Projection-Based Approximate Benders Decomposition for Scalable and Privacy-Preserving Optimization

Benders decomposition (BD) is a classical technique for solving large-scale mixed-integer linear programs. BD is an **iterative technique** where the **master problem makes a decision**, and one or more **subproblems evaluate the consequences of that decision**. The subproblems provide feedback that helps the master problem improve its approximation of the true objective function. This feedback takes the form of cuts that are added to the master problem, allowing it to adjust its choices in the next iteration. The process continues until the master problem converges to an optimal solution.

BD has been successfully applied to a range of domains, including **facility location, stochastic programming, network design, and supply chain optimization**. In energy markets, for instance, Benders decomposition is used to model long-term investment and operational decisions under uncertainty. Similarly, in decentralized logistics, the method can be used to model supply chain coordination problems where multiple actors contribute data or decisions across different layers of the network.

Despite its flexibility, BD faces **two fundamental limitations**: **scalability and privacy**. As the number of subproblems grows, particularly in scenario-based stochastic programming, **the time spent solving subproblems can dominate the total computational cost**. This becomes a critical bottleneck when subproblems are high-dimensional or when they need to be solved many times in the master iteration process. Equally important, in decentralized or multi-agent contexts, each **subproblem may correspond to a different stakeholder** or scenario owner who may not be willing to share internal data such as costs, constraints, or operational models. This creates a tension between the centralized nature of cut generation in BD and the need to preserve data confidentiality.

To address these challenges, we propose an approach based on **projecting the space of the subproblems into a smaller space**. The idea is to reduce the dimensionality of the variables that link the master and subproblem using a **random or learned projection**. In the case of random projection, we leverage the **Johnson-Lindenstrauss Lemma to ensure that the geometry of the problem is approximately preserved**, thereby allowing the subproblem to be solved in a lower-dimensional space. In a second, more adaptive approach, we use **machine learning, specifically an encoder-decoder architecture, to learn instance-specific projections** that optimize the trade-off between subproblem fidelity and dimensionality reduction. In both cases, **the projection obfuscates the original variable values, thereby preserving the privacy** of the master-side decisions while also reducing the size and computational cost of the subproblem.

Timeline

The PhD project will be structured across **three main phases**, progressively building from theoretical foundations to practical implementations and real-world applications.

In the first year, the focus will be on developing a **baseline implementation of BD with random projections**. The main objective will be to understand how projecting the master-side decision variables into a lower-dimensional space impacts the feasibility, convergence, and overall quality of the decomposition process. This includes both **theoretical and empirical investigations**. A key research direction will be to **determine whether random projections preserve sufficient geometric structure to ensure meaningful cut generation**. More precisely, we aim to characterize the contexts in which random projections lead to a good estimation of the subproblem cost, and to develop methods for quantifying both the **underestimation and overestimation errors** introduced by the projection. It is important to emphasize that, regardless of the direction of the approximation error, **the solution produced by the master problem remains valid**. Another core aspect will be to analyze **to what extent random projections implicitly protect sensitive information** by obfuscating the original decision variables. In this regard, **a formal definition of privacy in the context of Benders decomposition will be introduced**, along with an associated measure to assess the effectiveness of projection-based privacy. By the end of this phase, we expect to deliver a working prototype of the projected Benders framework along with a systematic evaluation of its scalability, approximation quality, and privacy-preserving properties, ideally resulting in a **journal submission**.

The second year will shift attention to **learning projections** adapted to the structure of the problem inspired by an encoder-decoder approach. The objective is to **replace fixed random projections with trainable mappings** that can automatically discover low-dimensional representations of master problem decisions, while preserving the ability of the subproblems to produce valid and informative feedback. We will also investigate **hybrid approaches**—combining learned and random projections— to introduce useful randomness that enhances both performance and privacy. More concretely, the encoder network will **learn a projection** from the original decision space to a compressed space, which will be used to define the subproblem. The decoder may be used to recover approximate primal or dual quantities in the original space. This will involve designing a loss function that penalizes cut inaccuracy or subproblem suboptimality, possibly informed by the dual gap or violation of the true cost function. We expect to **submit a publication** summarizing the methodological and empirical advances made in this second year.

The third year will focus on consolidating the findings from the previous stages and **formalizing privacy guarantees**. A central component of this phase will be a systematic **comparison of the two projection strategies**—random versus learned—in terms of solution quality, computational efficiency, generalizability, and privacy preservation. Particular attention will be given to **developing metrics for measuring information leakage from cuts and subproblem responses**. Applications in energy markets and decentralized logistics will serve as benchmark domains to demonstrate the practical relevance of the proposed methods, especially in scenarios where subproblems are managed by independent actors with sensitive data. The final year will also involve preparing the thesis manuscript. A **journal submission** synthesizing both the methodological innovations and **application-driven managerial insights** will be targeted, along with the defense of the dissertation.