
Implicit differentiation of Lasso-type models for hyperparameter optimization

Quentin Bertrand^{*1} Quentin Klopfenstein^{*2} Mathieu Blondel³ Samuel Vaiter⁴ Alexandre Gramfort¹
Joseph Salmon⁵

Abstract

Setting regularization parameters for Lasso-type estimators is notoriously difficult, though crucial for obtaining the best accuracy. The most popular hyperparameter optimization approach is grid-search on a held-out dataset. However, grid-search requires to choose a predefined grid of parameters and scales exponentially in the number of parameters. Another class of approaches casts hyperparameter optimization as a bi-level optimization problem, typically solved by gradient descent. The key challenge for these approaches is the estimation of the gradient w.r.t. the hyperparameters. Computing that gradient via forward or backward automatic differentiation usually suffers from high memory consumption, while implicit differentiation typically involves solving a linear system which can be prohibitive and numerically unstable. In addition, implicit differentiation usually assumes smooth loss functions, which is not the case of Lasso-type problems. This work introduces an efficient implicit differentiation algorithm, without matrix inversion, tailored for Lasso-type problems. Our proposal scales to high-dimensional data by leveraging the sparsity of the solutions. Empirically, we demonstrate that the proposed method outperforms a large number of standard methods for hyperparameter optimization.

1. Introduction

In many statistical applications, the number of parameters p is much larger than the number of observations n . In such scenarios, a popular approach to tackle linear regression problems is to consider convex ℓ_1 -type penalties, used in Lasso (Tibshirani, 1996), Group-Lasso (Yuan and Lin, 2006), Elastic-Net (Zou and Hastie, 2005) or adaptive Lasso (Zou, 2006). These *Lasso-type* estimators rely on regularization hyperparameters, trading data fidelity against sparsity. Unfortunately, setting these hyperparameters is hard in practice: estimators based on ℓ_1 -type penalties are indeed more sensitive to the choice of hyperparameters than ℓ_2 regularized estimators.

To control for overfitting, it is customary to use different datasets for model training (*i.e.*, computing the regression coefficients) and hyperparameter selection (*i.e.*, choosing the best regularization parameters). A *metric*, *e.g.*, *hold-out loss*, is optimized on a validation dataset (Stone and Ramer, 1965). Alternatively one can rely on a statistical criteria that penalizes complex models such as AIC/BIC (Liu et al., 2011) or SURE (Stein Unbiased Risk Estimator, Stein 1981). In all cases, hyperparameters are tuned to optimize a chosen metric.

The canonical hyperparameter optimization method is *grid-search*. It consists in fitting and selecting the best model over a predefined grid of parameter values. The complexity of grid-search is exponential with the number of hyperparameters, making it only competitive when the number of hyperparameters is small. Other hyperparameter selection strategies include *random search* (Bergstra and Bengio, 2012) and Bayesian optimization (Brochu et al., 2010; Snoek et al., 2012) that aims to learn an approximation of the metric over the parameter space and rely on an exploration policy to find the optimum.

Another line of work for hyperparameter optimization (HO) relies on gradient descent in the hyperparameter space. This strategy has been widely explored for smooth objective functions (Larsen et al., 1996; Bengio, 2000; Larsen et al., 2012). The main challenge for this class of methods is estimating the gradient *w.r.t.* the hyperparameters. Gradient estimation techniques are mostly divided in two categories. *Implicit differentiation* requires the exact

^{*}Equal contribution ¹Université Paris-Saclay, Inria, CEA, Palaiseau, France ²Institut Mathématique de Bourgogne, Université de Bourgogne, Dijon, France ³Google Research, Brain team, Paris, France ⁴CNRS and Institut Mathématique de Bourgogne, Université de Bourgogne, Dijon, France ⁵IMAG, Université de Montpellier, CNRS, Montpellier, France. Correspondence to: Quentin Bertrand <quentin.bertrand@inria.fr>, Quentin Klopfenstein <quentin.klopfenstein@u-bourgogne.fr>.

solution of the optimization problem and involves the resolution of a linear system (Bengio, 2000). This can be expensive to compute and lead to numerical instabilities, especially when the system is ill-conditioned (Lorraine et al., 2019). Alternatively, *iterative differentiation* computes the gradient using the iterates of an optimization algorithm. Backward iterative differentiation (Domke, 2012) is computationally efficient when the number of hyperparameters is large. However it is memory consuming since it requires storing all intermediate iterates. In contrast, forward iterative differentiation (Deledalle et al., 2014; Franceschi et al., 2017) does not require storing the iterates but can be computationally expensive with a large number of hyperparameters; see Baydin et al. (2018) for a survey.

This article proposes to investigate the use of these methods to set the regularization hyperparameters in an automatic fashion for Lasso-type problems. To cover the cases of both low and high number of hyperparameters, two estimators are investigated, namely the Lasso and the weighted Lasso which have respectively one or as many parameters as features. Our contributions are as follows:

- We show that forward iterative differentiation of block coordinate descent (BCD), a state-of-the-art solver for Lasso-type problems, converges towards the true gradient. Crucially, we show that this scheme converges linearly once the support is identified and that its limit does **not** depend of the initial starting point.
- These results lead to the proposed algorithm (Algorithm 2) where the computation of the Jacobian is **decoupled** from the computation of the regression coefficients. The later can be done with state-of-the-art convex solvers, and interestingly, it does not require solving a linear system, potentially ill-conditioned.
- We show through an extensive benchmark on simulated and real high dimensional data that the proposed method outperforms state-of-the-art HO methods.

Our work is somewhat similar to Gregor and LeCun (2010); Xin et al. (2016); Borgerding et al. (2017); Liu et al. (2018); Wu et al. (2019), where the *solution* is differentiated w.r.t. optimization parameters instead of the regularization parameter. However the goal is very different as they want to accelerate the optimization algorithm whereas we provide an efficient algorithm to compute the gradient.

Notation The design matrix is $X \in \mathbb{R}^{n \times p}$ (corresponding to n samples and p features) and the observation vector is $y \in \mathbb{R}^n$. The regularization parameter, possibly multivariate, is denoted by $\lambda = (\lambda_1, \dots, \lambda_r)^\top \in \mathbb{R}^r$. We denote $\hat{\beta}^{(\lambda)} \in \mathbb{R}^p$ the regression coefficients associated to λ . We denote $\hat{\mathcal{J}}_{(\lambda)} \triangleq (\nabla_{\lambda} \hat{\beta}_1^{(\lambda)}, \dots, \nabla_{\lambda} \hat{\beta}_p^{(\lambda)})^\top \in \mathbb{R}^{p \times r}$ the weak

Jacobian (Evans and Gariepy, 1992) of $\hat{\beta}^{(\lambda)}$ w.r.t. λ . For a function $\psi : \mathbb{R}^p \times \mathbb{R}^r \rightarrow \mathbb{R}$ with weak derivatives of order two, we denote by $\nabla_{\beta} \psi(\beta, \lambda) \in \mathbb{R}^p$ (resp. $\nabla_{\lambda}(\beta, \lambda) \in \mathbb{R}^r$) its weak gradient w.r.t. the first parameter (resp. the second parameter). The weak Hessian $\nabla^2 \psi(\beta, \lambda)$ is a matrix in $\mathbb{R}^{(p+r) \times (p+r)}$ which has a block structure

$$\nabla^2 \psi(\beta, \lambda) = \begin{pmatrix} \nabla_{\beta}^2 \psi(\beta, \lambda) & \nabla_{\beta, \lambda}^2 \psi(\beta, \lambda) \\ \nabla_{\lambda, \beta}^2 \psi(\beta, \lambda) & \nabla_{\lambda}^2 \psi(\beta, \lambda) \end{pmatrix}.$$

The support of $\hat{\beta}^{(\lambda)}$ (the indices of non-zero coefficients) is denoted by $\hat{S}^{(\lambda)}$, and $\hat{s}^{(\lambda)}$ represents its cardinality (i.e., the number of non-zero coefficients). The sign vector $\text{sign} \hat{\beta}^{(\lambda)} \in \mathbb{R}^p$ is the vector of component-wise signs (with the convention that $\text{sign}(0) = 0$) of $\hat{\beta}^{(\lambda)}$. Note that to ease the reading, we drop λ in the notation when it is clear from the context and use $\hat{\beta}$, $\hat{\mathcal{J}}$, \hat{S} and \hat{s} . The Mahalanobis distance of a vector $x \in \mathbb{R}^p$ and a matrix $A \succ 0$ is noted $\|x\|_A \triangleq \sqrt{x^\top A^{-1} x}$.

