
Differentiating the Value Function by using Convex Duality

Sheheryar Mehmood

Department of Mathematics,
University of Tübingen,
Germany

Peter Ochs

Department of Mathematics,
University of Tübingen,
Germany

Abstract

We consider the differentiation of the value function for parametric optimization problems. Such problems are ubiquitous in Machine Learning applications such as structured support vector machines, matrix factorization and min-min or minimax problems in general. Existing approaches for computing the derivative rely on strong assumptions of the parametric function. Therefore, in several scenarios there is no theoretical evidence that a given algorithmic differentiation strategy computes the true gradient information of the value function. We leverage a well known result from convex duality theory to relax the conditions and to derive convergence rates of the derivative approximation for several classes of parametric optimization problems in Machine Learning. We demonstrate the versatility of our approach in several experiments, including non-smooth parametric functions. Even in settings where other approaches are applicable, our duality based strategy shows a favorable performance.

1 Introduction

Given a function $f : \mathbb{R}^N \times \mathbb{R}^P \rightarrow \mathbb{R}$ with values $f(\mathbf{x}, \mathbf{u})$, we consider the following parametric optimization problem:

$$p(\mathbf{u}) := \inf_{\mathbf{x} \in \mathbb{R}^N} f(\mathbf{x}, \mathbf{u}). \quad (\mathcal{P})$$

The optimal value of $f(\cdot, \mathbf{u})$, which we denote by $p(\mathbf{u})$, depends on the parameter \mathbf{u} and is commonly referred

to as the *value function* of (\mathcal{P}) or the *infimal projection* of f . When the minimum is attained at some $\mathbf{x}^*(\mathbf{u}) \in \mathbb{R}^N$ for a given $\mathbf{u} \in \mathbb{R}^P$, the value function is given by $p(\mathbf{u}) = f(\mathbf{x}^*(\mathbf{u}), \mathbf{u})$. For many applications, quantifying the change in p with respect to \mathbf{u} is key, which is achieved by computing gradient or subgradient information of p .

This is particularly true for Machine Learning applications, for which a parametric dependency occurs naturally, for example, when solving a min-min or minimax optimization problem, in Structured Support Vector Machines (Taskar et al., 2004; Tsochantaridis et al., 2005), Sparse Dictionary Learning (Mairal et al., 2010), Generative Adversarial Networks (Goodfellow et al., 2014) and Matrix Factorization. Another important area where such derivative information is crucial is the Sensitivity Analysis of an optimization problem, which finds applications in the shadow price problem (Still, 2018, Section 4.3) and also in bridge crane design or breakwater modeling (Castillo et al., 2008). The decision-making is based on a measure of how sensitive the model is when parameters \mathbf{u} are changed.

If $\mathbf{x}^*(\mathbf{u})$ is available and differentiable, the gradient information can be computed by differentiating $p(\mathbf{u}) = f(\mathbf{x}^*(\mathbf{u}), \mathbf{u})$ with respect to \mathbf{u} , i.e.,

$$\begin{aligned} \nabla p(\mathbf{u}) &= D_{\mathbf{u}} \mathbf{x}^*(\mathbf{u}) \nabla_{\mathbf{x}} f(\mathbf{x}^*(\mathbf{u}), \mathbf{u}) + \nabla_{\mathbf{u}} f(\mathbf{x}^*(\mathbf{u}), \mathbf{u}) \\ &= \nabla_{\mathbf{u}} f(\mathbf{x}^*(\mathbf{u}), \mathbf{u}). \end{aligned}$$

However, clearly, this approach demands for strong smoothness conditions of the parametric function f and the solution mapping $\mathbf{x}^*(\mathbf{u})$, which are not satisfied for common Machine Learning applications.

Consider for example the following sparsity constrained linear regression problem:

$$\min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2, \text{ s.t. } \|\mathbf{x}\|_1 \leq u, \quad (1)$$

where $\mathbf{A} \in \mathbb{R}^{M \times N}$, $\mathbf{b} \in \mathbb{R}^M$, and $u \geq 0$. As a constrained optimization problem, the objective (including the constraint in terms of an indicator function) is not differentiable. The value function, however, is

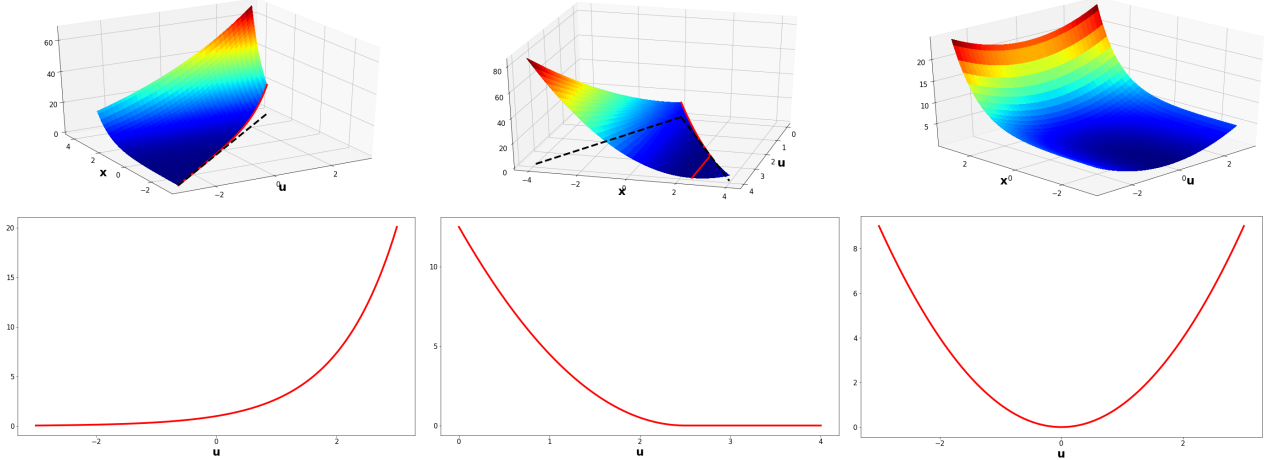


Figure 1: Graphical Depiction of some simple objective functions which have a smooth value function but either don't have a minimum or are not differentiable at the minimizer. (Top row) The objective f is plotted against x and u for some toy examples (left) $\mathbb{R} \times \mathbb{R} \ni (x, u) \mapsto f(x, u) = \exp(x) + \delta_{[u, +\infty)}(x)$, (middle) $\mathbb{R} \times [0, 4] \ni (x, u) \mapsto f(x, u) = (ax - b)^2/2 + \delta_{[-u, u]}(x)$ where $a, b \neq 0$ and (right) $\mathbb{R} \times \mathbb{R} \ni (x, u) \mapsto f(x, u) = \exp(x) + u^2/2$. The function $u \mapsto (u, x^*, f(x^*, u))$ is sketched in red in first two examples (left and middle). No such curve is shown for the third example because the minimum w.r.t. x is not attained. The black dashed-line (left and middle) shows the boundary of the feasible set in $\mathbb{R} \times \mathbb{R}$. (Bottom row) Value function p plotted against u for the corresponding examples.

continuously differentiable on $(0, \infty)$ and subdifferentiable at $u = 0$ (van den Berg and Friedlander, 2008). As noted in van den Berg and Friedlander (2008) and more generally in van den Berg and Friedlander (2011) and Aravkin et al. (2013), its gradient can be used to solve the following minimal norm problem:

$$\min_{\mathbf{x}} \|\mathbf{x}\|_1, \text{ s.t. } \|A\mathbf{x} - b\|_2 \leq w. \quad (2)$$

The problem in (1) is one of many instances where the parametric function f is jointly convex in its arguments. Yet algorithmic differentiation strategies based on differentiating approximations to the solution mapping cannot be applied. This is due to the fact that the boundary of the feasible set changes with \mathbf{u} and when the solution $\mathbf{x}^*(\mathbf{u})$ lies at the boundary for some \mathbf{u} , the subdifferential of f with respect to \mathbf{x} at $(\mathbf{x}^*(\mathbf{u}), \mathbf{u})$ is a shifted non-trivial cone, hence, in particular not single-valued. We explain this phenomenon more concretely in Section 2.2 (see also Figure 1).

