
000 BENCHMARKING STOCHASTIC APPROXIMATION ALGO- 001 RITHMS FOR FAIRNESS-CONSTRAINED TRAINING OF 002 DEEP NEURAL NETWORKS 003

004 **Anonymous authors**
005
006

007 Paper under double-blind review
008
009

010 ABSTRACT 011

012 The ability to train Deep Neural Networks (DNNs) with constraints is instrumental
013 in improving the fairness of modern machine-learning models. Many algorithms
014 have been analysed in recent years, and yet there is no standard, widely accepted
015 method for the constrained training of DNNs. In this paper, we provide a chal-
016 lenging benchmark of real-world large-scale fairness-constrained learning tasks,
017 built on top of the US Census (Folktbles, Ding et al. (2021)). We point out the
018 theoretical challenges of such tasks and review the main approaches in stochastic
019 approximation algorithms. Finally, we demonstrate the use of the benchmark by im-
020 plementing and comparing three recently proposed, but as-of-yet unimplemented,
021 algorithms both in terms of optimization performance, and fairness improvement.
022 We will release the code of the benchmark as a Python package after peer-review.
023
024

025 1 INTRODUCTION 026

027 There has been a considerable interest in detecting and mitigating bias in artificial intelligence (AI)
028 systems, recently. Multiple legislative frameworks, including the AI Act in the European Union,
029 require the bias to be removed, but there is no agreement on what the correct definition of bias is or
030 how to remove it. A natural translation of the requirement of removing bias into the formulation
031 of training of deep neural network (DNN) utilizes constraints bounding the difference in empirical
032 risk across multiple subgroups (Chen et al., 2018; Nandwani et al., 2019; Ravi et al., 2019). Over
033 the past five years, there have been numerous algorithms proposed to solve convex and non-convex
034 empirical-risk minimization (ERM) problems subject to constraints bounding the absolute value of
035 empirical risk (Fang et al., 2024; Berahas et al., 2021; Curtis et al., 2024a; Oztoprak et al., 2023;
036 Berahas et al., 2023; Na et al., 2023a;b; Bollapragada et al., 2023; Curtis et al., 2024b; Shi et al.,
037 2022; Facchinei & Kungurtsev, 2023; Huang et al., 2025; Huang & Lin, 2023). Numerous other
algorithms of this kind could be construed, based on a number of design choices, including:

- 038 • sampling techniques for the ERM objective and the constraints, either the same or different;
- 039 • use of first-order or higher-order derivatives, possibly in quasi-Newton methods;
- 040 • use of globalization strategies such as filters or line search;
- 041 • use of “true” globalization strategies including random initial points and random restarts in order
042 to reach global minimizers.

043 Nevertheless, there is no single toolkit implementing the algorithms, which would allow for their
044 easy comparison, and there is no benchmark to test the combinations of design choices on.
045

046 In this paper, we consider the constrained ERM problem:

$$047 \min_{x \in \mathbb{R}^n} \mathbb{E}[f(x, \xi)] \quad \text{s.t.} \quad \mathbb{E}[c(x, \zeta)] \leq 0, \quad (1)$$

048

049 where ξ and ζ are random variables. Further, we provide an automated way of constructing the ERM
050 formulations out of a computation graph of a neural network defined by PyTorch or TensorFlow, the
051 choice of the constraints (see Table 1), and a definition of the protected subgroups to apply constraints
052 to. Specifically, we provide means of utilizing the US Census data via the Python package Folktbles,
053 together with definitions of up to 5.7 billion protected subgroups. This presents a challenging
benchmark in stochastic approximation for the constrained training of deep neural networks.

Table 1: Particular formulations of the constraint function c to enforce fairness.

Model	Our formulation
Accuracy equality	$ \mathbb{E}_{\mathcal{D}[\text{group } A]}[\ell(f_\theta(X), Y)] - \mathbb{E}_{\mathcal{D}[\text{group } B]}[\ell(f_\theta(X), Y)] \leq \delta$
Equal opportunity Hardt et al. (2016)	$ \mathbb{E}_{\mathcal{D}[\text{group } A, Y=+]}[\ell(f_\theta(X), Y)] - \mathbb{E}_{\mathcal{D}[\text{group } B, Y=+]}[\ell(f_\theta(X), Y)] \leq \delta$
Equalized odds Hardt et al. (2016)	$\sum_{v \in \{+, -\}} \mathbb{E}_{\mathcal{D}[\text{group } A, Y=v]}[\ell(f_\theta(X), Y)] - \mathbb{E}_{\mathcal{D}[\text{group } B, Y=v]}[\ell(f_\theta(X), Y)] \leq \delta$

Our contributions. The contributions of this paper are:

- a literature review of algorithms subject to handling (1);
- a toolbox that (i) implements four algorithms applicable in real-world situations, and (ii) provides an easy-to-use benchmark on real-world fairness problems;
- numerical experiments that compare these algorithms on a real-world dataset, and a comparison with alternative approaches to fairness.

Paper structure. The rest of the paper is organized as follows. Section 2 reviews related works and presents the relevant notions of fairness. Section 3 introduces the algorithms. Section 4 reports on our experiments. Section 5 concludes.

2 RELATED WORK, AND BACKGROUND IN FAIRNESS

In the literature on fairness, one distinguishes among pre-processing, in-processing, and post-processing. Pre-processing methods focus on modifying the training data to mitigate biases (Tawakuli & Engel, 2024; Du et al., 2021). In-processing methods enforce fairness during the training process by modifying the learning algorithm itself (Wan et al., 2023). Post-processing methods adjust the model's predictions after training (Kim et al., 2019). The constrained ERM approach (1) belongs to the class of in-processing methods.

In-processing methods include several approaches. One trend consists in jointly learning a predictor function and an adversarial agent that aims to reconstitute the subgroups from the predictor (Adel et al., 2019; Louppe et al., 2017; Madras et al., 2018; Edwards & Storkey, 2016). Another approach consists in adding “penalization” terms to the empirical risk term. These additional penalization terms, commonly referred to as regularizers, promote models that are a compromise between fitting the training data, and optimizing a fairness metric. Differentiable regularizers include, among others, HSIC (Li et al., 2022), Fairret (Buyl et al., 2024), or Prejudice Remover (Kamishima et al., 2012).