2. Background

2.1. Problem setting

To favor sparse coefficients, we consider Lasso-type estimators based on non-smooth regularization functions. Such problems consist in finding:

$$\hat{\beta}^{(\lambda)} \in \arg \min_{\beta \in \mathbb{R}^p} \psi(\beta, \lambda). \quad (1)$$

The Lasso (Tibshirani, 1996) is recovered, with the number of hyperparameters set to $r = 1$:

$$\psi(\beta, \lambda) = \frac{1}{2n} \|y - X\beta\|_2^2 + e^\lambda \|\beta\|_1, \quad (2)$$

while the weighted Lasso (wLasso, Zou 2006, introduced to reduce the bias of the Lasso) has $r = p$ hyperparameters and reads:

$$\psi(\beta, \lambda) = \frac{1}{2n} \|y - X\beta\|_2^2 + \sum_{j=1}^p e^{\lambda_j} |\beta_j|. \quad (3)$$

Note that we adopt the hyperparameter parametrization of Pedregosa (2016), i.e., we write the regularization parameter as e^λ . This avoids working with a positivity constraint in the optimization process and fixes scaling issues in the line search. It is also coherent with the usual choice of a geometric grid for grid-search (Friedman et al., 2010).

Remark 1. Other formulations could be investigated like Elastic-Net or non-convex formulation, e.g., MCP (Zhang, 2010). Our theory does not cover non-convex cases, though we illustrate that it behaves properly numerically. Handling such non-convex cases is left as a question for future work.

The HO problem can be expressed as a nested *bi-level optimization* problem. For a given differentiable criterion $\mathcal{C} : \mathbb{R}^p \mapsto \mathbb{R}$ (e.g., hold-out loss or SURE), it reads:

$$\begin{aligned} \arg \min_{\lambda \in \mathbb{R}^r} \left\{ \mathcal{L}(\lambda) \triangleq \mathcal{C} \left(\hat{\beta}(\lambda) \right) \right\} \\ \text{s.t. } \hat{\beta}(\lambda) \in \arg \min_{\beta \in \mathbb{R}^p} \psi(\beta, \lambda) . \end{aligned} \quad (4)$$

Note that SURE itself is not necessarily weakly differentiable *w.r.t.* $\hat{\beta}(\lambda)$. However a weakly differentiable approximation can be constructed (Ramani et al., 2008; Deledalle et al., 2014). Under the hypothesis that Problem (1) has a unique solution for every $\lambda \in \mathbb{R}^r$, the function $\lambda \mapsto \hat{\beta}(\lambda)$ is weakly differentiable (Vaiteer et al., 2013). Using the chain rule, the gradient of \mathcal{L} *w.r.t.* λ then writes:

$$\nabla_{\lambda} \mathcal{L}(\lambda) = \hat{\mathcal{J}}_{(\lambda)}^{\top} \nabla \mathcal{C} \left(\hat{\beta}(\lambda) \right) . \quad (5)$$

Computing the weak Jacobian $\hat{\mathcal{J}}_{(\lambda)}$ of the inner problem is the main challenge, as once the *hypergradient* $\nabla_{\lambda} \mathcal{L}(\lambda)$ has been computed, one can use usual gradient descent, $\lambda^{(t+1)} = \lambda^{(t)} - \rho \nabla_{\lambda} \mathcal{L}(\lambda^{(t)})$, for a step size $\rho > 0$. Note however that \mathcal{L} is usually non-convex and convergence towards a global minimum is not guaranteed. In this work, we propose an efficient algorithm to compute $\hat{\mathcal{J}}_{(\lambda)}$ for Lasso-type problems, relying on improved forward differentiation.

2.2. Implicit differentiation (smooth case)

Implicit differentiation, which can be traced back to Larsen et al. (1996), is based on the knowledge of $\hat{\beta}$ and requires solving a $p \times p$ linear system (Bengio, 2000, Sec. 4). Since then, it has been extensively applied in various contexts. Chapelle et al. (2002); Seeger (2008) used implicit differentiation to select hyperparameters of kernel-based models. Kunisch and Pock (2013) applied it to image restoration. Pedregosa (2016) showed that each inner optimization problem could be solved only approximately, leveraging noisy gradients. Related to our work, Foo et al. (2008) applied implicit differentiation on a “weighted” Ridge-type estimator (*i.e.*, a Ridge penalty with one λ_j per feature).

Yet, all the aforementioned methods have a common drawback : they are limited to the smooth setting, since they rely on optimality conditions for smooth optimization. They proceed as follows: if $\beta \mapsto \psi(\beta, \lambda)$ is a smooth convex function (for any fixed λ) in Problem (1), then for all λ , the solution $\hat{\beta}(\lambda)$ satisfies the following fixed point equation:

$$\nabla_{\beta} \psi \left(\hat{\beta}(\lambda), \lambda \right) = 0 . \quad (6)$$

Then, this equation can be differentiated *w.r.t.* λ :

$$\nabla_{\beta, \lambda}^2 \psi \left(\hat{\beta}(\lambda), \lambda \right) + \hat{\mathcal{J}}_{(\lambda)}^{\top} \nabla_{\beta}^2 \psi \left(\hat{\beta}(\lambda), \lambda \right) = 0 . \quad (7)$$

Assuming that $\nabla_{\beta}^2 \psi \left(\hat{\beta}(\lambda), \lambda \right)$ is invertible this leads to a closed form solution for the weak Jacobian $\hat{\mathcal{J}}_{(\lambda)}$:

$$\hat{\mathcal{J}}_{(\lambda)}^{\top} = -\nabla_{\beta, \lambda}^2 \psi \left(\hat{\beta}(\lambda), \lambda \right) \underbrace{\left(\nabla_{\beta}^2 \psi \left(\hat{\beta}(\lambda), \lambda \right) \right)^{-1}}_{p \times p} , \quad (8)$$

which in practice is computed by solving a linear system. Unfortunately this approach cannot be generalized for non-smooth problems since Equation (6) no longer holds.

2.3. Implicit differentiation (non-smooth case)

Related to our work Mairal et al. (2012) used implicit differentiation with respect to the dictionary ($X \in \mathbb{R}^{n \times p}$) on Elastic-Net models to perform dictionary learning. Regarding Lasso problems, the literature is quite scarce, see (Dossal et al., 2013; Zou et al., 2007) and (Vaiteer et al., 2013; Tibshirani and Taylor, 2011) for a more generic setting encompassing weighted Lasso. General methods for gradient estimation of non-smooth optimization schemes exist (Vaiteer et al., 2017) but are not practical since they depend on a possibly ill-posed linear system to invert. Amos and Kolter (2017) have applied implicit differentiation on estimators based on quadratic objective function with linear constraints, whereas Niculae and Blondel (2017) have used implicit differentiation on a smooth objective function with simplex constraints. However none of these approaches leverages the sparsity of Lasso-type estimators.

