Big data and new methods allow firms to personalize their marketing actions for different customers. The growth in industry interest in personalization has been mirrored by rapid growth in academic interest in personalization (targeting). This academic attention has primarily focused on training personalization policies without constraints. However, in practice, internal business rules or society considerations often impose constraints on firms' targeting policies. In this paper, we investigate how to train scalable targeting policies in the presence of constraints.

Constraints on targeting policies are typically of two types. *Volume* constraints limit the total number of marketing actions that can be taken. This type of constraint may result from capacity constraints. For example, a firm's ability to make outbound phone calls may be limited by the availability of trained associates to make these calls. Budget constraints may also impose minimum and/or maximum limits on the total number of marketing actions. These constraints may operate in aggregate, or they may apply to specific customer segments. *Similarity* constraints limit differences in marketing actions taken with different customer segments. These constraints are often motivated by concerns for fairness. For example, a constraint might require that the firm takes similar marketing actions with neighboring zip codes, or that customers located near one store are treated with similar marketing actions as customers located near other stores.

While there are many standard methods for optimizing problems with constraints, large numbers of decision variables and large numbers of constraints can both make the problem challenging. In a personalization problem, the number of decision variables and the number of constraints can both potentially be large. For example, if a separate decision is made for each customer, the number of decision variables may be in the millions. Even where decisions are made at the customer segment level, if there are many segments, there will be many decision variables. Similarly, if constraints apply to specific segments, the number of constraints will be at least as large as the number of segments and may grow as a polynomial function of the number of segments. Standard optimization methods are not well-suited to solving optimization problems with either large numbers of decision variables, or large numbers of constraints.

We formulate the personalization problem as a linear programming problem with constraints and illustrate how to incorporate both volume constraints and similarity constraints. We then adapt and apply a recently developed algorithm, which is designed to solve large-scale linear programming problems. The algorithm belongs within the class of first-order methods, which use gradient information to construct algorithms to find optimal solutions. This class of methods scales very well, and is widely used in many applications, including many machine learning algorithms. The algorithm that we adapt leverages the primal-dual hybrid gradient. Similar methods are widely used in image processing and computer vision. Recent developments have made the algorithm especially suitable for large-scale linear programming problems.

We provide two theoretical results. The first result compares the proposed algorithm with state-of-the-art benchmarks: primal simplex, dual simplex and barrier methods. We prove that the algorithm can solve larger problems (in terms of customers and constraints) than any of these methods. The second theoretical result extends existing guarantees on optimality and computation speed, by adjusting the method (and existing theory) to accommodate the characteristics of personalization problems.

To illustrate the practical value of the method, we apply it to an actual personalization problem. The problem involves choosing which promotions to send to prospective customers. The response functions are estimated using a large-scale field experiment that includes approximately 2.4 million households and five marketing actions. The findings provide practical empirical evidence of how the proposed method extends the scale of solvable personalization problems in the presence of constraints. Our analysis recognizes that different firms that have access to different hardware resources. As we change the available hardware, our method consistently solves problems faster, yields higher profits, and accommodates both a larger number of customers and a larger number of constraints.