Closer to our setting, Cotter et al. (2019) consider minimizing the empirical risk subject to the so-called rate constraints based on the model’s prediction rates on different datasets. These rates, derived from a dataset, give rise to non-convex, non-smooth, and large-scale inequality constraints akin to (1). Cotter et al. (2019) argue that hard constraints, although leading to a more difficult optimization problem, offer advantages over using a weighted sum of multiple penalization terms. Indeed, while the choice of weights for the penalization terms may depend on the dataset, specifying one constraint for each goal is easier for practitioners. In addition, a penalization-based model provides a predictor that balances minimizing the data-fit term and penalties in an opaque way, whereas a constraint-based model allows for a clearer understanding of the model design: minimizing the data-fit term subject to “hard” fairness constraints. Rate constraints differ from those in (1) in that they are piecewise-constant, rendering first-order methods unsuitable for solving them. We refer to the recent work of Ramirez et al. (2025) for a detailed argument on why constraining ERM problems is preferable to penalizing the ERM with multiple terms.

Major toolboxes for evaluating the fairness of models or for training models with fairness guarantees include AIF360 (Bellamy et al., 2018) and FairLearn (Bird et al., 2020). Delaney et al. (2024) compute the Pareto front of accuracy and fairness metrics for high-capacity models, and Buyl et al. (2024) provides differentiable fairness-inducing penalization terms. We also note the recent Cooper toolbox, closest to our setting, with Lagrangian-based methods focus (Gallego-Posada et al., 2025).

Le Quy et al. (2022) provides a detailed survey of fairness-oriented datasets, and Ding et al. (2021) derives new datasets. The benchmark of Han et al. (2023) reviews the existence of biases in prominent datasets, finding that “not all widely used fairness datasets stably exhibit fairness issues”, and assesses

108 the performance of a range of in-processing methods in addressing biases, focusing on differentiable
 109 minimization only. Other benchmarks of fairness methods include Defrance et al. (2024); Fabris et al.
 110 (2022); Pessach & Shmueli (2022); Chen et al. (2024). Statistical aspects of the fairness-constrained
 111 Empirical Risk Minimization have only been considered recently; see e.g. Chamon et al. (2022).

112 The template problem (1) encompasses fairness-enforcing approaches that find applications in high-
 113 risk domains, such as credit scoring, hiring processes, medicine and healthcare (Chen et al., 2023),
 114 ranking and recommendation (Pitoura et al., 2022), but also in forecasting the observations of linear
 115 dynamical systems (Zhou et al., 2023b), or in two-sided economic markets (Zhou et al., 2023a). In
 116 addition, solving (1) is of interest in other fields, such as compression of neural networks (Chen
 117 et al., 2018), improving statistical performance of neural networks (Nandwani et al., 2019; Ravi et al.,
 118 2019), or the training of neural networks with constraints on the Lipschitz bound (Pauli et al., 2021).
 119 We note that all the aforementioned methodologies feature large-scale constraints.

120
 121 **Deep neural networks (DNNs).** Consider a dataset of N observations $\mathcal{D} = \{(X_i, Y_i), i = 1, \dots, N\}$. We seek some function f_θ such that $f_\theta(X_i) \approx Y_i$. A typical formulation of this task is the
 122 following regression problem:

$$124 \quad \min_{\theta \in \mathbb{R}^n} \frac{1}{N} \sum_{i=1}^N \ell(f_\theta(X_i), Y_i) + \mathcal{R}(\theta). \quad (2)$$

125 Here, $\ell : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ is a loss function, such as the logistic loss $\ell(y; z) = \log(1 + e^{-yz})$, the hinge
 126 loss $\ell(y; z) = \max\{0, 1 - yz\}$, the absolute deviation loss $\ell(y; z) = |y - z|$, or the square loss
 127 $\ell(y; z) = \frac{1}{2}(y - z)^2$. The term \mathcal{R} is a regularizer, and f_θ is a deep neural network (DNN) of depth L
 128 with parameters θ . The DNN f_θ is defined recursively, for some input X , as
 129

$$132 \quad a_0 = X, \quad a_i = \rho_i(V_i(\theta)a_{i-1}), \quad \text{for every } i = 1, \dots, L, \quad f_\theta(X) = a_L, \quad (3)$$

133 where $V_i(\cdot)$ are linear maps into the space of matrices, and ρ_i are activation functions applied
 134 coordinate-wise, such as ReLU $\max(0, t)$, quadratics t^2 , hinge losses $\max\{0, t\}$, and SoftPlus
 135 $\log(1 + e^t)$. A dataset \mathcal{D} is described by attributes (or features), such as age, income, gender, etc.
 136 The attribute which the DNN is trained to predict is called the class attribute. We denote the class
 137 attribute by Y , whereas the predicted value given by the DNN is denoted by \hat{Y} . Both Y and \hat{Y} are
 138 binary and take values in $\{+, -\}$.
 139

140 **Fairness-aware learning applied to DNNs.** The goal of this approach is to reduce discriminatory
 141 behavior in the predictions of a DNN across different demographic groups (e.g., male vs. female).
 142 The demographic groups are also referred to as subgroups. The attributes such as race or gender
 143 which must be handled cautiously are called protected. We denote by S the protected attribute or more
 144 generally the set of groups defined by multiple protected attributes, s_1, \dots, s_m its possible values or
 145 the indicator of membership in the groups, and $\mathcal{D}[s_i]$ the observations in \mathcal{D} such that $S = s_i$. A way
 146 to impose fairness on the learned predictor is to equip (2) with suitable constraints. Some possible
 147 constraint choices are shown in Table 1. Choosing loss difference bound as the constraint, denoting
 148 $\ell^{s_i}(\theta) = \frac{1}{|\mathcal{D}[s_i]|} \sum_{X, Y \in \mathcal{D}[s_i]} \ell(f_\theta(X), Y)$ for $i = 1, \dots, m$, and setting $\delta > 0$ yields formulation:
 149

$$150 \quad \min_{\theta \in \mathbb{R}^n} \frac{1}{N} \sum_{i=1}^N \ell(f_\theta(X_i), Y_i) + \mathcal{R}(\theta) \\ 151 \quad \text{s.t.} \quad -\delta \leq \ell^{s_i}(\theta) - \frac{1}{m} \sum_{j=1}^m \ell^{s_j}(\theta) \leq \delta, \quad i = 1, \dots, m. \quad (4)$$