3. Hypergradients for Lasso-type problems

To tackle hyperparameter optimization of non-smooth Lasso-type problems, we propose in this section an efficient algorithm for hypergradient estimation. Our algorithm relies on implicit differentiation, thus enjoying low-memory cost, yet does not require to naively solve a (potentially ill-conditioned) linear system of equations. In the sequel, we assume access to a (weighted) Lasso solver, such as ISTA (Daubechies et al., 2004) or Block Coordinate Descent (BCD, Tseng and Yun 2009, see also Algorithm 5).

3.1. Implicit differentiation

Our starting point is the key observation that Lasso-type solvers induce a fixed point iteration that we can leverage to compute a Jacobian. Indeed, proximal BCD algorithms (Tseng and Yun, 2009), consist in a local gradient step composed with a soft-thresholding step (ST), *e.g.*, for the Lasso, for $j \in 1, \dots, p$:

$$\beta_j \leftarrow \text{ST} \left(\beta_j - \frac{X_{:,j}^{\top} (X\beta - y)}{\|X_{:,j}\|^2}, \frac{ne^{\lambda}}{\|X_{:,j}\|^2} \right) \quad (9)$$

where $\text{ST}(t, \tau) = \text{sign}(t) \cdot (|t| - \tau)_+$ for any $t \in \mathbb{R}$ and $\tau \geq 0$ (extended for vectors component-wise). The solution of

the optimization problem satisfies, for any $\alpha > 0$, the fixed-point equation (Combettes and Wajs, 2005, Prop. 3.1), for $j \in 1, \dots, p$:

$$\hat{\beta}_j^{(\lambda)} = \text{ST} \left(\hat{\beta}_j^{(\lambda)} - \frac{1}{\alpha} X_{j,:}^\top (X \hat{\beta}^{(\lambda)} - y), \frac{ne^\lambda}{\alpha} \right). \quad (10)$$

The former can be differentiated *w.r.t.* λ , see Lemma A.1 in Appendix, leading to a closed form solution for the Jacobian $\mathcal{J}_{(\lambda)}$ of the Lasso and the weighted Lasso.

Proposition 1 (Adapting Vaiter et al. 2013, Thm. 1). Let \hat{S} be the support of the vector $\hat{\beta}^{(\lambda)}$. Suppose that $X_{\hat{S}}^\top X_{\hat{S}} \succ 0$, then a weak Jacobian $\hat{\mathcal{J}} = \hat{\mathcal{J}}_{(\lambda)}$ of the Lasso writes:

$$\hat{\mathcal{J}}_{\hat{S}} = -ne^\lambda (X_{\hat{S}}^\top X_{\hat{S}})^{-1} \text{sign} \hat{\beta}_{\hat{S}}, \quad (11)$$

$$\hat{\mathcal{J}}_{\hat{S}^c} = 0, \quad (12)$$

and for the weighted Lasso:

$$\hat{\mathcal{J}}_{\hat{S}, \hat{S}} = -(X_{\hat{S}}^\top X_{\hat{S}})^{-1} \text{diag}(ne^{\lambda_{\hat{S}}} \odot \text{sign} \hat{\beta}_{\hat{S}}) \quad (13)$$

$$\hat{\mathcal{J}}_{j_1, j_2} = 0 \quad \text{if } j_1 \notin \hat{S} \text{ or if } j_2 \notin \hat{S}. \quad (14)$$

The proof of Proposition 1 can be found in Appendix A.1. Note that the positivity condition in Proposition 1 is satisfied if the (weighted) Lasso has a unique solution. Moreover, even for multiple solutions cases, there exists at least one satisfying the positivity condition (Vaiter et al., 2013).

Proposition 1 shows that the Jacobian of the weighted Lasso $\hat{\mathcal{J}}_{(\lambda)} \in \mathbb{R}^{p \times p}$ is row and column sparse. This is key for algorithmic efficiency. Indeed, *a priori*, one has to store a possibly dense $p \times p$ matrix, which is prohibitive when p is large. Proposition 1 leads to a simple algorithm (see Algorithm 1) to compute the Jacobian in a *cheap* way, as it *only* requires storing and inverting an $\hat{s} \times \hat{s}$ matrix. Even if the linear system to solve is of size $\hat{s} \times \hat{s}$, instead of $p \times p$ for smooth objective function, the system to invert can be ill-conditioned, especially when a large support size \hat{s} is encountered. This leads to numerical instabilities and slows down the resolution (see an illustration in Figure 2). Forward (Algorithm 3 in Appendix) and backward (Algorithm 4 in Appendix) iterative differentiation, which do not require solving linear systems, can overcome these issues.

3.2. Link with iterative differentiation

Iterative differentiation in the field of hyperparameter setting can be traced back to Domke (2012) who derived a backward differentiation algorithm for gradient descent, heavy ball and L-BFGS algorithms applied to smooth loss functions. Agrawal et al. (2019) generalized it to a specific subset of convex programs. Maclaurin et al. (2015) derived a backward differentiation for stochastic gradient

Algorithm 1 IMPLICIT DIFFERENTIATION

input : $X \in \mathbb{R}^{n \times p}$, $y \in \mathbb{R}^n$, $\lambda \in \mathbb{R}$, $n_{\text{iter}} \in \mathbb{N}$
 // jointly compute coef. and Jacobian
if Lasso **then**
 Get $\hat{\beta} = \text{Lasso}(X, y, \lambda, n_{\text{iter}})$ and its support \hat{S} .
 $\hat{\mathcal{J}} = 0_p$
 $\hat{\mathcal{J}}_{\hat{S}} = -ne^\lambda (X_{\hat{S}}^\top X_{\hat{S}})^{-1} \text{sign} \hat{\beta}_{\hat{S}}$
if wLasso **then**
 Get $\hat{\beta} = \text{wLasso}(X, y, \lambda, n_{\text{iter}})$ and its support \hat{S} .
 $\hat{\mathcal{J}} = 0_{p \times p}$
 $\hat{\mathcal{J}}_{\hat{S}, \hat{S}} = -(X_{\hat{S}}^\top X_{\hat{S}})^{-1} \text{diag}(ne^{\lambda_{\hat{S}}} \odot \text{sign} \hat{\beta}_{\hat{S}})$
return $\hat{\beta}, \hat{\mathcal{J}}$

descent. On the other hand Deledalle et al. (2014) used forward differentiation of (accelerated) proximal gradient descent for hyperparameter optimization with non-smooth penalties. Franceschi et al. (2017) proposed a benchmark of forward mode versus backward mode, varying the number of hyperparameters to learn. Frecon et al. (2018) cast the problem of inferring the groups in a group-Lasso model as a bi-level optimization problem and solved it using backward differentiation.