As a remedy, we invoke standard results from convex duality of the function f to derive the above mentioned differentiability property of the value function for a large class of optimization problems including (1). In fact, beyond differentiability, we explore the formula

$$\partial p(\mathbf{u}) = \arg \max_{\mathbf{y} \in \mathbb{R}^P} \langle \mathbf{u}, \mathbf{y} \rangle - f^*(0, \mathbf{y}),$$

which expresses the convex (Fréchet) subdifferential of the value function p at \mathbf{u} as the set of solutions to a

certain optimization problem that depends on the convex conjugate f^* of f . For a jointly convex function f in (\mathbf{x}, \mathbf{u}) , the validity of the formula is asserted under the weak assumptions that $p(\mathbf{u})$ is finite (i.e. the infimum of $f(\cdot, \mathbf{u})$ in (\mathcal{P}) is finite) and $\mathbf{u} \in \text{ri}(\text{dom } p)$ lies in the relative interior of the domain of p . Therefore, in these situations, the problem of differentiating the value function p is equivalent to solving a convex optimization problem, which allows us to explore the large literature on convex optimization algorithms.

Since single-valuedness of the subdifferential of p implies differentiability, for example, strict convexity of $\mathbf{y} \mapsto f^*(0, \mathbf{y})$ implies differentiability of p without the need for $f(\mathbf{x}, \mathbf{u})$ to be differentiable. Therefore our approach allows us to compute the variation (gradient) of the value function p in situations for which commonly used direct differentiation strategies, for example based on automatic differentiation, cannot be applied. Nevertheless, even if the parametric function f is sufficiently smooth, the flexibility to apply various (optimal) convex optimization algorithms for computing this derivative information compares favorably with those direct differentiation strategies.

For the large class of optimization problems that we consider, we summarize algorithms with their convergence guarantees based on the properties of the objective function.

Remark. Differentiation of the value function p in (P) is not to be confused with differentiating the optimal solution mapping $\mathbf{x}^*(\mathbf{u})$ with respect to \mathbf{u} . Besides its usage in computing automatic and implicit gradient estimator, the argmin derivative is used in optimization layers, that is, neural networks whose output is given by solving an optimization problem (Amos and Kolter, 2017; Agrawal et al., 2019). It is also required in bilevel optimization (Dempe et al., 2015), a most well known application of which is gradient-based hyperparameter optimization or parameter learning (Domke, 2012; Kunisch and Pock, 2013; Deledalle et al., 2014).

Another problem which is similar to ours is the differentiation of a function $g(\mathbf{x}, \mathbf{u})$ with respect to the parameter \mathbf{u} evaluated at a solution $\mathbf{x}^*(\mathbf{u})$ of a system of parametric nonlinear equations $h(\mathbf{x}, \mathbf{u}) = 0$, where $g : \mathbb{R}^N \times \mathbb{R}^P \rightarrow \mathbb{R}^M$ and $h : \mathbb{R}^N \times \mathbb{R}^P \rightarrow \mathbb{R}^N$ satisfy some regularity conditions. We can also replace the function g with a functional and the non-linear system with a parametric ordinary differential equation. The two problems are related (but not equivalent) because when f in (P) is continuously differentiable and has a minimum $\mathbf{x}^*(\mathbf{u})$ in \mathbf{x} for a given \mathbf{u} , then $h = \nabla_{\mathbf{x}} f$ and $g = f$ with $M = 1$ and our goal is to differentiate $p(\mathbf{u}) = f(\mathbf{x}^*(\mathbf{u}), \mathbf{u})$. To differentiate $g(\mathbf{x}^*(\mathbf{u}), \mathbf{u})$ with respect to \mathbf{u} , we can make use of Piggyback differentiation (Griewank and Faure, 2003) or the Adjoint-state method (Pontryagin et al., 1962; Pollini et al., 2018). These techniques find their use in solving constrained optimization problems (where constraints are often given as ODE's or PDE's) with various applications in Geophysics (Plessix, 2006), Medicine (Knopoff et al., 2012) and Neural Networks (Chen et al., 2018).

2 Problem Setting

We consider parametric optimization problems of type (P) and seek for computing $\nabla p(\mathbf{u})$, i.e., the variation of the value function p with respect to \mathbf{u} . One of our major goals is to characterize the properties of f for which various numeric differentiation strategies with theoretical convergence guarantees can be used. We emphasize differentiation strategies based on iterative algorithms and provide convergence rates.

First, in Section 2.1, we recall the most widely used approaches for smooth parametric functions f , and demonstrate their limitations for several examples in Section 2.2. Therefore, in Section 3, as a remedy for such situations, we leverage a well known result from convex duality theory for numerical estimation of the variation of the value function, which allows us to classify problem classes with convergence rates.

2.1 Analytical, Automatic and Implicit Gradient Estimator

Ablin et al. (2020) analyze three different methods for iterative derivative approximation of smooth parametric functions f , provide convergences rates and enlighten a super-efficiency phenomenon for the automatic differentiation strategy. We recall their results.

Let f be twice continuously differentiable on $\mathbb{R}^N \times \mathbb{R}^P$ and $\mathbf{x}^*(\mathbf{u})$ be the unique minimizer for every $\mathbf{u} \in \mathbb{R}^P$ such that $\nabla_{\mathbf{x}}^2 f(\mathbf{x}^*(\mathbf{u}), \mathbf{u})$ is positive definite. From the Implicit Function Theorem, we derive that $\mathbf{x}^* : \mathbb{R}^P \rightarrow \mathbb{R}^N$ is continuously differentiable with derivative $D_{\mathbf{u}} \mathbf{x}^*(\mathbf{u}) = \varphi(\mathbf{x}^*(\mathbf{u}), \mathbf{u})$, where we define the mapping $\varphi : \mathbb{R}^N \times \mathbb{R}^P \rightarrow \mathbb{R}^{N \times P}$ as:

$$\varphi(\mathbf{x}, \mathbf{u}) = -[\nabla_{\mathbf{x}}^2 f(\mathbf{x}, \mathbf{u})]^{-1} \nabla_{\mathbf{x}} f(\mathbf{x}, \mathbf{u}). \quad (3)$$

The value function and its gradient are then given by:

$$p(\mathbf{u}) = f(\mathbf{x}^*(\mathbf{u}), \mathbf{u}) \text{ and } \nabla p(\mathbf{u}) = \nabla_{\mathbf{u}} f(\mathbf{x}^*(\mathbf{u}), \mathbf{u}).$$

The expression for ∇p follows from chain rule and the optimality condition $\nabla_{\mathbf{x}} f(\mathbf{x}^*(\mathbf{u}), \mathbf{u}) = 0$. The minimizer $\mathbf{x}^*(\mathbf{u})$ is estimated by an iterative optimization method which yields a sequence $(\mathbf{x}^{(k)})_{k \in \mathbb{N}}$ with limit $\mathbf{x}^*(\mathbf{u})$. In a realistic setting such a process is terminated after K iterations to yield a so-called sub-optimal solution $\mathbf{x}^{(K)}(\mathbf{u})$ for each \mathbf{u} . To compute ∇p , we either substitute $\mathbf{x}^{(K)}(\mathbf{u})$ in place of $\mathbf{x}^*(\mathbf{u})$ in the expression for ∇p to obtain the **analytic gradient estimator**:

$$\mathbf{g}_1^{(K)}(\mathbf{u}) := \nabla_{\mathbf{u}} f(\mathbf{x}^{(K)}, \mathbf{u}) \quad (\text{AnG})$$

or in the expression for p and then differentiate it with respect to \mathbf{u} , assuming that the sequence $(\mathbf{x}^{(k)})_{k \in \mathbb{N}}$ is differentiable, meaning that the mapping between successive iterates and the dependence on the parameter \mathbf{u} is differentiable, giving us the **automatic gradient estimator**:

$$\mathbf{g}_2^{(K)}(\mathbf{u}) := [D_{\mathbf{u}} \mathbf{x}^{(K)}]^T \nabla_{\mathbf{x}} f(\mathbf{x}^{(K)}, \mathbf{u}) + \nabla_{\mathbf{u}} f(\mathbf{x}^{(K)}, \mathbf{u}). \quad (\text{AuG})$$

The term $D_{\mathbf{u}} \mathbf{x}^{(K)}$ is an estimator of $D_{\mathbf{u}} \mathbf{x}^*$ and is obtained by applying automatic differentiation on $\mathbf{x}^{(K)}$, hence the name. Using the expression in (3) to estimate $D_{\mathbf{u}} \mathbf{x}^*$ yields the **implicit gradient estimator**, i.e.,

$$\mathbf{g}_3^{(K)}(\mathbf{u}) := [\varphi(\mathbf{x}^{(K)}, \mathbf{u})]^T \nabla_{\mathbf{x}} f(\mathbf{x}^{(K)}, \mathbf{u}) + \nabla_{\mathbf{u}} f(\mathbf{x}^{(K)}, \mathbf{u}). \quad (\text{IG})$$

Ablin et al. (2020) provide following error bounds for (AnG), (AuG) and (IG).