152 Bounding the distance between subgroup losses yields m constraints. Other formulations are
 153 possible, such as bounding the distance between every pair of subgroup, providing simpler individual
 154 constraints, but in greater number ($m(m + 1)/2$). Formulation (4) extends to several protected
 155 attributes by adding the corresponding set of equations; we omit this direct generalisation for clarity.
 156 Note that we employ constraints based on the loss function, as it is a continuous function, amenable
 157 to nonsmooth nonconvex optimization. For constraints involving directly discontinuous quantities
 158 such as accuracies and rates, see Cotter et al. (2019).
 159

Table 2: Three elementary notions of fairness

164	Independence		Separation	Sufficiency
165	$\frac{1}{m} \sum_{i=1}^m P_i^{\text{ind}} - \frac{1}{m} \sum_j P_j^{\text{ind}} $	$\frac{1}{2} \sum_{v \in \{+,-\}} \frac{1}{m} \sum_{i=1}^m P_{i,v}^{\text{Sp}} - \frac{1}{m} \sum_j P_{j,v}^{\text{Sp}} $	$\frac{1}{2} \sum_{v \in \{+,-\}} \frac{1}{m} \sum_{i=1}^m P_{i,v}^{\text{Sf}} - \frac{1}{m} \sum_j P_{j,v}^{\text{Sf}} $	
166				

Fairness metrics. There exist tens of fairness metrics (Verma & Rubin, 2018). However, Barocas et al. (2023, Ch. 3) pointed out that most fairness metrics are combinations of three elementary fairness criteria: independence, separation, and sufficiency. These criteria cannot be minimized simultaneously, and there is a trade-off between attaining the elementary fairness metrics and the prediction accuracy, i.e., the probability that the predicted value is equal to the actual value. Thus, we seek an optimal trade-off between attaining the fairness metrics and minimizing the prediction inaccuracy. We next recall the definitions of these three relevant metrics, following Barocas et al. (2023), and provide the formula for computing them in Table 2.

Independence (Ind) This fairness criterion requires the prediction \hat{Y} to be statistically independent of the protected attribute S . Equivalent definitions of independence for a binary classifier \hat{Y} are referred to as statistical parity (SP), demographic parity, and group fairness. Independence is the simplest criterion to work with, both mathematically and algorithmically. In a binary classification task, independence implies the equality of $P_i^{\text{ind}} = P(\hat{Y} = + | S = s_i)$ for all $i = 1, \dots, m$.

Separation (Sp) Unlike independence, the separation criterion requires the prediction \hat{Y} to be statistically independent of the protected attribute S , given the true label Y . The separation criterion also appears under the name Equalized odds (EO). In a binary classification task, the separation criterion requires that all groups experience the same true negative rate and the same true positive rate. Formally, we require the equality of $P_{i,v}^{\text{Sp}} = P(\hat{Y} = + \mid S = s_i, Y = v)$ for every $i = 1, \dots, m$, and $v \in \{+, -\}$.

Sufficiency (Sf) The sufficiency criterion is satisfied if the true label Y is statistically independent of the protected attribute S , given the prediction \hat{Y} . In a binary classification task, the sufficiency criterion requires a parity of positive and negative predictive values across the groups. Formally, we require the equality of $P_{v,s}^{\text{Sf}} = P(Y = + \mid \hat{Y} = v, S = s)$, for every $i = 1, \dots, m$, and $v \in \{+, -\}$.

3 ALGORITHMS

We recall that we consider the optimization problem

$$\min_{x \in \mathbb{P}^n} F(x) \quad \text{s.t.} \quad C(x) \leq 0, \quad (5)$$

where the functions $F : \mathbb{R}^n \rightarrow \mathbb{R}$ and $C : \mathbb{R}^n \rightarrow \mathbb{R}^m$ are defined as expectations of functions f and c , which depend on random variables ξ and ζ , respectively. Solving (5) has the following challenges:

- large-scale objective and constraint functions, which require sampling schemes,
- the necessity of incorporating inequality constraints, not merely equality constraints (see fairness formulations in Table 1),
- the necessity to cope with the nonconvexity and nonsmoothness of F and C , due to the presence of neural networks.

In this section, we identify the algorithms that address these challenges most precisely. However, we note that there exists currently no algorithm with guarantees for such a general setting.

Recalls and notation. We denote the projection of a point x onto a set \mathcal{X} by $\text{proj}_{\mathcal{X}}(x) = \arg \min_{v \in \mathcal{X}} \|x - v\|^2$. We denote by $N \sim \mathcal{G}(p_0)$ sampling a random variable from the geometric distribution with a parameter p_0 , i.e., the probability that $N = n$ equals $(1 - p_0)^n p_0$ for $n \geq 0$. We distinguish between the random variable ξ associated with the objective function and the random variable ζ associated with the constraint function. Their probability distributions are denoted by \mathcal{P}_ξ and \mathcal{P}_ζ . For an integer $J \in \mathbb{N}$, a set $\{\xi_j\}_{j=1}^J$ of independent and identically distributed random

216 Table 3: Assumptions on objective and constraint functions, F and C , which allow for theoretical
217 convergence proofs.