Forward differentiation consists in differentiating each step of the algorithm (*w.r.t.* λ in our case). For the Lasso solved with BCD it amounts differentiating Equation (9), and leads to the following recursive equation for the Jacobian, for $j \in 1, \dots, p$, with $z_j = \beta_j - X_{:,j}^\top (X\beta - y) / \|X_{:,j}\|^2$:

$$\mathcal{J}_j \leftarrow \partial_1 \text{ST} \left(z_j, \frac{ne^\lambda}{\|X_{:,j}\|^2} \right) \left(\mathcal{J}_j - \frac{1}{\|X_{:,j}\|^2} X_{:,j}^\top X \mathcal{J} \right) + \partial_2 \text{ST} \left(z_j, \frac{ne^\lambda}{\|X_{:,j}\|^2} \right) \frac{ne^\lambda}{\|X_{:,j}\|^2}, \quad (15)$$

see Algorithm 3 (in Appendix) for full details. Our proposed algorithm uses the fact that after a finite number of epochs $\partial_1 \text{ST}(z_j, ne^\lambda / \|X_{:,j}\|^2)$ and $\partial_2 \text{ST}(z_j, ne^\lambda / \|X_{:,j}\|^2)$ are **constant** (they no longer depends on the current β). Indeed, the sign of $\hat{\beta}$ is identified after a finite number of iterations thus the partial derivatives are constant. It is then possible to **decouple** the computation of the Jacobian by only solving Problem (1) in a first step and then apply the forward differentiation recursion steps, see Algorithm 2. This can be seen as the forward counterpart in a non-smooth case of the recent paper Loraine et al. (2019). An additional benefit of such updates is that they can be restricted to the (current) support, which leads to faster Jacobian computation.

We now show that the Jacobian computed using forward differentiation and our method, Algorithm 2, converges toward the true Jacobian.

Proposition 2. Assuming the Lasso solution (Problem (2)) (or weighted Lasso Problem (3)) is unique, then Algorithms 2 and 3 converge toward the Jacobian $\hat{\mathcal{J}}$ defined in Proposition 1. Algorithm 3 computes the Jacobian along with the regression coefficients, once the support has been identified, the Jacobian converges linearly. Algorithm 2 computes first the coefficients $\hat{\beta}$ and then the Jacobian $\hat{\mathcal{J}}$, provided that the support has been identified in the first step, the convergence is linear in the second, with the same rate as Algorithm 3:

$$\|\mathcal{J}_{\hat{s}}^{(k+1)} - \hat{\mathcal{J}}\|_{(X_{:, \hat{s}}^\top X_{:, \hat{s}})^{-1}} \leq C^k \|\mathcal{J}_{\hat{s}}^{(k)} - \hat{\mathcal{J}}\|_{(X_{:, \hat{s}}^\top X_{:, \hat{s}})^{-1}}$$

where $C = \|A^{(j_{\hat{s}})} \dots A^{(j_1)}\|_2 < 1$, $j_1, \dots, j_{\hat{s}}$ are the indices of the support of $\hat{\beta}$ in increasing order and

$$A^{(j_s)} = \text{Id}_{\hat{s}} - \frac{\left(X_{:, \hat{s}}^\top X_{:, \hat{s}}\right)^{1/2}}{\|X_{:, j_s}\|} \frac{\left(X_{:, \hat{s}}^\top X_{:, \hat{s}}\right)^{1/2}}{\|X_{:, j_s}\|} \in \mathbb{R}^{\hat{s} \times \hat{s}}.$$

□

Proof of Proposition 2 can be found in Appendix A.2 and A.3.

Remark 3. Uniqueness. As proved in Tibshirani (2013, Lem. 3 and 4) the set of (pathological) lambdas where the Lasso solution is not unique is typically empty. Moreover if the Lasso solution is not unique, there could be a non-continuous solution path $\lambda \mapsto \hat{\beta}^{(\lambda)}$, leaving only non-gradient based methods available. Even if Proposition 2 does not provide theoretical guarantees in such a pathological setting, one can still apply Algorithms 2 and 3, see Appendix E.1 for experiments in this settings.

Remark 4. Rate for the backward differentiation. The backward and forward differentiation compute the same quantity: $\nabla_{\lambda} \mathcal{L}(\lambda)$, but the backward differentiation directly computes the product given in Equation (5) leading to the gradient of $\mathcal{L}(\lambda)$. Proposition 2 provides rates for the convergence of the Jacobian \mathcal{J} which leads to rates for the gradient *i.e.*, for the backward algorithm as well.

As an illustration, Figure 1 shows the times of computation of a single gradient $\nabla_{\lambda} \mathcal{L}(\lambda)$ and the distance to “optimum” of this gradient as a function of the number of iterations in the inner optimization problem for the forward iterative differentiation (Algorithm 3), the backward iterative differentiation (Algorithm 4), and the proposed algorithm (Algorithm 2). The backward iterative differentiation is several order of magnitude slower than the forward and our implicit forward method. Moreover, once the support has been identified (after 20 iterations) the proposed implicit forward method converges faster than other methods. Note also that in Propositions 1 and 2 the Jacobian for the




Figure 1. Time to compute a single gradient (Synthetic data, Lasso, $n, p = 1000, 2000$). Influence on the number of iterations of BCD (in the inner optimization problem of Problem (4)) on the computation time (left) and the distance to “optimum” of the gradient $\nabla_{\lambda} \mathcal{L}(\lambda)$ (right) for the Lasso estimator. The “optimum” is here the gradient given by implicit differentiation (Algorithm 1).

Lasso only depends on the *support* (*i.e.*, the indices of the non-zero coefficients) of the regression coefficients $\hat{\beta}^{(\lambda)}$. In other words, once the support of $\hat{\beta}^{(\lambda)}$ is correctly identified, even if the value of the non-zero coefficients are not correctly estimated, the Jacobian is exact, see Sun et al. (2019) for support identification guarantees.

4. Experiments

Our Python code is released as an open source package: <https://github.com/QB3/sparse-ho>. All the experiments are written in Python using Numba (Lam et al., 2015) for the critical parts such as the BCD loop. We compare our gradient computation technique against other competitors (see the competitors section) on the HO problem (Problem (4)).

Solving the inner optimization problem. Note that our proposed method, implicit forward differentiation, has the appealing property that it can be used with any solver. For instance for the Lasso one can combine the proposed algorithm with state of the art solver such as Massias et al. (2018) which would be tedious to combine with iterative differentiation methods. However for the comparison to be fair, for all methods we have used the same vanilla BCD algorithm (recalled in Algorithm 5). We stop the Lasso-type solver when $\frac{f(\beta^{(k+1)}) - f(\beta^{(k)})}{f(0)} < \epsilon^{\text{tol}}$, where f is the cost function of the Lasso or wLasso and ϵ^{tol} a given tolerance. The tolerance is fixed at $\epsilon^{\text{tol}} = 10^{-5}$ for all methods throughout the different benchmarks.

Line search. For each hypergradient-based method, the gradient step is combined with a line-search strategy following the work of Pedregosa (2016)¹.

Initialization. Since the function to optimize \mathcal{L} is not con-

¹see <https://github.com/fabianp/hoag> for details

Table 1. Summary of cost in time and space for each method

Mode	Computed quantity	Space (Lasso)	Time (Lasso)	Space (wLasso)	Time (wLasso)
F. Iterdiff.	\mathcal{J}	$\mathcal{O}(p)$	$\mathcal{O}(2npn_{\text{iter}})$	$\mathcal{O}(p^2)$	$\mathcal{O}(np^2n_{\text{iter}})$
B. Iterdiff.	$\mathcal{J}^\top v$	$\mathcal{O}(2pn_{\text{iter}})$	$\mathcal{O}(nnpn_{\text{iter}} + np^2n_{\text{iter}})$	$\mathcal{O}(p^2n_{\text{iter}})$	$\mathcal{O}(nnpn_{\text{iter}} + np^2n_{\text{iter}})$
Implicit	$\mathcal{J}^\top v$	$\mathcal{O}(p)$	$\mathcal{O}(nnpn_{\text{iter}} + \hat{s}^3)$	$\mathcal{O}(p + \hat{s}^2)$	$\mathcal{O}(nnpn_{\text{iter}} + \hat{s}^3)$
Imp. F. Iterdiff.	\mathcal{J}	$\mathcal{O}(p)$	$\mathcal{O}(nnpn_{\text{iter}} + n\hat{s}n_{\text{iter_jac}})$	$\mathcal{O}(p + \hat{s}^2)$	$\mathcal{O}(nnpn_{\text{iter}} + n\hat{s}^2n_{\text{it_jac}})$