Theorem 1. Let $D := D_x \times D_u \subset \mathbb{R}^N \times \mathbb{R}^P$ be compact and f be m -strongly convex with respect to \mathbf{x} and twice differentiable over D with second derivatives $\nabla_{\mathbf{x}\mathbf{u}}f$ and $\nabla_{\mathbf{x}}^2f$ respectively L_{xu} and L_{xx} -Lipschitz continuous. Then the first derivatives $\nabla_{\mathbf{u}}f$ and $\nabla_{\mathbf{x}}f$ are respectively L_u and L_x -Lipschitz continuous and for $\mathbf{x}^{(k)}$ produced by $\mathbf{x}^{(k+1)} := \mathbf{x}^{(k)} - \tau \nabla_{\mathbf{x}}f(\mathbf{x}^{(k)}, \mathbf{u})$ with $\tau \leq 1/L_x$ and $\omega := 1 - m\tau$, following statements hold:

(a) The analytic estimator converges and we have:

$$\left\| \mathbf{g}_1^{(K)} - \nabla p(\mathbf{u}) \right\|_2 \leq L_x \left\| \mathbf{x}^{(0)} - \mathbf{x}^*(\mathbf{u}) \right\|_2 \omega^K$$

(b) The automatic estimator converges and for $C_k := \tau(L_x k + \omega/2)(L_{xu} + L_1 L_{xx})$ with $\left\| D_{\mathbf{u}}\mathbf{x}^{(k)} \right\|_2 \leq L_1$ we have:

$$\left\| \mathbf{g}_2^{(K)} - \nabla p(\mathbf{u}) \right\|_2 \leq C_K \left\| \mathbf{x}^{(0)} - \mathbf{x}^*(\mathbf{u}) \right\|_2 \omega^{2K-1}$$

(c) The implicit estimator converges and for $C := (L_{xu} + L_1 L_{xx})/2 + L_2 L_x$ with $\left\| \varphi(\mathbf{x}^{(K)}, \mathbf{u}) \right\|_2 \leq L_2$ we have:

$$\left\| \mathbf{g}_3^{(K)} - \nabla p(\mathbf{u}) \right\|_2 \leq C \left\| \mathbf{x}^{(0)} - \mathbf{x}^*(\mathbf{u}) \right\|_2 \omega^{2K}$$

Theorem 1 shows faster convergence of automatic and implicit estimators as compared to analytical estimator. The automatic estimator is more stable than implicit estimator as depicted experimentally in Ablin et al. (2020). This makes the automatic method a strong contender for estimating ∇p . It is also not computationally expensive thanks to reverse mode AD. The memory overhead is overcome by discarding the iterates $\mathbf{x}^{(k)}$ for $k = 0, \dots, K-1$ and using $\mathbf{x}^{(K)}$ only in all the calculations when going backward (Christianson, 1994; Mehmood and Ochs, 2020). Ablin et al. (2020) also study these methods under weaker conditions, for instance, when $f(\cdot, \mathbf{u})$ is μ -Łojasiewicz (Attouch and Bolte, 2009) which generalizes strong convexity. However, their results depend on strong smoothness assumptions for f .

2.2 Problems with Direct Differentiation

Obviously, the settings for which Ablin et al. (2020) provide convergence rate guarantees is quite limited. We would like to emphasize the fact that differentiability of the parametric function f is not required for that of the value function p . In this section, we show with simple examples that the necessary conditions like differentiability of the objective and the existence of the minimizer are the key disadvantages of the above methods.

Example 2. Let $f : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ be defined as:

$$f(x, u) = \exp(x) + \delta_{[u, +\infty)}(x),$$

then (AnG) and (IG) fail to converge for all $u \in \mathbb{R}$ while (AuG) converges only when $x^{(k)} > u$ for all k (eventually) and $D_{\mathbf{u}}\mathbf{x}^{(k)}$ converges to $D_{\mathbf{u}}\mathbf{x}^*(u)$.

Detail. f is jointly convex in x and u and for all $u \in \mathbb{R}$, $f(\cdot, u)$ is $\exp(u)$ -strongly convex and $x^* : \mathbb{R} \rightarrow \mathbb{R}$ and $p : \mathbb{R} \rightarrow \mathbb{R}$, given by $x^*(u) = u$ and $p(u) = \exp(u)$ respectively, are continuously differentiable on \mathbb{R} . On the other hand, f is neither differentiable with respect to x nor u at $(x^*(u), u)$ for any $u \in \mathbb{R}$. To see why f is not differentiable with respect to u , note that f is alternatively written as $f(x, u) = \exp(x) + \delta_{(-\infty, x]}(u)$. The subdifferential of f with respect to u is \mathbb{R}_+ when $x = u$ and $\{0\}$ when $x > u$. Thus when $x^{(k)} = u$ for some $k \in \mathbb{N}$, none of the above methods is useful here. If $x^{(k)} > u$ for all k , we get $g_1^{(k)}(u) = 0$ and $g_3^{(k)}(u) = 0$ since $\partial f / \partial u(x^{(k)}, u) = 0$ and $\partial^2 f / \partial u^2(x^{(k)}, u) = 0$. The automatic estimator is given by $g_2^{(k)}(u) = D_{\mathbf{u}}\mathbf{x}^{(k)}(u) \exp(u)$ because $\partial f / \partial x(x^{(k)}, u) = \exp(u)$. It converges to $p'(u) = \exp(u)$ only if $D_{\mathbf{u}}\mathbf{x}^{(k)}(u)$ converges to $D_{\mathbf{u}}\mathbf{x}^*(u) = 1$.

The convergence of $D_{\mathbf{u}}\mathbf{x}^{(k)}(u)$ to $D_{\mathbf{u}}\mathbf{x}^*(u)$ in Example 2 is possible only under limited conditions which we do not establish here. The next example considers the non-smooth parametric objective of (1) in an analogue 1D setting.

Example 3. Let $f : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ be defined as:

$$f(x, u) = \frac{1}{2}(ax - b)^2 + \delta_{[-u, u]}(x),$$

where $a, b \in \mathbb{R} \setminus \{0\}$, then for all $u \in (0, |b/a|)$, (AnG) and (IG) fail to converge while (AuG) converges only when $x^{(k)} \in (-u, u)$ for all k (eventually) and $D_{\mathbf{u}}\mathbf{x}^{(k)}$ converges to $D_{\mathbf{u}}\mathbf{x}^*(u)$.

Detail. f is jointly convex in x and u and for all $u \in (0, |b/a|)$, $f(\cdot, u)$ is a^2 -strongly convex and we have $x^*(u) = \text{sgn}(b/a)u$ and $p(u) = (ax^*(u) - b)^2/2$. Since $f(x, u) = (ax - b)^2/2 + \delta_{[|x|, |b/a|]}(u)$ and $\partial_{\mathbf{u}}f(x, u) = N_{[|x|, |b/a|]}(u)$, f is not differentiable with respect to u at $(x^*(u), u)$ for any $u \in (0, |b/a|)$. Given a sequence $x^{(k)} \in (-u, u)$ with limit $x^*(u)$ we have $g_1^{(k)} = 0$ and $g_3^{(k)} = 0$ because $\partial f / \partial u(x^{(k)}, u) = 0$ and $\partial^2 f / \partial u^2(x^{(k)}, u) = 0$. The automatic estimator $g_2^{(k)}(u) = D_{\mathbf{u}}\mathbf{x}^{(k)}(u)a(ax^{(k)} - b)$ converges to $p'(u) = D_{\mathbf{u}}\mathbf{x}^*(u)a(ax^* - b)$ when $D_{\mathbf{u}}\mathbf{x}^{(k)}(u)$ converges to $D_{\mathbf{u}}\mathbf{x}^*(u)$.