Algorithm	Objective function F				Constraint function C			
	stochastic	weakly convex	C^1 with Lipschitz ∇F	tame loc. Lipschitz	stochastic	$C(x) = 0$	$C(x) = 0$ and $C(x) \leq 0$	tame loc. Lipschitz
SGD	✓	(✓)	(✓)	✓				
Berahas et al. (2023)	✓	—	✓	—	—	✓	—	✓
Fang et al. (2024)	✓	—	✓(C^3)	—	—	✓	—	✓(C^3)
Na et al. (2023a)	✓	—	✓	—	—	✓	—	✓(C^3)
Shi et al. (2022)	✓	—	✓	—	—	(✓)	✓	✓
Curtis et al. (2024b)	✓	—	✓	—	—	(✓)	✓	✓
Na et al. (2023b)	✓	—	✓(C^2)	—	—	(✓)	✓	✓(C^2)
Bollapragada et al. (2023)	✓	—	✓(+ cvx)	—	—	✓	✓	—
Oztoprak et al. (2023)	✓	—	✓	—	✓	✓	—	✓
SSL-ALM Huang et al. (2025)	✓	—	✓	—	✓	(✓)	✓	—
Stoch. Ghost Facchinei & Kungurtsev (2023)	✓	—	✓	—	✓	(✓)	✓	—
Stoch. Switch. Subg. Huang & Lin (2023)	✓	✓	—	—	✓	(✓)	✓	✓

variables $\xi_1, \dots, \xi_J \stackrel{iid}{\sim} \mathcal{P}_\xi$ is called a mini-batch. Inspired by Na et al. (2023a), we use the following notation for the stochastic estimates computed from a mini-batch of size J :

$$\bar{\nabla}^J f(x) = \frac{1}{J} \sum_{j=1}^J \nabla f(x, \xi_j), \quad \bar{c}^J(x) = \frac{1}{J} \sum_{j=1}^J c(x, \xi_j), \quad \bar{\nabla}^J c(x) = \frac{1}{J} \sum_{j=1}^J \nabla c(x, \xi_j). \quad (6)$$

3.1 REVIEW OF METHODS FOR CONSTRAINED ERM

We compare recent constrained optimization algorithms considering a stochastic objective function in Table 3. We note that most of them do not consider the case of stochastic constraints. Among those which do consider stochastic constraints, only three admit inequality constraints. Moreover, with the exception of Huang & Lin (2023), all the algorithms in Table 3 assume F to be at least C^1 , which makes addressing the challenge of nonsmoothness of F infeasible. Davis et al. (2018) leads us to conclude that assuming the objective and constraint functions to be tame and locally Lipschitz is a suitable requirement for solving (5) with theoretical guarantees of convergence. At this point, however, no such algorithm exists, to the best of our knowledge.

Consequently, we consider the practical performance of the algorithms that address the challenges of solving (5) most closely: Stochastic Ghost, SSL-ALM, and Stochastic Switching Subgradient.

3.2 STOCHASTIC GHOST METHOD (STGH)

Facchinei & Kungurtsev (2023) propose the Stochastic Ghost method, that combines a deterministic method for solving (1) (Facchinei et al., 2021) with a stochastic sampling approach for nonlinear maps (Blanchet et al., 2019). The deterministic method of Facchinei et al. (2021) consists in solving subproblem (7) to obtain a direction d , and then to perform a line search. Here, $e \in \mathbb{R}^m$ is a vector with all elements equal to one, τ and $\beta > 0$ are user-prescribed constants, and κ_k is defined as a certain convex combination of optimization subproblems related to C and ∇C . The definition of κ_k enables to expand the feasibility region so that (7) is always feasible. As the problem (1) is stochastic, the subproblem (7) is modified to a stochastic version (8), using the notation in (6):

$$\begin{aligned} \min_d \quad & \nabla F(x_k)^\top d + \frac{\tau}{2} \|d\|^2, \\ \text{s.t.} \quad & C(x_k) + \nabla C(x_k)^\top d \leq \kappa_k e, \\ & \|d\|_\infty \leq \beta, \end{aligned} \quad (7) \quad \begin{aligned} \min_d \quad & \bar{\nabla}^J f(x_k)^\top d + \frac{\tau}{2} \|d\|^2, \\ \text{s.t.} \quad & \bar{c}^J(x_k) + \bar{\nabla}^J c(x_k)^\top d \leq \bar{\kappa}_k^J e, \\ & \|d\|_\infty \leq \beta. \end{aligned} \quad (8)$$

In the stochastic setting (8), an unbiased estimate $d(x_k)$ of the line search direction d is computed using four particular mini-batches as follows. To facilitate comprehension, we denote $X_k^J = \{X_{k,j}\}_{j=1}^J$ a mini-batch of size J with the j -th element $X_{k,j} = (\nabla f(x_k, \xi_{k,j}), c(x_k, \xi_{k,j}), \nabla c(x_k, \xi_{k,j}))$. First,

we sample a random variable $N \sim \mathcal{G}(p_0)$ from the geometric distribution. Then we sample the mini-batches X_k^1 and $X_k^{2^{N+1}}$ and we partition the mini-batch $X_k^{2^{N+1}}$ of size 2^{N+1} into two mini-batches $\text{odd}(X_k^{2^{N+1}})$ and $\text{even}(X_k^{2^{N+1}})$ of size 2^N . Finally, we solve (8) for each of the four mini-batches, denoting by $d(x_k; X_k^J)$ the solution of (8) for the corresponding mini-batch X_k^J . We obtain

$$d(x_k) = \frac{d(x_k; X_k^{2^{N+1}}) - \frac{1}{2} \left(d(x_k; \text{odd}(X_k^{2^{N+1}})) + d(x_k; \text{even}(X_k^{2^{N+1}})) \right)}{(1 - p_0)^N p_0} + d(x_k; X_k^1). \quad (9)$$

An update between the iterations x_k and x_{k+1} is then computed as $x_{k+1} = x_k + \alpha_k d(x_k)$, where the deterministic stepsize α_k should be square-summable $\sum_{k=1}^{\infty} \alpha_k^2 < \infty$ but not summable $\sum_{k=1}^{\infty} \alpha_k = \infty$. For more details, see Algorithm 1 (Appendix C).