Algorithm 2 IMP. F. ITERDIFF. (proposed)

```

input :  $X \in \mathbb{R}^{n \times p}, y \in \mathbb{R}^n, \lambda \in \mathbb{R}, n_{\text{iter}}, n_{\text{iter\_jac}} \in \mathbb{N}$ 
init   :  $\mathcal{J} = 0$ 
// sequentially compute coef. & Jacobian
if Lasso then
    Get  $\hat{\beta} = \text{Lasso}(X, y, \lambda, n_{\text{iter}})$  and its support  $\hat{S}$ .
     $dr = -X_{:, \hat{S}} \mathcal{J}_{\hat{S}}$  // trick for cheap updates
if wLasso then
    Get  $\hat{\beta} = \text{wLasso}(X, y, \lambda, n_{\text{iter}})$  and its support  $\hat{S}$ .
     $dr = -X_{:, \hat{S}} \mathcal{J}_{\hat{S}, \hat{S}}$ 
for  $k = 0, \dots, n_{\text{iter\_jac}} - 1$  do
    for  $j \in \hat{S}$  do
        if Lasso then
             $\mathcal{J}_{\text{old}} = \mathcal{J}_j$  // trick for cheap update
            // diff. Equation (9) w.r.t.  $\lambda$ 
             $\mathcal{J}_j += \frac{X_{:,j}^\top dr}{\|X_{:,j}\|^2} - \frac{ne^\lambda}{\|X_{:,j}\|^2} \text{sign } \hat{\beta}_j$  //  $\mathcal{O}(n)$ 
             $dr -= X_{:,j} (\mathcal{J}_j - \mathcal{J}_{\text{old}})$  //  $\mathcal{O}(n)$ 
        if wLasso then
             $\mathcal{J}_{\text{old}} = \mathcal{J}_{j, \cdot}$  // trick for cheap update
            // diff. Equation (9) w.r.t.  $\lambda$ 
             $\mathcal{J}_{j, \hat{S}} += \frac{1}{\|X_{:,j}\|^2} X_{:,j}^\top dr$  //  $\mathcal{O}(n \times \hat{s})$ 
             $\mathcal{J}_{j,j} -= \frac{ne^{\lambda_j}}{\|X_{:,j}\|^2} \text{sign } \hat{\beta}_j$  //  $\mathcal{O}(1)$ 
             $dr -= X_{:,j} \otimes (\mathcal{J}_{j, \cdot} - \mathcal{J}_{\text{old}})$  //  $\mathcal{O}(n \times \hat{s})$ 
return  $\hat{\beta}, \mathcal{J}$ 
    
```

vex, initialization plays a crucial role in the final solution as well as the convergence of the algorithm. For instance, initializing $\lambda = \lambda_{\text{init}}$ in a flat zone of $\mathcal{L}(\lambda)$ could lead to slow convergence. In the numerical experiments, the Lasso is initialized with $\lambda_{\text{init}} = \lambda_{\text{max}} - \log(10)$, where λ_{max} is the smallest λ such that 0 is a solution of Problem (2).

Competitors. In this section we compare the empirical performance of implicit forward differentiation algorithm to different competitors. Competitors are divided in two categories. Firstly, the ones relying on hyperparameter gradient:

- **Imp. F. Iterdiff.:** implicit forward differentiation (proposed) described in Algorithm 2.
- **Implicit:** implicit differentiation, which requires solv-

ing a $\hat{s} \times \hat{s}$ linear system as described in Algorithm 1.

- **F. Iterdiff.:** forward differentiation (Deledalle et al., 2014; Franceschi et al., 2017) which jointly computes the regression coefficients $\hat{\beta}$ as well as the Jacobian $\hat{\mathcal{J}}$ as shown in Algorithm 3.

Secondly, the ones not based on hyperparameter gradient:

- **Grid-search:** as recommended by Friedman et al. (2010), we use 100 values on a uniformly-spaced grid from λ_{max} to $\lambda_{\text{max}} - 4 \log(10)$.
- **Random-search:** we sample uniformly at random 100 values taken on the same interval as for the Grid-search $[\lambda_{\text{max}} - 4 \log(10); \lambda_{\text{max}}]$, as suggested by Bergstra et al. (2013).
- **Lattice Hyp.:** lattice hypercube sampling (Bousquet et al., 2017), combines the idea of grid-search and random-search. We used the sampling scheme of Bouhlel et al. (2019) and their code ² to sample the points to evaluate the function on.
- **Bayesian:** sequential model based optimization (SMBO) using a Gaussian process to model the objective function. We used the implementation of Bergstra et al. (2013).³ The constraints space for the hyperparameter search was set in $[\lambda_{\text{max}} - 4 \log(10); \lambda_{\text{max}}]$, and the expected improvement (EI) was used as acquisition function.

The cost and the quantity computed by each algorithm can be found in Table 1. The backward differentiation (Domke, 2012) is not included in the benchmark in Figure 2 since it was several orders of magnitude slower than the other techniques (see Figure 1). This is due to the high cost of the BCD algorithm in backward mode, see Table 1.

4.1. Application to held-out loss

When using the held-out loss, each dataset (X, y) is split in 3 equal parts: the training set $(X^{\text{train}}, y^{\text{train}})$, the validation set $(X^{\text{val}}, y^{\text{val}})$ and the test set $(X^{\text{test}}, y^{\text{test}})$.

²<https://github.com/SMTorg/smt>

³<https://github.com/hyperopt/hyperopt>

(Lasso, held-out criterion). For the Lasso and the held-out loss, the bilevel optimization [Problem \(4\)](#) reads:

$$\begin{aligned} & \arg \min_{\lambda \in \mathbb{R}} \|y^{\text{val}} - X^{\text{val}} \hat{\beta}^{(\lambda)}\|^2 & (16) \\ \text{s.t. } & \hat{\beta}^{(\lambda)} \in \arg \min_{\beta \in \mathbb{R}^p} \frac{1}{2n} \|y^{\text{train}} - X^{\text{train}} \beta\|_2^2 + e^\lambda \|\beta\|_1 . \end{aligned}$$

[Figure 2](#) (top) shows on 3 datasets (see [Appendix D](#) for dataset details) the distance to the ‘‘optimum’’ of $\|y^{\text{val}} - X^{\text{val}} \hat{\beta}^{(\lambda)}\|^2$ as a function of time. Here the goal is to find λ solution of [Problem \(16\)](#). The ‘‘optimum’’ is chosen as the minimum of $\|y^{\text{val}} - X^{\text{val}} \hat{\beta}^{(\lambda)}\|^2$ among all the methods. [Figure 2](#) (bottom) shows the loss $\|y^{\text{test}} - X^{\text{test}} \hat{\beta}^{(\lambda)}\|^2$ on the test set (independent from the training set and the validation set). This illustrates how well the estimator generalizes. Firstly, it can be seen that on all datasets the proposed implicit forward differentiation outperforms forward differentiation which illustrates [Proposition 2](#) and corroborates the cost of each algorithm in [Table 1](#). Secondly, it can be seen that on the *20news* dataset ([Figure 2](#), top) the implicit differentiation ([Algorithm 1](#)) convergence is slower than implicit forward differentiation, forward differentiation, and even slower than the grid-search. In this case, this is due to the very slow convergence of the conjugate gradient algorithm ([Nocedal and Wright, 2006](#)) when solving the ill-conditioned linear system in [Algorithm 1](#).