Example 4. Let $f : \mathbb{R} \times \mathcal{U} \rightarrow \mathbb{R}$ be defined as:

$$f(x, u) = \exp(x) + \frac{1}{2}u^2,$$

then (AnG), (AuG) and (IG) fail to converge for all $u \in \mathbb{R}$.

Detail. This case is obvious because $\inf_x \exp(x) = 0$ with minimum not attained, giving us $p(u) = u^2/2$ and $\arg \min_x f(x, u) = \emptyset$.

While one may still use the methods of Ablin et al. (2020) or cvxpylayers (Amos and Kolter, 2017; Agrawal et al., 2019) to efficiently estimate ∇p in situations like those presented above, it should be noted that differentiating the solution mapping is not a strictly more general approach than differentiating the value function. This is due to the failure of the applicability of the chain rule for evaluating $\nabla_u[f(\mathbf{x}^*(\mathbf{u}), \mathbf{u})]$ for non-smooth functions in general (Bolte and Pauwels, 2020). The concept of a subgradient is not defined for a vector-valued non-smooth function and must be replaced by graphical derivatives and coderivatives; see Section 9.D in Rockafellar and Wets (1998) and the chain rule in Theorem 10.49 in Rockafellar and Wets (1998). This calls for a theoretically justified approach for estimating ∇p beyond those which are currently available (Ablin et al., 2020; Amos and Kolter, 2017; Agrawal et al., 2019).

3 Dual Gradient Estimator

The discussion in the previous section suggests that a different method is needed which is independent of directly differentiating the parametric objective function f . Trading the differentiability assumption for a joint convexity assumption of f in (\mathbf{x}, \mathbf{u}) , we invoke the powerful convex duality for computing derivative information of the value function p in cases beyond differentiability of f . Moreover, the same statement provides an expression for the convex subdifferential of p .

Denoting the convex conjugate of a function p by

$$p^*(\mathbf{y}) := \sup_{\mathbf{u}} \langle \mathbf{y}, \mathbf{u} \rangle - p(\mathbf{u}),$$

and its biconjugate by $p^{**} := (p^*)^*$, the following result can be derived when strong duality, i.e., $p^{**} = p$, holds. The dual of the problem defined in (P) is given by:

$$p^{**}(\mathbf{u}) = \sup_{\mathbf{y} \in \mathbb{R}^P} \langle \mathbf{u}, \mathbf{y} \rangle - f^*(0, \mathbf{y}). \quad (\mathcal{D})$$

and for $\mathbf{u} \in \text{ri}(\text{dom } p)$

$$\partial p(\mathbf{u}) = \arg \max_{\mathbf{y} \in \mathbb{R}^P} \langle \mathbf{u}, \mathbf{y} \rangle - f^*(0, \mathbf{y}).$$

When $\partial p(\mathbf{u})$ is single-valued, p is differentiable at \mathbf{u} and therefore solving (D) yields the gradient of the value function, which does not require differentiability

of f . These results rely on the following standard convex duality result, which we state from (Drusvyatskiy, 2020, Theorem 4.1) (Rockafellar and Wets, 1998, Section 11.H).

Theorem 5. For $\mathcal{X} \subset \mathbb{R}^N$ and $\mathcal{U} \subset \mathbb{R}^P$, let $f : \mathcal{X} \times \mathcal{U} \rightarrow \overline{\mathbb{R}}$ be a proper, lower semi-continuous and convex function. Then following are true for all $\mathbf{u} \in \mathcal{U}$:

- (a) **Weak Duality:** $p^{**}(\mathbf{u}) \leq p(\mathbf{u})$.
- (b) **Subdifferential:** If $p(\mathbf{u})$ is finite, then:

$$\partial p(\mathbf{u}) \subset \arg \max_{\mathbf{y} \in \mathbb{R}^P} \langle \mathbf{u}, \mathbf{y} \rangle - f^*(0, \mathbf{y}). \quad (4)$$

If in addition, the inclusion $\mathbf{u} \in \text{ri}(\text{dom } p)$ holds, then (4) holds with equality.

- (c) **Strong Duality:** If the subdifferential $\partial p(\mathbf{u})$ is nonempty, then the equality $p^{**}(\mathbf{u}) = p(\mathbf{u})$ holds and the supremum $p^{**}(\mathbf{u})$ is attained.

Therefore, (D) is key for computing the variation of p with respect to \mathbf{u} . Our goal is reduced to the problem of solving (D), for which the machinery of convex optimization can be invoked to state algorithms and convergence rates. Moreover, in contrast to the automatic differentiation strategy (backpropagation), there is no need to store the iterates, which dramatically reduces the memory requirements.

Let $(\mathbf{y}^{(K)})_{K \in \mathbb{N}}$ be a sequence generated by an algorithm for solving (D), we call the gradient computed by this method the **dual gradient estimator**:

$$\mathbf{g}_4^{(K)}(\mathbf{u}) = \mathbf{y}^{(K)}. \quad (\text{DG})$$

The dual estimator is computationally efficient since it requires solving an optimization problem and does not require computing any additional gradient and Hessian terms. The computational expenses depend only on the method used to solve the problem and the rate of convergence. Such an estimator also does not have a memory overhead like storing the iterates $(\mathbf{y}^{(k)})_{k \in \mathbb{N}}$.

3.1 A Large Class of Parametric Optimization Problems

As an application of our approach, we consider the following class of parametric optimization problem:

$$f(\mathbf{x}, \mathbf{u}) = \langle \mathbf{c}, \mathbf{x} \rangle + h(\mathbf{b} - A\mathbf{x} + \mathbf{u}) + k(\mathbf{x}), \quad (\mathcal{R}_f)$$

where $h : \mathbb{R}^P \rightarrow \overline{\mathbb{R}}$ and $k : \mathbb{R}^N \rightarrow \overline{\mathbb{R}}$ are proper, lower semi-continuous and convex, $A : \mathbb{R}^N \rightarrow \mathbb{R}^P$ is a linear map and $(\mathbf{c}, \mathbf{b}) \in \mathbb{R}^N \times \mathbb{R}^P$. The convex conjugate of f is given by:

$$f^*(\mathbf{v}, \mathbf{y}) = -\langle \mathbf{b}, \mathbf{y} \rangle + k^*(A^*\mathbf{y} - \mathbf{c} + \mathbf{v}) + h^*(\mathbf{y}),$$

which yields the conjugate of the value function as:

$$p^*(\mathbf{y}) = -\langle \mathbf{b}, \mathbf{y} \rangle + k^*(A^*\mathbf{y} - \mathbf{c}) + h^*(\mathbf{y}). \quad (\mathcal{R}_{p^*})$$

Therefore, in order to compute the variation of p with respect to \mathbf{u} by using (4), we must solve a problem of the form:

$$\min_{\mathbf{y} \in \mathbb{R}^P} k^*(A^*\mathbf{y} - \mathbf{c} + \mathbf{v}) + h^*(\mathbf{y}) - \langle \mathbf{b} + \mathbf{u}, \mathbf{y} \rangle. \quad (5)$$

In the following section, depending on the properties of k , h , and A , we provide algorithms and convergence rates for solving (5) and, hence, for approximating the variation of the value function p . As a generic algorithm for solving (5), we mention the Primal–Dual Hybrid Gradient Algorithm by Chambolle and Pock [Chambolle and Pock \(2011\)](#) here. A sufficient condition for uniqueness of solution of (5) is strong convexity of h^* which follows from the Lipschitz continuity of ∇h . For a weaker condition, we state the following result:

Proposition 6. *Let h, k, A and \mathbf{c} in (\mathcal{R}_f) be such that h is differentiable on $\text{int}(\text{dom } h)$ and there exist $(\mathbf{x}, \mathbf{u}) \in \text{dom } k \times \text{int}(\text{dom } h)$ with $A^*\nabla h(\mathbf{u}) - \mathbf{c} \in \partial k(\mathbf{x})$, then $\partial p(\mathbf{u}) = \{\nabla p(\mathbf{u})\}$ is single-valued for all $\mathbf{u} \in \text{ri}(\text{Adom } k + \text{dom } h - \mathbf{b})$.*

Proof. The condition $A^*\nabla h(\mathbf{u}) - \mathbf{c} \in \partial k(\mathbf{x})$ guarantees the existence of some $\mathbf{y} \in \text{dom } h^*$ with $A^*\mathbf{y} - \mathbf{c} \in \text{dom } k^*$. In such case, the expressions $h^*(\mathbf{y})$ and $k^*(A^*\mathbf{y} - \mathbf{c})$ are finite-valued and $\text{dom } p^*$ is non-empty. Since p is proper, lower semi-continuous and convex ([Hoheisel, 2019](#), Theorem 3.101), for every $\mathbf{u} \in \text{ri}(\text{dom } p)$ with $\text{dom } p = \text{Adom } k + \text{dom } h - \mathbf{b}$ ([Rockafellar and Wets, 1998](#), Example 11.41), $\partial p(\mathbf{u})$ is non-empty ([Rockafellar, 1970](#), Theorem 23.4). The single-valuedness of $\partial p(\mathbf{u})$ then follows from the strict convexity of h^* (see Lemma 7(e)). \square

To understand how this works we consider the example where $h = \|\cdot\|_2^2/2$ and $k = \lambda \|\cdot\|_2^2/2 + \gamma \|\cdot\|_1$ for $\lambda > 0$ and $\gamma \geq 0$. By choosing $\mathbf{u} = 0$ and \mathbf{x} as:

$$\mathbf{x}_i = \begin{cases} (-\mathbf{c}_i - \gamma)/\lambda & , \quad \mathbf{c}_i < -\gamma \\ 0 & , -\gamma \leq \mathbf{c}_i \leq \gamma \\ (-\mathbf{c}_i + \gamma)/\lambda & , \quad \gamma < \mathbf{c}_i, \end{cases}$$

we observe that $A^*\nabla h(\mathbf{u}) - \mathbf{c} = -\mathbf{c} \in \partial k(\mathbf{x})$.

Parametric Optimization problems of the form (\mathcal{R}_f) are ubiquitous in Machine Learning, Computer Vision, and Signal Processing. In Signal and Image Processing, the parameter \mathbf{u} represents the observed variable while the mapping A represents the operation performed on the optimal hidden variable \mathbf{x} (which is to

be determined) to obtain the observed variable. In Machine Learning, \mathbf{u} is the target or the label vector, A represents the feature matrix obtained from the independent variable and \mathbf{x} denotes the weights of the mapping to be learned which fits the training set (A, \mathbf{u}) . In this model, h measures the dissimilarity between $A\mathbf{x}$ and \mathbf{u} . The second term puts the penalty on \mathbf{x} and therefore indicates a prior information of the optimal \mathbf{x} which is necessary when $P < N$. In many applications like supervised learning, image denoising and segmentation, $N \leq P$, while in those like compressed sensing and deconvolution, $N > P$.

Other than the above applications, (\mathcal{R}_f) also generalizes the classical infimal convolution ([Rockafellar, 1970](#)). For example, such expressions occur in Image Processing applications in the context of regularization via Total Generalized Variation ([Bredies et al., 2010](#)). Moreover, the Moreau envelope ([Rockafellar and Wets, 1998](#)) of a non-smooth function is of the presented form and is employed for solving non-smooth optimization problems and is key for interpretation of many convex optimization algorithms such as proximal splitting methods ([Lions and Mercier, 1979](#); [Combettes and Pesquet, 2011](#)). Also, the penalty approaches for approximating the minimization of $f(\mathbf{x}) + g(\mathbf{x})$ via $\min_{\mathbf{x}}(f(\mathbf{x}) + \min_{\mathbf{z}} g(\mathbf{z}) + 1/2 \|\mathbf{x} - \mathbf{z}\|^2)$ have the same form, which shows relations to alternating minimization approaches and has been employed in real world machine learning problems ([Laude et al., 2019](#)).

3.2 Rate of Convergence

By invoking convex duality, as described in the previous section, the computation of the value function's variation is reduced to solving problems of type (5) for which a large literature of optimization algorithms is available for several special cases.

We consider the following situations:

- (a) Let (5) be a quadratic problem with matrix $Q \in \mathbb{R}^{P \times P}$ and let $L = \lambda_{\max}(Q)$ and $m = \lambda_{\min}(Q)$, then (DG) computed by using conjugate-gradient method converges like $\mathcal{O}(\omega^K)$ with $\omega := (\sqrt{L} - \sqrt{m})/(\sqrt{L} + \sqrt{m})$ ([Bertsekas, 1999](#), Section 1.6).

Discussion. Depending on whether P is smaller (resp. larger) than N , this rate is better (resp. worse) than those provided in Theorem 1 (see the first column of Figure 2 for a comparison).

- (b) Let h^* be possibly non-smooth with efficiently computable proximal mapping and $k^* \circ A^*$ has an L -Lipschitz continuous gradient then the following are true for solving (5) by using ISTA and accelerated proximal gradient descent (FISTA):

- (DG) converges with ISTA (Chambolle and Pock, 2016, Theorem 4.9) and FISTA (Chambolle and Dossal, 2015, Theorem 3) to $\nabla p(\mathbf{u})$.
- If h^* and $k^* \circ A^*$ are strongly convex with parameters $\delta \geq 0$ and $\gamma \geq 0$ and $\mu = \delta + \gamma > 0$, (DG) converges to $\nabla p(\mathbf{u})$ like $\mathcal{O}(\omega_1^K)$ with ISTA (Chambolle and Pock, 2016, Theorem 4.9) and like $\mathcal{O}(\omega_2^K)$ with FISTA (Chambolle and Pock, 2016, Theorem 4.10) where $\omega_1 = (1 - \tau\gamma)/(1 + \tau\delta)$ and $\omega_2 = 1 - \sqrt{\tau\mu/(1 + \tau\delta)}$.

Discussion. This general setting is beyond the theory that is provided by Theorem 1.

- (c) Let h^* and k be possibly non-smooth and their respective proximal mappings can be computed efficiently, then the following are true for solving (5) by using Primal–Dual Hybrid Gradient Algorithm with $L = \|A^*\|$:

- (DG) converges to $\nabla p(\mathbf{u})$ (Chambolle and Pock, 2016, Theorem 5.1).
- If either h^* or k are strongly convex, then (DG) converges to $\nabla p(\mathbf{u})$ like $\mathcal{O}(1/K^2)$ to $\nabla p(\mathbf{u})$ (Chambolle and Pock, 2011, Theorem 2).
- If h^* and k are both strongly convex with parameters δ and γ respectively, then (DG) converges to $\nabla p(\mathbf{u})$ like $\mathcal{O}(\omega^{K/2})$ to $\nabla p(\mathbf{u})$ with $\omega = (1 + \theta)/(2 + \mu)$ and $\mu = 2\sqrt{\delta\gamma}/L$ (Chambolle and Pock, 2011, Theorem 3).

Discussion. Similarly, this setting is more general than that of Theorem 1.

Remark. For non-strongly convex settings in (5), the sequence of values $p^*(\mathbf{y}^{(K)}) - \langle \mathbf{u}, \mathbf{y}^{(K)} \rangle$ converges like $\mathcal{O}(1/K)$ with ISTA (Chambolle and Pock, 2016, Theorem 4.9) and PDHG (Chambolle and Pock, 2011, Theorem 1) and like $\mathcal{O}(1/K^2)$ with FISTA (Beck and Teboulle, 2009, Theorem 4.4) i.e., we have a potentially accelerated rate of convergence of the objective values. However this rate does not directly translate to a convergence rate of the iterates and hence, to a rate for the convergence of the dual gradient estimator. Such a conclusion requires additional properties of the optimization problem, such as local strong convexity, error bounds, or growth conditions (Attouch and Bolte, 2009; Frankel et al., 2015; Drusvyatskiy and Lewis, 2018).