3.3 STOCHASTIC SMOOTHED AND LINEARIZED AL METHOD (SSL-ALM)

The Stochastic Smoothed and Linearized AL Method (SSL-ALM) was described in Huang et al. (2025) for optimization problems with stochastic linear constraints. Although problem (1) has non-linear inequality constraints, we use the SSL-ALM due to the lack of algorithms in the literature dealing with stochastic non-linear constraints; see Table 3. The transition between equality and inequality constraints is handled with slack variables. Following the structure of Huang et al. (2025), we minimize over the set $\mathcal{X} = \mathbb{R}^n \times \mathbb{R}_{\geq 0}^m$. The method is based on the augmented Lagrangian (AL) function $L_\rho(x, y) = F(x) + y^\top C(x) + \frac{\rho}{2} \|C(x)\|^2$; see e.g., (Bertsekas & Rheinboldt, 2014). Adding a smoothing term with an additional variable $z \in \mathbb{R}^n$ yields the proximal AL function

$$K_{\rho, \mu}(x, y, z) = L_\rho(x, y) + \frac{\mu}{2} \|x - z\|^2.$$

The SSL-ALM method was originally proposed in Huang et al. (2025) where it is interpreted as an inexact gradient descent step on the Moreau envelope. An important property of the Moreau envelope is that its stationary points coincide with those of the original function.

The strength of this method is that, as opposed to the Stochastic Ghost method, it does not use large mini-batch sizes. In each iteration, we sample $\xi \stackrel{iid}{\sim} \mathcal{P}_\xi$ to evaluate the objective and $\zeta_1, \zeta_2 \stackrel{iid}{\sim} \mathcal{P}_\zeta$ to evaluate the constraint function and its Jacobian matrix, respectively. The function

$$G(x, y, z; \xi, \zeta_1, \zeta_2) = \nabla f(x, \xi) + \nabla c(x, \zeta_1)^\top y + \rho \nabla c(x, \zeta_1)^\top c(x, \zeta_2) + \mu(x - z) \quad (10)$$

is defined so that, in iteration k , $\mathbb{E}_{\xi, \zeta_1, \zeta_2} [G(x_k, y_{k+1}, z_k; \xi, \zeta_1, \zeta_2)] = \nabla K_{\rho, \mu}(x_k, y_{k+1}, z_k)$. Denoting η, τ , and β positive parameters, the update is

$$\begin{aligned} y_{k+1} &= y_k + \eta c(x, \zeta_1), \\ x_{k+1} &= \text{proj}_{\mathcal{X}}(x_k - \tau G(x_k, y_{k+1}, z_k; \xi, \zeta_1, \zeta_2)), \\ z_{k+1} &= z_k + \beta(x_k - z_k). \end{aligned} \quad (11)$$

For more details, see Algorithm 2 (Appendix C).

3.4 STOCHASTIC SWITCHING SUBGRADIENT METHOD (SSw)

The Stochastic Switching Subgradient method was described in Huang & Lin (2023) for optimization over a closed convex set $\mathcal{X} \subset \mathbb{R}^d$ which is easy to project on. It allows for weakly-convex, possibly nonsmooth, objective and constraint functions. They consider subgradients instead of gradients.

The algorithm relies on a prescribed sequence of infeasibility tolerances ϵ_k and of stepsizes η_k^f and η_k^c . At iteration k , we sample $\zeta_1, \dots, \zeta_J \stackrel{iid}{\sim} \mathcal{P}_\zeta$ to compute $\bar{c}^J(x_k)$. If $\bar{c}^J(x_k)$ is smaller than ϵ_k , we sample $\xi \stackrel{iid}{\sim} \mathcal{P}_\xi$ and update using a stochastic estimate $S^f(x_k, \xi)$ of a subgradient of F :

$$x_{k+1} = \text{proj}_{\mathcal{X}}(x_k - \eta_k^f S^f(x_k, \xi)).$$

If not, we sample $\zeta \stackrel{iid}{\sim} \mathcal{P}_\zeta$ and update using a stochastic estimate $S^c(x_k, \zeta)$ of a subgradient of C :

$$x_{k+1} = \text{proj}_{\mathcal{X}}(x_k - \eta_k^c S^c(x_k, \zeta)).$$

324 In either case, the updates are only saved starting from a prescribed index k_0 and the final output is
325 sampled randomly from the saved updates. The algorithm presented here is slightly more general
326 than the one presented in Huang & Lin (2023): we allow for different stepsizes for the objective and
327 the constraint update, while the original method employs equal stepsizes $\eta_k^f = \eta_k^c$. For more details,
328 see Algorithm 3 (Appendix C).
329

330 **4 EXPERIMENTAL EVALUATION**
331

332 In this section, we illustrate the presented algorithms on a real-world instance of the ACS dataset,
333 comparing how they fare with optimization and fairness metrics.
334

335 **4.1 DATASET FOR FAIR ML**
336

337 Ding et al. (2021) proposed a large-scale dataset for fair Machine Learning, based on the ACS PUMS
338 data sample (American Community Survey Public Use Microdata Sample). The ACS survey is sent
339 annually to approximately 3.5 million US households in order to gather information on features such
340 as ancestry, citizenship, education, employment, or income. Therefore, it has the potential to give rise
341 to large-scale learning and optimization problems.

342 In our experiments, we use the ACSIncome dataset, and choose the binary classification task of
343 predicting whether an individual’s income is over \$50,000.
344

345 **4.2 EXPERIMENTS**
346

347 **Numerical setup.** Experiments are conducted on an Asus Zenbook UX535 laptop with AMD
348 Ryzen 7 5800H CPU, and 16GB RAM, using Python with the PyTorch package (Paszke et al., 2019).
349

350 **4.2.1 BINARY PROTECTED ATTRIBUTE**
351

352 **Dataset and problems** We use the ACSIncome dataset over the state of Oklahoma. The dataset
353 contains 9 features and 17,917 data points, and may be accessed via the Python package Folktables.
354 We choose race (**RAC1P**) as the protected attribute. In the original dataset, it is a categorical variable
355 with 9 values. For the purposes of this experiment, we binarize it to obtain the non-protected group of
356 “white” people and the protected group of “non-white” people. The dataset is split randomly into train
357 (80%, 14,333 points) validation (10%, 1,792), and test (10%, 1,792 points) subsets and it is stratified
358 with respect to the protected attribute, i.e., the proportion of “white” and “non-white” samples in the
359 training, validation, and test sets is equivalent to that in the full dataset (30.8% of positive labels in
360 group “white”, 20.7% in the group “non-white”). The protected attributes are then removed from the
361 data so that the model cannot learn from them directly. The data is normalized using Scikit-Learn
362 StandardScaler.