(MCP, held-out criterion). We also applied our algorithm on an estimator based on a non-convex penalty: the MCP ([Zhang, 2010](#)) with 2 hyperparameters. Since the penalty is non-convex the estimator may not be continuous *w.r.t.* hyperparameters and the theory developed above does not hold. However experimentally implicit forward differentiation outperforms forward differentiation for the HO, see [Appendix C](#) for full details.

4.2. Application to another criterion: SURE

Evaluating models on held-out data makes sense if the design is formed from random samples as it is often considered in supervised learning. However, this assumption does not hold for certain kinds of applications in signal or image processing. For these applications, the held-out loss cannot be used as the criterion for optimizing the hyperparameters of a given model. In this case, one may use a proxy of the prediction risk, like the Stein Unbiased Risk Estimation (SURE, [Stein \(1981\)](#)). The SURE is an unbiased estimator of the prediction risk under weak differentiable conditions. The drawback of this criterion is that it requires the knowledge of the variance of the noise. The SURE is defined as follows: $\text{SURE}(\lambda) = \|y - X \hat{\beta}^{(\lambda)}\|^2 - n\sigma^2 + 2\sigma^2 \text{dof}(\hat{\beta}^{(\lambda)})$, where the degrees of freedom ([dof Efron 1986](#)) is defined as $\text{dof}(\hat{\beta}^{(\lambda)}) = \sum_{i=1}^n \text{cov}(y_i, (X \hat{\beta}^{(\lambda)})_i) / \sigma^2$. The dof can be seen a measure of the complexity of the model, for instance for the Lasso $\text{dof}(\hat{\beta}^{(\lambda)}) = \hat{s}$, see [Zou et al. \(2007\)](#).

The SURE can thus be seen as a criterion trading data-fidelity against model complexity. However, the dof is not differentiable (not even continuous in the Lasso case), yet it is possible to construct a weakly differentiable approximation of it based on Finite Differences Monte-Carlo (see [Deledalle et al. 2014](#) for full details), with $\epsilon > 0$ and $\delta \sim \mathcal{N}(0, \text{Id}_n)$:

$$\text{dof}_{\text{FDMC}}(y, \lambda, \delta, \epsilon) = \frac{1}{\epsilon} \langle X \hat{\beta}^{(\lambda)}(y + \epsilon \delta) - X \hat{\beta}^{(\lambda)}(y), \delta \rangle .$$

We use this smooth approximation in the bi-level optimization problem to find the best hyperparameter. The bi-level optimization problem then reads:

$$\begin{aligned} & \arg \min_{\lambda \in \mathbb{R}} \|y - X \hat{\beta}^{(\lambda)}\|^2 + 2\sigma^2 \text{dof}_{\text{FDMC}}(y, \lambda, \delta, \epsilon) & (17) \\ \text{s.t. } & \hat{\beta}^{(\lambda)}(y) \in \arg \min_{\beta \in \mathbb{R}^p} \frac{1}{2n} \|y - X \beta\|_2^2 + e^\lambda \|\beta\|_1 \\ & \hat{\beta}^{(\lambda)}(y + \epsilon \delta) \in \arg \min_{\beta \in \mathbb{R}^p} \frac{1}{2n} \|y + \epsilon \delta - X \beta\|_2^2 + e^\lambda \|\beta\|_1 \end{aligned}$$

Note that solving this problem requires the computation of two (instead of one for the held-out loss) Jacobians *w.r.t.* λ of the solution $\hat{\beta}^{(\lambda)}$ at the points y and $y + \epsilon \delta$.

(Lasso, SURE criterion). To investigate the estimation performance of the implicit forward differentiation in comparison to the competitors described above, we used as metric the (normalized) Mean Squared Error (MSE) defined as $\text{MSE} \triangleq \|\hat{\beta} - \beta^*\|^2 / \|\beta^*\|^2$. The entries of the design matrix $X \in \mathbb{R}^{n \times p}$ are i.i.d. random Gaussian variables $\mathcal{N}(0, 1)$. The number of rows is fixed to $n = 100$. Then, we generated β^* with 5 non-zero coefficients equals to 1. The vector y was computed by adding to $X \beta^*$ additive Gaussian noise controlled by the Signal-to-Noise Ratio: $\text{SNR} \triangleq \|X \beta^*\| / \|y - X \beta^*\|$ (here $\text{SNR} = 3$). Following [Deledalle et al. \(2014\)](#), we set $\epsilon = 2\sigma/n^{0.3}$. We varied the number of features p between 200 and 10,000 on a linear grid of size 10. For a fixed number of features, we performed 50 repetitions and each point of the curves represents the mean of these repetitions. Comparing efficiency in time between methods is difficult since they are not directly comparable. Indeed, grid-search and random-search discretize the HO space whereas others methods work in the continuous space which is already an advantage. However, to be able to compare the hypergradient methods and possibly compare them to the others, we computed the total amount of time for a method to return its optimal value of λ . In order to have a *fair* comparison, we compared 50 evaluations of the line-search for each hypergradient methods, 50 evaluations of the Bayesian methods and finally 50 evaluations on fixed or random grid. We are aware that the cost of each of these evaluations is not the same but it allows to see that our method stays competitive in time with optimizing one parameter. Moreover we will also see that our method scales better with a large number of hyperparameters to optimize.

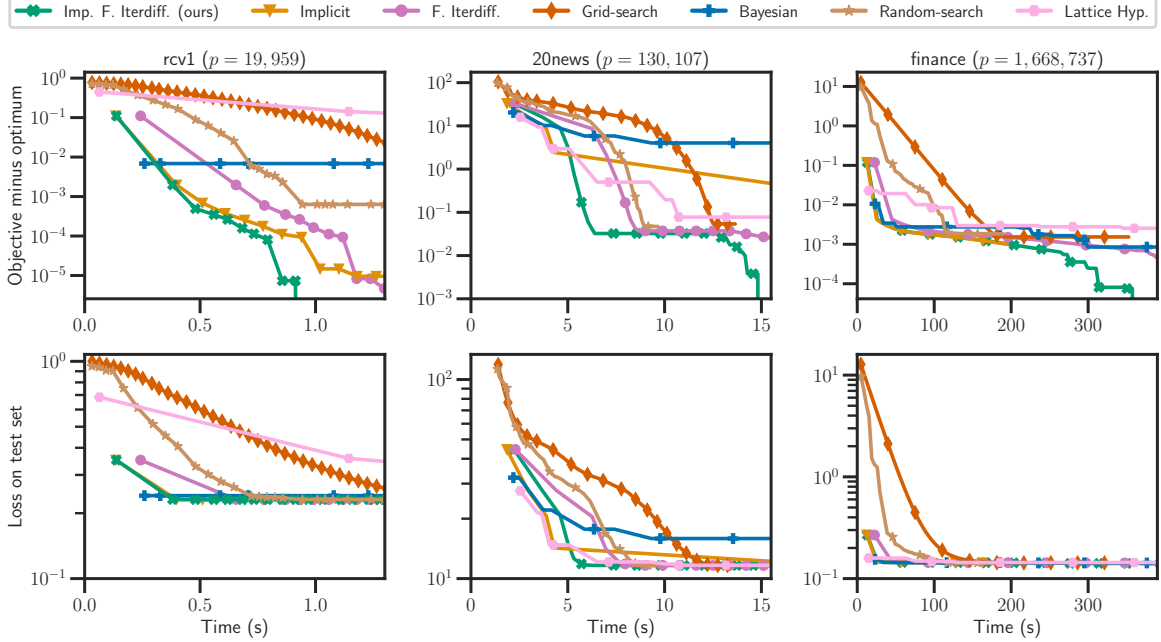


Figure 2. **Computation time for the HO of the Lasso on real data.** Distance to “optimum” (top) and performance (bottom) on the test set for the Lasso for 3 different datasets: *rcv1*, *20news* and *finance*.