In order to recognize the potential of the dual gradient approach based on properties of the primal functions in (\mathcal{R}_f) , we trace the conditions for the various convergence rates listed above back to properties of the

primal functions. These results are based on the following lemma. Their proofs can be found in most standard texts on Convex Analysis, e.g., Hiriart-Urruty and Lemaréchal (2012) or Rockafellar (1970).

Lemma 7. Let $g, h : \mathbb{R}^P \rightarrow \bar{\mathbb{R}}$ be proper, lower semi-continuous and convex functions, $\mathbf{d} \in \mathbb{R}^P$ and $B : \mathbb{R}^N \rightarrow \mathbb{R}^P$ be a linear mapping. Let $l : \mathbb{R}^N \rightarrow \bar{\mathbb{R}}$ be defined by $l(\mathbf{x}) = g(B\mathbf{x} + \mathbf{d})$. Then the following results hold:

- If g is m_g -strongly convex on \mathbb{R}^P for $m_g \geq 0$, then l is $\lambda_{\min}(B^*B)m_g$ -strongly convex on \mathbb{R}^N .
- If g and h are strongly convex on \mathbb{R}^P with parameters m_g and m_h respectively, then $g + h$ is $m_g + m_h$ -strongly convex on \mathbb{R}^P .
- If g has an L_g -Lipschitz continuous gradient on \mathbb{R}^P for $L_g \in (0, +\infty)$, then l has a $\lambda_{\max}(B^*B)m_g$ -Lipschitz continuous gradient on \mathbb{R}^N .
- If g and h have Lipschitz continuous gradients on \mathbb{R}^P with parameters L_g and L_h respectively, then $g + h$ has an $L_g + L_h$ -Lipschitz continuous gradient on \mathbb{R}^P .
- If g is differentiable on $\Omega := \text{int dom } g$, then g^* is strictly convex on each convex subset $C \subset \nabla g(\Omega)$.
- g is m_g -strongly convex on \mathbb{R}^P if and only if g^* has a $1/m_g$ -Lipschitz continuous gradient on \mathbb{R}^P .
- g has an L_g -Lipschitz continuous gradient on \mathbb{R}^P if and only if g^* is $1/L_g$ -strongly convex on \mathbb{R}^P .

Let us look at (\mathcal{R}_f) when the regularity conditions given in Theorem 1 are satisfied for f . Let h and k be strongly convex with parameters $m_h > 0$ and $m_k > 0$ respectively and twice differentiable with Lipschitz continuous first and second derivatives. Let L_h and L_k be Lipschitz constants of ∇h and ∇k and let $L_A = \lambda_{\max}(A^*A)$, $m_p = \lambda_{\min}(A^*A)$ and $m_d = \lambda_{\min}(AA^*)$ then $f(\cdot, \mathbf{u})$ is $m_h m_p + m_k$ -strongly convex and has an $L_h L_A + L_k$ -Lipschitz continuous gradient. Using these parameters, the optimal convergence rate for gradient descent is given by $(L - m)/(L + m)$ Polyak (1987), which along with Theorem 1 gives us the rates for the analytical, automatic and implicit estimators as $\mathcal{O}(\omega_p^K)$, $\mathcal{O}(K\omega_p^{2K})$ and $\mathcal{O}(\omega_p^{2K})$ respectively where we have:

$$\omega_p = \frac{(L_h L_A - m_h m_p) + (L_k - m_k)}{(L_h L_A + m_h m_p) + (L_k + m_k)}.$$

The strong convexity parameters of $k^* \circ A^*$ and h^* are m_d/L_k and $1/L_h$. The Lipschitz constants of gradients of these functions are L_A/m_k and $1/m_h$. These

parameters similarly give us the convergence rate for the dual estimator as $\mathcal{O}(\omega_d^K)$ for:

$$\omega_d = \frac{L_h m_h (L_k L_A - m_k m_d) + L_k m_k (L_h - m_h)}{L_h m_h (L_k L_A + m_k m_d) + L_k m_k (L_h + m_h)}.$$

Assuming A is full rank, the convergence rates depend on whether P is larger or smaller than N .

The condition of strong convexity of h or k can be relaxed to non-strong convexity. The expression for convergence rate for primal problem will stay the same with m_h or m_k set to 0. For the dual problem, we make use of the results listed in (b) and (c) to compute the rate. We note that the theoretical guarantees for the primal gradient estimators are difficult to establish beyond the strong convexity and twice continuous differentiability of f in (\mathcal{R}_f) . On the other hand, the dual gradient estimator is quite powerful as it converges in a very broad setting and the convergence rates are theoretically justified.

4 Experiments

We compare the performance of the four different gradient estimators, i.e., (AnG), (AuG), (IG) and (DG) for estimation of $\nabla p(\mathbf{u})$ in different settings. Therefore we fix N and run these methods for different values of P and for different choices of h and k in (\mathcal{R}_f) . Changing P will affect L_A , $\lambda_{\min}(A^*A)$ and $\lambda_{\min}(AA^*)$ while changing h and k will modify L_h , L_k , m_h and m_k . This also includes cases of non-differentiability of k and non-strong convexity of h . For each problem and for each P , we generate error plots of the sequences $\mathbf{g}_i^{(n)}$ for a given \mathbf{u} . Since the convergence rates for methods (AnG), (AuG) and (IG) depend on that of the original sequence, we also show the plots for $\mathbf{x}^{(n)}$ for each of the examples.

We consider following four examples to experimentally verify our observations:

$$\begin{aligned} f_1(\mathbf{x}, \mathbf{u}) &= \frac{1}{2} \|\mathbf{u} - A\mathbf{x}\|_2^2 + \frac{\lambda}{2} \|\mathbf{x}\|_2^2 \\ f_2(\mathbf{x}, \mathbf{u}) &= h_\delta(\mathbf{u} - A\mathbf{x}) + \frac{\lambda}{2} \|\mathbf{x}\|_2^2 \\ f_3(\mathbf{x}, \mathbf{u}) &= \frac{1}{2} \|\mathbf{u} - A\mathbf{x}\|_2^2 + \frac{\lambda}{2} \|\mathbf{x}\|_2^2 + \gamma \|\mathbf{x}\|_1 \\ f_4(\mathbf{x}, \mathbf{u}) &= h_\delta(\mathbf{u} - A\mathbf{x}) + \frac{\lambda}{2} \|\mathbf{x}\|_2^2 + \gamma \|\mathbf{x}\|_1, \end{aligned} \tag{6}$$

where $h_\delta : \mathbb{R}^P \rightarrow \mathbb{R}$ in second and fourth equations in (6) is the Huber function defined by:

$$h_\delta(\mathbf{u}) := \begin{cases} \frac{1}{2} \|\mathbf{u}\|_2^2 & , \quad \|\mathbf{u}\|_2 \leq \delta \\ \delta \left(\|\mathbf{u}\|_2 - \frac{\delta}{2} \right) & , \quad \|\mathbf{u}\|_2 > \delta. \end{cases}$$