363 Note that ACSIncome is a real-world dataset for which ERM-based predictors without fairness
364 safeguards are known to learn biases (Han et al., 2023). Accordingly, Table 4 (line 1) shows that an
365 ERM predictor without fairness safeguards has poor fairness metrics; see also Figure 4.
366

367 **Problems.** We consider the constrained ERM problem (4) without any regularization $\mathcal{R} = 0$. As
368 our data is divided into just two groups, we constrain the difference between the loss values ℓ^{si}
369 directly, instead of taking the average. In addition, we consider as baselines the ERM problem (2)
370 without any regularization, $\mathcal{R} = 0$, and with a fairness inducing regularizer \mathcal{R} that promotes small
371 difference in accuracy between groups, provided by the Fairret library (Buyl et al., 2024). In all
372 problems, we take as loss function the Binary Cross Entropy with Logits Loss

373
$$\ell(f_\theta(X_i), Y_i) = -Y_i \cdot \log \sigma(f_\theta(X_i)) - (1 - Y_i) \cdot \log(1 - \sigma(f_\theta(X_i))), \quad (12)$$

374 where $\sigma(z) = \frac{1}{1+e^{-z}}$ is the sigmoid function, and the prediction function f_θ is a neural network with
375 2 interconnected hidden layers of sizes 64 and 32 and ReLU activation, with a total of 194 parameters.
376

377 **Algorithms and parameters.** We assess the performance of four algorithms for solving the
378 constrained problem (4):




Figure 1: Train (blue) and test (orange) statistics over time (s) on the ACS Income dataset for each algorithm: SGD (column 1), fairret-regularized SGD (column 2), SSL-ALM (column 3), ALM (column 4) Switching Subgradient (column 5), and Stochastic Ghost (column 6). The plots depict the mean values for loss (first row) and the constraint at each timestamp, rounded to the nearest 0.5 seconds, over 10 runs. The shaded area depicts the region between the first and third quartiles.

- Stochastic Ghost (StGh) (Sec. 3.2 - parameters $\alpha_0 = 0.05$, $\rho = 0.8$, $\tau = 1$, $\beta = 20$, $\lambda = 0.5$, $\hat{\alpha} = 0.05$)
- SSL-ALM (Sec. 3.3 - parameters $\mu = 2.0$, $\rho = 1.0$, $\tau = 0.01$, $\eta = 0.05$, $\beta = 0.5$, $M_y = 10$),
- plain Augmented Lagrangian Method ALM (Sec. 3.3, smoothing term removed $\mu = 0$, otherwise the same setting as SSL-ALM),
- Stochastic Switching Subgradient (SSw) (Sec. 3.4 - $\eta^f = 0.05$, $\eta^c = 0.04$, $\epsilon_0 = 0.01$, $\epsilon_k = \frac{\epsilon_0}{\sqrt{k+1}}$).

The hyperparameter values were chosen by grid-search on a validation set; we report the performance of the algorithms under different hyperparameter choices in appendix B.

We also provide the behavior of SGD for solving the ERM problem, both with no fairness safeguards (SGD), and with fairness regularization on accuracy as provided by the Fairret library (Buyl et al., 2024) (SGD-Fairret). These methods serve as baselines. When estimating the constraints, we sample an equal number of data points for every subgroup.

Optimization performance. Figure 1 present the evolution of loss and constraint values over the train and test datasets for the four algorithms addressing the constrained problem (columns 3–6), as well as for the two baselines: SGD without fairness (col. 1), and SGD with fairness regularization (col. 2). Each algorithm is run 10 times, and the plots display the mean and quartiles values.

To a certain extent, the four algorithms (col. 3–6) succeed in minimizing the loss and satisfying the constraints on the train set. The AL-based methods (col. 3 and 4) demonstrate a better behavior compared to StGh (col. 5) and SSw (col. 6). Indeed, StGh exhibits higher variability in both loss and constraint values. **SSw satisfies the constraint the best out of all methods, but fails to minimize the objective function to the extent that other algorithms do. See appendix B for exploration of the algorithms’ behaviour under different hyperparameter choices.**

SGD (col. 1) exhibits the lowest variability in the trajectory and minimizes the loss in least time, but, as expected, does not satisfy the constraints. Fairret-regularized SGD (col. 2) minimizes the objective function slower than some constrained algorithms, while maintaining a small, but increasing, constraint violation. Note that the penalty parameter was optimized for constraint satisfaction; other penalty values would allow a faster minimization of the objective at the cost of higher constraint violation; see appendix B more for details. This observation is consistent with Ramirez et al. (2025).

The ALM and SSL-ALM schemes satisfy the constraints on the train set. On the test set, however, they are slightly biased towards negative values. Such bias is expected on unseen data and reflects the generalization behavior of fairness-constrained estimators. This is beyond the scope of the current work; see e.g. Chamon et al. (2022).

Fairness performance. Table 4 displays the fairness metrics presented in Section 2: independence (Ind), separation (Sp), and sufficiency (Sf), along with inaccuracy (Ina). The mean value and standard deviation over 10 runs are presented for the four fairness-constrained models and the two baselines,

432 Table 4: Fairness metrics (independence, separation, sufficiency), inaccuracy, and Wasserstein
433 distances between groups (Wd) for the four constrained estimators and the two baselines.