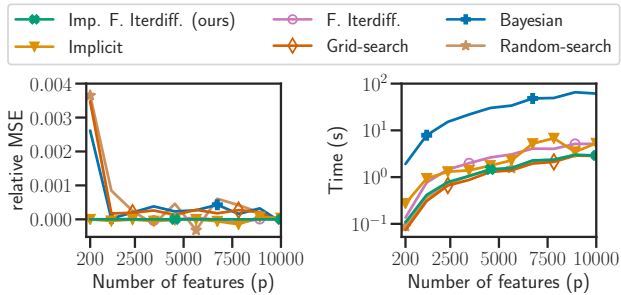


Figure 3. **Lasso: estimation performance.** Estimation relative Mean Squared Error (left) and running time (right) as a function of the number of features for the Lasso model.

Figure 3 shows the influence of the number of features on the relative MSE (ie. MSE of a method minus the MSE of our implicit forward method) and the computation time. First, MSE of all gradient based methods is lower than the other methods which means that $\hat{\beta}^{(\lambda)}$ leads to a better estimation when λ is chosen via the gradient based methods. This illustrates that continuous optimization for hyperparameter selection leads to better estimation performance than discrete or Bayesian optimization. Yet, the running time of our proposed method is the lowest of all hypergradient-based strategies and competes with the grid-search and the random-search.

(*Weighted Lasso vs Lasso, SURE criterion*). As our method

leverages the sparsity of the solution, it can be used for HO with a large number of hyperparameters, contrary to classical forward differentiation. The weighted Lasso (wLasso, Zou 2006) has p hyperparameters and was introduced to reduce the bias of the Lasso. However setting the p hyperparameters is impossible with grid-search.

Figure 4 shows the estimation MSE and the running time of the different methods to obtain the hyperparameter values as a function of the number of features used to simulate the data. The simulation setting is here the same as for the Lasso problems investigated in Figure 3 ($n = 100$, SNR = 3). We compared the classical Lasso estimator and the weighted Lasso estimator where the regularization hyperparameter was chosen using implicit forward differentiation and the forward iterative differentiation as described in Algorithm 3. Problem (4) is not convex for the weighted Lasso and a descent algorithm like ours can be trapped in local minima, crucially depending on the starting point λ_{init} . To alleviate this problem, we introduced a regularized version of Problem (4):

$$\begin{aligned} \arg \min_{\lambda \in \mathbb{R}} \mathcal{C}(\hat{\beta}^{(\lambda)}) + \gamma \sum_j^p \lambda_j^2 \\ \text{s.t. } \hat{\beta}^{(\lambda)} \in \arg \min_{\beta \in \mathbb{R}^p} \triangleq \psi(\beta, \lambda) . \end{aligned} \quad (18)$$

The solution obtained by solving Equation (18) is then used as the initialization $\lambda^{(0)}$ for our algorithm. In this experiment the regularization term is constant $\gamma =$




Figure 4. **Lasso vs wLasso.** Estimation Mean Squared Error (left) and running (right) of competitors as a function of the number of features for the weighted Lasso and Lasso models.

$C(\beta^{(\lambda_{\max})})/10$. We see in Figure 4 that the weighted Lasso gives a lower MSE than the Lasso and allows for a better recovery of β^* . This experiment shows that the amount of time needed to obtain the vector of hyperparameters of the weighted Lasso via our algorithm is in the same range as for obtaining the unique hyperparameter of the Lasso problem. It also shows that our proposed method is much faster than the *naive* way of computing the Jacobian using forward or backward iterative differentiation. The implicit differentiation method stays competitive for the wLasso due to the small support of the solution and hence a small matrix to inverse. A maximum running time threshold was used for this experiment checking the running time at each line-search iteration, explaining why the forward differentiation and backward differentiation of the wLasso does not explode in time on Figure 4.

Conclusion

In this work we studied the performance of several methods to select hyperparameters of Lasso-type estimators showing results for the Lasso and the weighted Lasso, which have respectively one or p hyperparameters. We exploited the sparsity of the solutions and the specific structure of the iterates of forward differentiation, leading to our implicit forward differentiation algorithm that computes efficiently the full Jacobian of these estimators *w.r.t.* the hyperparameters. This allowed us to select them through a standard gradient descent and have an approach that scales to a high number of hyperparameters. Importantly, contrary to a classical implicit differentiation approach, the proposed algorithm does not require solving a linear system. Finally, thanks to its two steps nature, it is possible to leverage in the first step the availability of state-of-the-art Lasso solvers that make use of techniques such as active sets or screening rules. Such algorithms, that involve calls to inner solvers run on subsets of features, are discontinuous

w.r.t. hyperparameters which would significantly challenge a single step approach based on automatic differentiation.

Acknowledgments This work was funded by ERC Starting Grant SLAB ERC-StG-676943 and ANR GraVa ANR-18-CE40-0005.

References

- A. Agrawal, B. Amos, S. Barratt, S. Boyd, S. Diamond, and J. Z. Kolter. Differentiable convex optimization layers. In *Advances in neural information processing systems*, pages 9558–9570, 2019.
- B. Amos and J. Z. Kolter. Optnet: Differentiable optimization as a layer in neural networks. In *ICML*, volume 70, pages 136–145, 2017.
- A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind. Automatic differentiation in machine learning: a survey. *J. Mach. Learn. Res.*, 18(153):1–43, 2018.
- Y. Bengio. Gradient-based optimization of hyperparameters. *Neural computation*, 12(8):1889–1900, 2000.
- J. Bergstra and Y. Bengio. Random search for hyperparameter optimization. *J. Mach. Learn. Res.*, 13:281–305, 2012.
- J. Bergstra, D. Yamins, and D. D. Cox. Hyperopt: A python library for optimizing the hyperparameters of machine learning algorithms. In *Proceedings of the 12th Python in science conference*, pages 13–20, 2013.
- M. Borgerding, P. Schniter, and S. Rangan. Amp-inspired deep networks for sparse linear inverse problems. *IEEE Transactions on Signal Processing*, 65(16):4293–4308, 2017.
- M. A. Bouhlef, J. T. Hwang, N. Bartoli, R. Lafage, J. Morlier, and J. R. R. A. Martins. A python surrogate modeling framework with derivatives. *Advances in Engineering Software*, page 102662, 2019. ISSN 0965-9978. doi: <https://doi.org/10.1016/j.advengsoft.2019.03.005>.
- O. Bousquet, S. Gelly, K. Kurach, O. Teytaud, and D. Vincent. Critical hyper-parameters: No random, no cry. *arXiv preprint arXiv:1706.03200*, 2017.
- P. Breheny and J. Huang. Coordinate descent algorithms for nonconvex penalized regression, with applications to biological feature selection. *Ann. Appl. Stat.*, 5(1):232, 2011.
- E. Brochu, V. M. Cora, and Nando De Freitas. A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*, 2010.
- O. Chapelle, V. Vapnik, O. Bousquet, and S. Mukherjee. Choosing multiple parameters for support vector machines. *Machine learning*, 46(1-3):131–159, 2002.
- P. L. Combettes and V. R. Wajs. Signal recovery by proximal forward-backward splitting. *Multiscale Modeling & Simulation*, 4(4):1168–1200, 2005.
- I. Daubechies, M. Defrise, and C. De Mol. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Comm. Pure Appl. Math.*, 57(11):1413–1457, 2004.
- C.-A. Deledalle, S. Vaiteer, J. Fadili, and G. Peyré. Stein Unbiased GrAdient estimator of the Risk (SUGAR) for multiple parameter selection. *SIAM J. Imaging Sci.*, 7(4):2448–2487, 2014.
- J. Domke. Generic methods for optimization-based modeling. In *AISTATS*, volume 22, pages 318–326, 2012.
- C. Dossal, M. Kachour, M.J. Fadili, G. Peyré, and C. Chesneau. The degrees of freedom of the lasso for general design matrix. *Statistica Sinica*, 23(2):809–828, 2013.
- B. Efron. How biased is the apparent error rate of a prediction rule? *J. Amer. Statist. Assoc.*, 81(394):461–470, 1986.
- L. C. Evans and R. F. Gariepy. *Measure theory and fine properties of functions*. CRC Press, 1992.
- C. S. Foo, C. B. Do, and A. Y. Ng. Efficient multiple hyperparameter learning for log-linear models. In *Advances in neural information processing systems*, pages 377–384, 2008.
- L. Franceschi, M. Donini, P. Frasconi, and M. Pontil. Forward and reverse gradient-based hyperparameter optimization. In *ICML*, pages 1165–1173, 2017.
- J. Frecon, S. Salzo, and M. Pontil. Bilevel learning of the group lasso structure. In *Advances in Neural Information Processing Systems*, pages 8301–8311, 2018.
- J. Friedman, T. J. Hastie, and R. Tibshirani. Regularization paths for generalized linear models via coordinate descent. *J. Stat. Softw.*, 33(1):1–22, 2010.
- K. Gregor and Y. LeCun. Learning fast approximations of sparse coding. In *ICML*, pages 399–406, 2010.
- E. Hale, W. Yin, and Y. Zhang. Fixed-point continuation for ℓ_1 -minimization: Methodology and convergence. *SIAM J. Optim.*, 19(3):1107–1130, 2008.
- K. Kunisch and T. Pock. A bilevel optimization approach for parameter learning in variational models. *SIAM J. Imaging Sci.*, 6(2):938–983, 2013.
- S. K. Lam, A. Pitrou, and S. Seibert. Numba: A LLVM-based Python JIT Compiler. In *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*, pages 1–6. ACM, 2015.

