource. We critically review previous studies, particularly those involving MOO in deep learning (using Physics-Informed Neural Networks (PINNs) as a guiding example), and identify misconceptions that highlight the need for a better grasp of MOO principles in ML. Using MOO of PINNs as a case study, we demonstrate the interplay between the *data loss* and the *physics loss* terms. We highlight the most common pitfalls one should avoid while using MOO techniques in ML. We begin by establishing the groundwork for MOO, focusing on well-known approaches such as the weighted sum (WS) method, alongside more complex techniques like the multiobjective gradient descent algorithm (MGDA). We emphasize the importance of understanding the specific problem, the objective space, and the selected MOO method, while also noting that neglecting factors such as convergence can result in inaccurate outcomes and, consequently, a non-optimal solution. Our goal is to offer a clear and practical guide for ML practitioners to effectively apply MOO, particularly in the context of DL.