Algname	Train					Test				
	Ind	Sp	Ina	Sf	Wd	Ind	Sp	Ina	Sf	Wd
SGD	0.095 \pm 0.003	0.124 \pm 0.006	0.186\pm0.023	0.065 \pm 0.004	0.062 \pm 0.012	0.098 \pm 0.005	0.155 \pm 0.015	0.209\pm0.019	0.061 \pm 0.005	0.183 \pm 0.015
StGh	0.082 \pm 0.020	0.103 \pm 0.052	0.230 \pm 0.025	0.048\pm0.022	0.037 \pm 0.038	0.086 \pm 0.024	0.123 \pm 0.050	0.239 \pm 0.020	0.057 \pm 0.021	0.161 \pm 0.053
ALM	0.083 \pm 0.009	0.112 \pm 0.024	0.210 \pm 0.010	0.057 \pm 0.011	0.063 \pm 0.021	0.058\pm0.012	0.114 \pm 0.014	0.244 \pm 0.007	0.221 \pm 0.017	0.158 \pm 0.027
SSL-ALM	0.074\pm0.005	0.091\pm0.010	0.208 \pm 0.009	0.050 \pm 0.009	0.054 \pm 0.014	0.083 \pm 0.006	0.108\pm0.018	0.223 \pm 0.013	0.046\pm0.012	0.170 \pm 0.033
SSw	0.096 \pm 0.010	0.139 \pm 0.011	0.191 \pm 0.020	0.064 \pm 0.007	0.001\pm0.001	0.103 \pm 0.018	0.168 \pm 0.036	0.212 \pm 0.020	0.066 \pm 0.019	0.018\pm0.005
SGD-Fairret	0.091 \pm 0.011	0.128 \pm 0.016	0.190 \pm 0.020	0.059 \pm 0.010	0.004 \pm 0.003	0.091 \pm 0.016	0.141 \pm 0.028	0.211 \pm 0.019	0.056 \pm 0.014	0.059 \pm 0.021

441 both on train and test sets. For all metrics, smaller is better. [We provide additional details in](#)
442 [appendices A and B.1.](#)

443 Among the four fairness-constrained models, StGh performs best in terms of sufficiency, but worst in
444 terms of accuracy. Overall, the constrained optimization models improve on the fairness metrics of
445 the unconstrained SGD and the regularized SGD-Fairret models. SSw has both fairness metrics and
446 inaccuracy comparable to that of the unconstrained SGD model. This is consistent with the observation
447 that the optimization method, with our choice of parameters, favored minimizing the objective over
448 satisfying the constraints. The ALM and SSL-ALM methods provide the best compromise: they
449 improve independence, separation, and sufficiency relative to the SGD model, while moderately
450 degrading accuracy. SGD-Fairret slightly improves sufficiency relative to the SGD model. Similar
451 observations hold for metrics on the test set.

452 4.2.2 MULTI-VALUED PROTECTED ATTRIBUTE

453 **Dataset and problems.** We again consider the dataset ACSIncome, but over the state of Virginia,
454 and choose *Mariage* as the protected attribute. This attribute takes five values, as opposed to the
455 binary attribute setup of section 4.2.1.

456 We consider three optimization problems as approaches to tackle the learning task. First, we consider
457 the *constrained* learning problem as described in eq. (4), with $m = 5$. Second, we consider the
458 unconstrained, but *penalized*, problem

$$461 \min_{\theta \in \mathbb{R}^n} \frac{1}{N} \sum_{i=1}^N \ell(f_\theta(X_i), Y_i) + \mathcal{R}(\theta) + \lambda \sum_{i=1}^m \left| \ell^{s_i}(\theta) - \frac{1}{m} \sum_{j=1}^m \ell^{s_j}(\theta) \right|, \quad (13)$$

462 where λ is a penalization weight. Third, we consider the unconstrained and unpenalized problem, as
463 described in eq. (2), for comparison.

464 **Algorithms and parameters.** We solve the constrained learning problem (4) with Stochastic Ghost,
465 Switching Subgradient, and SSL-ALM. We solve the baseline penalized problem (13) and the basic
466 unconstrained unpenalized problem (2) using SGD.

467 The hyperparameters for each algorithm were tuned on the validation set; we picked the values
468 resulting in lowest loss and constraint satisfaction after 60 seconds. Our hyperparameter choices are:

- 469 • Regularized SGD: $\lambda = 0.4$.
- 470 • SSL-ALM: $\tau = \eta = 0.01, \beta = 0.5, \mu = 2, \rho = 1$.
- 471 • Stochastic Ghost: $\beta = 1.0, \gamma_0 = 0.005, \zeta = 0.05, \rho = 0.1, \tau = 1, \lambda = 0.5$.
- 472 • SSw: $\eta^f = 0.05$ constant, η_k^c diminishing with $\eta_0^c = 0.25, \eta_k^c = \frac{\eta_{k-1}^c}{\sqrt{k}}, k > 0$; constraint tolerance
473 ϵ diminishing with $\epsilon_0 = 0.01, \epsilon_k = \frac{\epsilon_{k-1}}{\sqrt{k}}, k > 0$.

474 **Optimization performance.** We report in Figure 2 the evolution of the mean and quartiles of the
475 train and test values over 10 runs. SGD on the unconstrained and unpenalized problem (2) (first row)
476 converges to a model such that the constraint are consistently above the constraint bound for three
477 values of the protected attribute (Wid, Div, and Nev). SGD on the penalized problem (second row)
478 manages to meet all constraints for the training set, but constraint Div on the test set is eventually
479 violated. The three constrained methods minimize function values while keeping with the constraint
480 bounds. The performance of SGD on penalized problem and the three constrained algorithms is
481 comparable.




Figure 2: Train (blue) and test (orange) statistics over time (s) on the ACS Income dataset for each algorithm: SGD (column 1), regularized SGD (column 2), SSL-ALM (column 3), Switching Subgradient (column 4), and Stochastic Ghost (column 5). The plots depict the mean values for loss (first row) and constraints (second to last row) at each timestamp, rounded to the nearest 0.5 seconds, over 10 runs. The shaded area depicts the region between the first and third quartiles.

Note that the penalized problem required consequent preliminary computations in order to tune the penalization parameter λ . We found that the performance of the estimator was sensitive to the value of λ in two aspects: (i) on the difficulty of the optimization problem for SGD, and (ii) in the trade-off between minimization of the empirical risk and satisfaction of constraints; see appendix B.2 for details. In contrast, the algorithms for constrained minimization showed (i) a smaller sensitivity to their hyperparameters, particularly so for SSL-ALM, and (ii) a controllable trade-off between minimization of the empirical risk, and constraint satisfaction; again, see appendix B.2 for details. This observation is consistent with the argument of Ramirez et al. (2025).

5 CONCLUSION

To the best of our knowledge, this paper provides the first benchmark for assessing the performance of optimization methods on real-world instances of fairness constrained training of models. We highlight the challenges of this approach, namely that objective and constraints are non-convex, non-smooth, and large-scale, and review the performance of four practical algorithms.