The conjugate of the corresponding value functions is given by:

$$\begin{aligned} p_1^*(\mathbf{y}) &= \frac{1}{2\lambda} \|A^T \mathbf{y}\|_2^2 + \frac{1}{2} \|\mathbf{y}\|_2^2 \\ p_2^*(\mathbf{y}) &= \frac{1}{2\lambda} \|A^T \mathbf{y}\|_2^2 + h_\delta^*(\mathbf{y}) \\ p_3^*(\mathbf{y}) &= k^*(A^T \mathbf{y}) + \frac{1}{2} \|\mathbf{y}\|_2^2 \\ p_4^*(\mathbf{y}) &= k^*(A^T \mathbf{y}) + h_\delta^*(\mathbf{y}), \end{aligned}$$

with conjugate of elastic-net term $k := \lambda \|\cdot\|_2^2 + \gamma \|\cdot\|_1$ given by:

$$k^*(\mathbf{v}) = \sum_{i=1}^N \max(0, |v_i| - \gamma)^2 / (2\lambda).$$

For evaluation we set N to 50 and choose P from $\{10, 30, 50, 70, 90\}$. This gives us five different plots for each problem. Three of them are shown in Figure 2 (with each row in the figure corresponding to a $P \in \{10, 50, 90\}$) while rest are given in the supplementary material (with additional two rows corresponding to 30 and 70). We set λ to 2, γ to 0.1 and δ to 0.1. We keep δ small because for sufficiently large values of δ , f_2 will behave like f_1 and f_4 will behave like f_3 . Each element of A and \mathbf{u} is drawn from a normal distribution with mean 0 and standard deviation 1. Thus A is full-rank almost surely. We also scale each column of A differently to introduce ill-conditioning. For all our problems and methods, we use gradient descent when the problem is entirely smooth and proximal gradient descent when the problem has a non-smooth component with an optimal step size $2/(L + m)$. To study the effect of inertia, we additionally employ the Heavy-ball method Polyak (1964) and iPiasco Ochs et al. (2015) with optimal step size $4/(\sqrt{L} + \sqrt{m})^2$ and momentum parameter $(\sqrt{L} - \sqrt{m})^2/(\sqrt{L} + \sqrt{m})^2$ on these problems. The gradients and the Hessians are computed by using the autograd package (Maclaurin et al., 2015). For f_1 , an analytical expression exists for both $\mathbf{x}^*(\mathbf{u})$ and $\nabla p(\mathbf{u})$. In order to compute a good estimate of such terms for the remaining problems, we run the primal and dual problems respectively for a large number of iterations. We verify the correctness of the obtained estimate of $\nabla p(\mathbf{u})$ by comparing it with the numerical gradient computed using central differences. Then we run each algorithm for 250 iterations for each P and f_j and generate the respective plots.

Each cell in Figure 2 displays plots for $\|\mathbf{x}^{(n)}(\mathbf{u}) - \mathbf{x}^*(\mathbf{u})\|_2$ and $\|\mathbf{g}_i^{(n)}(\mathbf{u}) - \nabla p(\mathbf{u})\|_2$ against the number of iterations K . For f_1 (first column), we note that all methods converge since they are backed by theoretical results. e see that for $P > N$

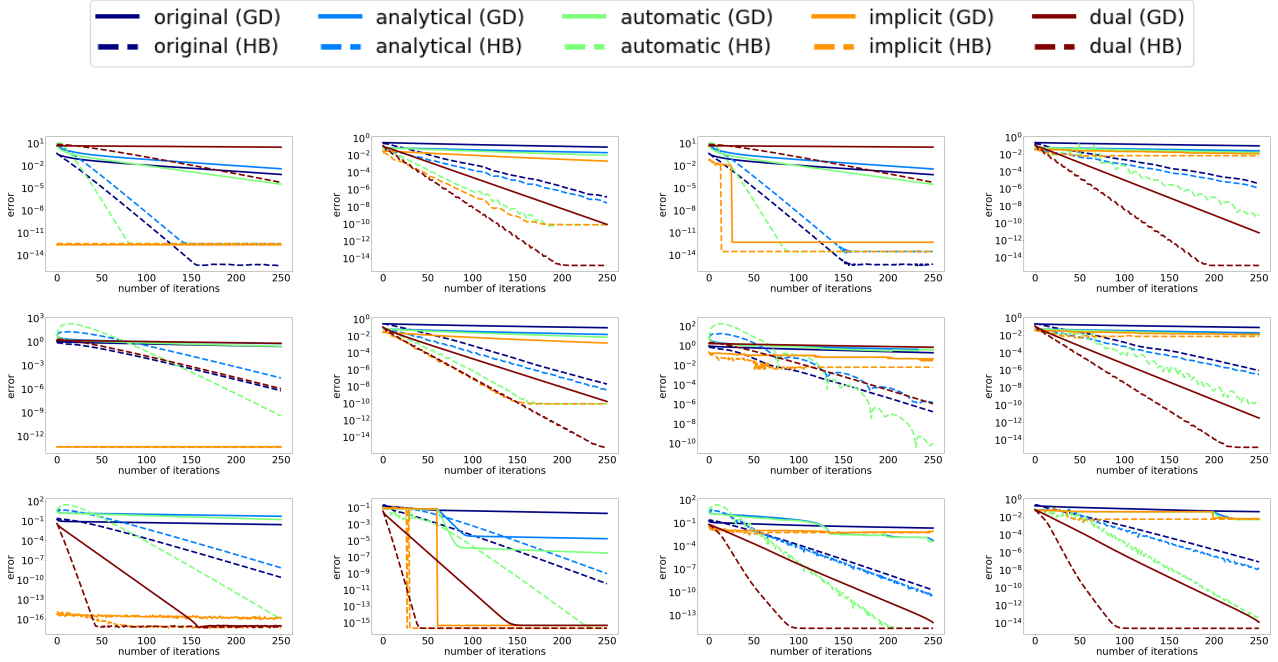


Figure 2: Error plots shown for the gradient sequences computed by the four methods, i.e., (AnG), (AuG), (IG) and (DG), as well as the original (primal variable) sequence computed using (proximal) gradient descent (solid lines) and its inertial variant (dashed lines). The sequences are evaluated on four problems f_1, f_2, f_3 and f_4 (f_j changes from left to right in the given order) for three different values of P , i.e., 90, 50 and 10 (P changes from top to bottom in the given order). Each cell shows the error plots for the ten sequences for a fixed P and f_j ; five for gradient descent or ISTA and five for the Heavy-ball method (Polyak, 1964) or iPiasco (Ochs et al., 2015).

(first column; first row), the dual method is slowest to converge and for $P < N$ (first column; last row), it outperforms the analytical and automatic methods. Since the problem is quadratic therefore the implicit method yields ∇p in one step. For f_2 , the dual method shows faster convergence than all the methods for every choice of P . The remaining two problems (third and forth columns) are not continuously differentiable and therefore (AnG), (AuG) and (IG) show an erratic behavior. The implicit method (orange) performs very poorly in most cases. The analytical (blue) and automatic (green) gradient estimators manage to converge but do so in an irregular manner. Like f_1 , the dual method converges slowly when $P \geq N$ for f_3 and quickly when $P < N$. Similarly, just like f_2 , the performance of the dual method is better than all other methods for every P . The difference between the error plots generated by gradient descent or ISTA (solid lines) and the Heavy-ball method or iPiasco (dashed lines) is also visible. We observe that all the methods benefit from inertia. The fast convergence of automatic method is because of the fact that the acceleration in the convergence of $\mathbf{x}^{(K)}$ is also reflected in that of $D_{\mathbf{u}}\mathbf{x}^{(K)}$ (Ablin et al., 2020; Mehmood and Ochs, 2020). In conclusion, we note

that for the given non-smooth problems, especially f_4 , the dual gradient estimator is not only stable but also performs better than its primal counterparts.

5 Conclusion

The variation of the value function of a parametric optimization problem is desirable in a wide range of Machine Learning and Image Processing applications. The methods for computing this gradient usually rely on directly differentiating the objective and are thus limited to the settings when the objective satisfies strong smoothness conditions. We emphasize that the gradient of the value function can also be computed by using a well-known result from convex duality. This method provides an enormous flexibility for numerical approximation of the value functions derivative, allows to leverage convergence rate results from convex optimization algorithms, and does not rely on differentiability; It can compute a subgradient of the value function.

Acknowledgments

Sheheryar Mehmood and Peter Ochs are supported by the German Research Foundation (DFG Grant OC 150/4-1).