- J. Larsen, L. K. Hansen, C. Svarer, and M. Ohlsson. Design and regularization of neural networks: the optimal use of a validation set. In *Neural Networks for Signal Processing VI. Proceedings of the 1996 IEEE Signal Processing Society Workshop*, pages 62–71, 1996.
- J. Larsen, C. Svarer, L. N. Andersen, and L. K. Hansen. Adaptive regularization in neural network modeling. In *Neural Networks: Tricks of the Trade - Second Edition*, pages 111–130. Springer, 2012.
- J. Liu, X. Chen, Z. Wang, and W. Yin. Alista: Analytic weights are as good as learned weights in lista. In *International Conference on Learning Representations*, 2018.
- W. Liu, Y. Yang, et al. Parametric or nonparametric? a parametricness index for model selection. *Ann. Statist.*, 39(4):2074–2102, 2011.
- J. Lorraine, P. Vicol, and D. Duvenaud. Optimizing millions of hyperparameters by implicit differentiation. *arXiv preprint arXiv:1911.02590*, 2019.
- D. Maclaurin, D. Duvenaud, and Ryan Adams. Gradient-based hyperparameter optimization through reversible learning. In *ICML*, volume 37, pages 2113–2122, 2015.
- J. Mairal, F. Bach, and J. Ponce. Task-driven dictionary learning. *IEEE Trans. Pattern Anal. Mach. Intell.*, 34(4): 791–804, 2012.
- M. Massias, A. Gramfort, and J. Salmon. Celer: a Fast Solver for the Lasso with Dual Extrapolation. In *ICML*, volume 80, pages 3315–3324, 2018.
- M. Massias, S. Vaiter, A. Gramfort, and J. Salmon. Dual extrapolation for sparse generalized linear models. *arXiv preprint arXiv:1907.05830*, 2019.
- V. Niculae and M. Blondel. A regularized framework for sparse and structured neural attention. In *Advances in neural information processing systems*, pages 3338–3348, 2017.
- J. Nocedal and S. J. Wright. *Numerical optimization*. Springer Series in Operations Research and Financial Engineering. Springer, New York, second edition, 2006.
- F. Pedregosa. Hyperparameter optimization with approximate gradient. In *ICML*, volume 48, pages 737–746, 2016.
- S. Ramani, T. Blu, and M. Unser. Monte-Carlo SURE: a black-box optimization of regularization parameters for general denoising algorithms. *IEEE Trans. Image Process.*, 17(9):1540–1554, 2008.
- M. W. Seeger. Cross-validation optimization for large scale structured classification kernel methods. *J. Mach. Learn. Res.*, 9:1147–1178, 2008.
- J. Snoek, H. Larochelle, and R. P. Adams. Practical bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems*, pages 2951–2959, 2012.
- E. Soubies, L. Blanc-Féraud, and G. Aubert. A unified view of exact continuous penalties for ℓ_2 - ℓ_0 minimization. *SIAM J. Optim.*, 27(3):2034–2060, 2017.
- C. M. Stein. Estimation of the mean of a multivariate normal distribution. *Ann. Statist.*, 9(6):1135–1151, 1981.
- L. R. A. Stone and J.C. Ramer. Estimating WAIS IQ from Shipley Scale scores: Another cross-validation. *Journal of clinical psychology*, 21(3):297–297, 1965.
- Y. Sun, H. Jeong, J. Nutini, and M. Schmidt. Are we there yet? manifold identification of gradient-related proximal methods. In *AISTATS*, volume 89, pages 1110–1119, 2019.
- R. Tibshirani. Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 58(1):267–288, 1996.
- R. J. Tibshirani. The lasso problem and uniqueness. *Electron. J. Stat.*, 7:1456–1490, 2013.
- R. J. Tibshirani and J. Taylor. The solution path of the generalized lasso. *Ann. Statist.*, 39(3):1335–1371, 2011.
- P. Tseng and S. Yun. Block-coordinate gradient descent method for linearly constrained nonsmooth separable optimization. *J. Optim. Theory Appl.*, 140(3):513, 2009.
- S. Vaiter, C.-A. Deledalle, G. Peyré, C. Dossal, and J. Fadili. Local behavior of sparse analysis regularization: Applications to risk estimation. *Appl. Comput. Harmon. Anal.*, 35(3):433–451, 2013.
- S. Vaiter, C.-A. Deledalle, G. Peyré, J. M. Fadili, and C. Dossal. The degrees of freedom of partly smooth regularizers. *Ann. Inst. Stat. Math.*, 69(4):791–832, 2017.
- K. Wu, Y. Guo, Z. Li, and C. Zhang. Sparse coding with gated learned ista. In *International Conference on Learning Representations*, 2019.
- B. Xin, Y. Wang, W. Gao, D. Wipf, and B. Wang. Maximal sparsity with deep networks? In *Advances in Neural Information Processing Systems*, pages 4340–4348, 2016.
- M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 68(1):49–67, 2006.

- C.-H. Zhang. Nearly unbiased variable selection under minimax concave penalty. *Ann. Statist.*, 38(2):894–942, 2010.
- H. Zou. The adaptive lasso and its oracle properties. *J. Amer. Statist. Assoc.*, 101(476):1418–1429, 2006.
- H. Zou and T. J. Hastie. Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 67(2):301–320, 2005.
- H. Zou, T. J. Hastie, and R. Tibshirani. On the “degrees of freedom” of the lasso. *Ann. Statist.*, 35(5):2173–2192, 2007.