References

- Ablin, P., Peyré, G., and Moreau, T. (2020). Super-efficiency of automatic differentiation for functions defined as a minimum. *arXiv preprint arXiv:2002.03722*.
- Agrawal, A., Amos, B., Barratt, S., Boyd, S., Diamond, S., and Kolter, J. Z. (2019). Differentiable convex optimization layers. In *Advances in Neural Information Processing Systems 32*, pages 9562–9574. Curran Associates, Inc.
- Amos, B. and Kolter, J. Z. (2017). OptNet: Differentiable optimization as a layer in neural networks. In Precup, D. and Teh, Y. W., editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 136–145, International Convention Centre, Sydney, Australia. PMLR.
- Aravkin, A. Y., Burke, J. V., and Friedlander, M. P. (2013). Variational properties of value functions. *SIAM Journal on optimization*, 23(3):1689–1717.
- Attouch, H. and Bolte, J. (2009). On the convergence of the proximal algorithm for nonsmooth functions involving analytic features. *Mathematical Programming*, 116(1–2):5–16.
- Beck, A. and Teboulle, M. (2009). A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202.
- Bertsekas, D. P. (1999). *Nonlinear Programming*. Athena Scientific.
- Bolte, J. and Pauwels, E. (2020). Conservative set valued fields, automatic differentiation, stochastic gradient methods and deep learning. *Mathematical Programming*.
- Bredies, K., Kunisch, K., and Pock, T. (2010). Total generalized variation. *SIAM Journal of Imaging Sciences*, 3(3):492–526.
- Castillo, E., Mínguez, R., and Castillo, C. (2008). Sensitivity analysis in optimization and reliability problems. *Reliability Engineering & System Safety*, 93(12):1788–1800.
- Chambolle, A. and Dossal, C. H. (2015). On the convergence of the iterates of “FISTA”. *Journal of Optimization Theory and Applications*, 166(3):25.
- Chambolle, A. and Pock, T. (2011). A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145.
- Chambolle, A. and Pock, T. (2016). An introduction to continuous optimization for imaging. *Acta Numerica*, 25:161–319.
- Chen, R. T. Q., Rubanova, Y., Bettencourt, J., and Duvenaud, D. K. (2018). Neural ordinary differential equations. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R., editors, *Advances in Neural Information Processing Systems 31*, pages 6571–6583. Curran Associates, Inc.
- Christianson, B. (1994). Reverse accumulation and attractive fixed points. *Optimization Methods and Software*, 3(4):311–326.
- Combettes, P. L. and Pesquet, J.-C. (2011). *Proximal Splitting Methods in Signal Processing*, pages 185–212. Springer New York, New York, NY.
- Deledalle, C.-A., Vaiter, S., Fadili, J., and Peyré, G. (2014). Stein unbiased gradient estimator of the risk (sugar) for multiple parameter selection. *SIAM Journal on Imaging Sciences*, 7(4):2448–2487.
- Dempe, S., Kalashnikov, V., Prez-Valds, G. A., and Kalashnykova, N. (2015). *Bilevel Programming Problems: Theory, Algorithms and Applications to Energy Networks*. Springer Publishing Company, Incorporated.
- Domke, J. (2012). Generic methods for optimization based modeling. In *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics*, volume 22 of *Proceedings of Machine Learning Research*, pages 318–326, La Palma, Canary Islands. PMLR.
- Drusvyatskiy, D. (2020). Convex analysis and non-smooth optimization. University Lecture. ”URL: https://sites.math.washington.edu/~ddrusv/crs/Math_516_2020/bookwithindex.pdf.
- Drusvyatskiy, D. and Lewis, A. S. (2018). Error bounds, quadratic growth, and linear convergence of proximal methods. *Mathematics of Operations Research*, 43(3):919–948.
- Frankel, P., Garrigos, G., and Peypouquet, J. (2015). Splitting methods with variable metric for kurdyka-łojasiewicz functions and general convergence rates. *Journal of Optimization Theory and Applications*, 165(3):874–900.
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative adversarial nets. In *Proceedings of the 27th International Conference on*

- Neural Information Processing Systems - Volume 2*, NIPS'14, page 2672–2680, Cambridge, MA, USA. MIT Press.
- Griewank, A. and Faure, C. (2003). Piggyback differentiation and optimization. In Biegler, L. T., Heinkenschloss, M., Ghattas, O., and van Bloemen Waanders, B., editors, *Large-Scale PDE-Constrained Optimization*, pages 148–164, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Hiriart-Urruty, J.-B. and Lemaréchal, C. (2012). *Fundamentals of convex analysis*. Springer Science & Business Media.
- Hoheisel, T. (2019). Topics in convex analysis in matrix space. "URL: <https://www.math.mcgill.ca/hoheisel/Paseky2019.pdf>.
- Knopoff, D., Fernandez, D., Torres, G., and Cristina, T. (2012). Adjoint method for a tumor growth pde-constrained optimization problem. *Computers and Mathematics with Applications*, 66:1104–1119.
- Kunisch, K. and Pock, T. (2013). A bilevel optimization approach for parameter learning in variational models. *SIAM Journal on Imaging Sciences*, 6(2):938–983.
- Laude, E., Wu, T., and Cremers, D. (2019). Optimization of inf-convolution regularized nonconvex composite problems. In Chaudhuri, K. and Sugiyama, M., editors, *Proceedings of Machine Learning Research*, volume 89 of *Proceedings of Machine Learning Research*, pages 547–556. PMLR.
- Lions, P.-L. and Mercier, B. (1979). Splitting algorithms for the sum of two nonlinear operators. *SIAM Journal on Numerical Analysis*, 16(6):964–979.
- Maclaurin, D., Duvenaud, D., and Adams, R. P. (2015). Autograd: Effortless gradients in numpy. In *ICML 2015 AutoML Workshop*, volume 238, page 5.
- Mairal, J., Bach, F., Ponce, J., and Sapiro, G. (2010). Online learning for matrix factorization and sparse coding. *Journal of Machine Learning Research*, 11:19–60.
- Mehmood, S. and Ochs, P. (2020). Automatic differentiation of some first-order methods in parametric optimization. In Chiappa, S. and Calandra, R., editors, *The 23rd International Conference on Artificial Intelligence and Statistics, AISTATS 2020, 26-28 August 2020, Online [Palermo, Sicily, Italy]*, volume 108 of *Proceedings of Machine Learning Research*, pages 1584–1594. PMLR.
- Ochs, P., Brox, T., and Pock, T. (2015). ipiasco: Inertial proximal algorithm for strongly convex optimization. *Journal of Mathematical Imaging and Vision*, 53(2):171–181.
- Plessix, R.-E. (2006). A review of the adjoint-state method for computing the gradient of a functional with geophysical applications. *Geophysical Journal International*, 167(2):495–503.
- Pollini, N., Lavan, O., and Amir, O. (2018). Adjoint sensitivity analysis and optimization of hysteretic dynamic systems with nonlinear viscous dampers. *Structural and Multidisciplinary Optimization*, 57(6):2273–2289.
- Polyak, B. T. (1964). Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 4(5):1–17.
- Polyak, B. T. (1987). *Introduction to Optimization*. Optimization Software.
- Pontryagin, L. S., Mishchenko, E. F., Boltyanskii, V. G., and Gamkrelidze, R. V. (1962). *Mathematical theory of optimal processes*. Interscience, New York.
- Rockafellar, R. and Wets, R. J.-B. (1998). *Variational Analysis*. Springer Verlag, Heidelberg, Berlin, New York.
- Rockafellar, R. T. (1970). *Convex analysis*. Princeton Mathematical Series. Princeton University Press, Princeton, N. J.
- Still, G. (2018). Lectures on parametric optimization: An introduction. *Optimization Online*.
- Taskar, B., Guestrin, C., and Koller, D. (2004). Max-margin markov networks. In Thrun, S., Saul, L. K., and Schölkopf, B., editors, *Advances in Neural Information Processing Systems 16*, pages 25–32. MIT Press.
- Tsochantaridis, I., Joachims, T., Hofmann, T., and Altun, Y. (2005). Large margin methods for structured and interdependent output variables. *Journal of Machine Learning Research*, 6:1453–1484.
- van den Berg, E. and Friedlander, M. P. (2008). Probing the pareto frontier for basis pursuit solutions. *SIAM Journal on Scientific Computing*, 31(2):890–912.
- van den Berg, E. and Friedlander, M. P. (2011). Sparse optimization with least-squares constraints. *SIAM Journal on Optimization*, 21(4):1201–1229.