LIMITATIONS

Our work identifies that there is currently no algorithm with guarantees for solving the fairness constrained problem. Above all, we hope that this work, along with the Python toolbox for easy benchmarking of new optimization methods, will stimulate further interest in this topic. Also, we caution readers that the method present here is not a silver-bullet that handles all biases and ethical issues of training ML models. In particular, care must be taken that fair ML is part of a

540 interdisciplinary pipeline that integrates the specifics of the use-case, and that it does not serve as an
541 excuse for pursuing Business-As-Usual policies that fail to tackle ethical issues (Balayn et al., 2023;
542 Wachter et al., 2021).

543

544 REPRODUCIBILITY

545 Code to reproduce the experiments is provided in the Supplementary Material. This includes a readme
546 file with instructions to reproduce experiments. Details on the computing environment are provided
547 in Section 4.

549

550 REFERENCES

551 Tameem Adel, Isabel Valera, Zoubin Ghahramani, and Adrian Weller. One-network adversarial
552 fairness. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 2412–
553 2420, 2019.

554 Agathe Balayn, Mireia Yurrita, Jie Yang, and Ujwal Gadiraju. “ \checkmark Fairness Toolkits, A Checkbox
555 Culture?” On the Factors that Fragment Developer Practices in Handling Algorithmic Harms. In
556 *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, AIES ’23, pp. 482–495,
557 New York, NY, USA, August 2023. Association for Computing Machinery. ISBN 9798400702310.
558 doi: 10.1145/3600211.3604674.

559 Solon Barocas, Moritz Hardt, and Arvind Narayanan. *Fairness and Machine Learning: Limitations
560 and Opportunities*. MIT Press, 2023.

561 Rachel K. E. Bellamy, Kuntal Dey, Michael Hind, Samuel C. Hoffman, Stephanie Houde, Kalapriya
562 Kannan, Pranay Lohia, Jacquelyn Martino, Sameep Mehta, Aleksandra Mojsilovic, Seema Nagar,
563 Karthikeyan Natesan Ramamurthy, John Richards, Diptikalyan Saha, Prasanna Sattigeri, Moninder
564 Singh, Kush R. Varshney, and Yunfeng Zhang. AI Fairness 360: An Extensible Toolkit for
565 Detecting, Understanding, and Mitigating Unwanted Algorithmic Bias, October 2018.

566 Albert Berahas, Frank E. Curtis, Daniel Robinson, and Baoyu Zhou. Sequential quadratic optimization
567 for nonlinear equality constrained stochastic optimization. *SIAM Journal on Optimization*, 31:
568 1352–1379, 05 2021. doi: 10.1137/20M1354556.

569 Albert S. Berahas, Frank E. Curtis, Michael J. O’Neill, and Daniel P. Robinson. A stochastic
570 sequential quadratic optimization algorithm for nonlinear equality constrained optimization with
571 rank-deficient jacobians, 2023. URL <https://arxiv.org/abs/2106.13015>.

572 D.P. Bertsekas and W. Rheinboldt. *Constrained Optimization and Lagrange Multiplier Methods*.
573 Computer science and applied mathematics. Academic Press, 2014. ISBN 9781483260471.

574 Sarah Bird, Miro Dudík, Richard Edgar, Brandon Horn, Roman Lutz, Vanessa Milan, Mehrnoosh
575 Sameki, Hanna Wallach, and Kathleen Walker. Fairlearn: A toolkit for assessing and improving
576 fairness in AI. Technical Report MSR-TR-2020-32, Microsoft, May 2020.

577 José H. Blanchet, Peter W. Glynn, and Yanan Pei. Unbiased multilevel monte carlo: Stochastic
578 optimization, steady-state simulation, quantiles, and other applications. *arXiv: Statistics Theory*,
579 2019. URL <https://api.semanticscholar.org/CorpusID:127952798>.

580 Raghu Bollapragada, Cem Karanli, Brendan Keith, Boyan Lazarov, Socratis Petrides, and Jingyi
581 Wang. An adaptive sampling augmented lagrangian method for stochastic optimization with
582 deterministic constraints. *Computers and Mathematics with Applications*, 149:239–258, 2023.
583 ISSN 0898-1221. doi: <https://doi.org/10.1016/j.camwa.2023.09.014>. URL <https://www.sciencedirect.com/science/article/pii/S0898122123003991>.

584 Maarten Buyl, Marybeth Defrance, and Tijl De Bie. fairret: a framework for differentiable fairness
585 regularization terms. In *International Conference on Learning Representations*, 2024.

586 Luiz FO Chamon, Santiago Paternain, Miguel Calvo-Fullana, and Alejandro Ribeiro. Constrained
587 learning with non-convex losses. *IEEE Transactions on Information Theory*, 69(3):1739–1760,
588 2022.

594 Changan Chen, Frederick Tung, Naveen Vedula, and Greg Mori. Constraint-aware deep neural
595 network compression. In *Computer Vision – ECCV 2018: 15th European Conference, Munich,
596 Germany, September 8–14, 2018, Proceedings, Part VIII*, pp. 409–424, Berlin, Heidelberg, 2018.
597 Springer-Verlag. ISBN 978-3-030-01236-6. doi: 10.1007/978-3-030-01237-3_25. URL https://doi.org/10.1007/978-3-030-01237-3_25.

599 Richard J. Chen, Judy J. Wang, Drew F. K. Williamson, Tiffany Y. Chen, Jana Lipkova, Ming Y. Lu,
600 Sharifa Sahai, and Faisal Mahmood. Algorithmic fairness in artificial intelligence for medicine and
601 healthcare. *Nature Biomedical Engineering*, 7(6):719–742, Jun 2023. ISSN 2157-846X. doi: 10.
602 1038/s41551-023-01056-8. URL <https://doi.org/10.1038/s41551-023-01056-8>.

604 Zhenpeng Chen, Jie M. Zhang, Max Hort, Mark Harman, and Federica Sarro. Fairness testing: A
605 comprehensive survey and analysis of trends. *ACM Trans. Softw. Eng. Methodol.*, 33(5), June 2024.
606 ISSN 1049-331X. doi: 10.1145/3652155. URL <https://doi.org/10.1145/3652155>.

607 Andrew Cotter, Heinrich Jiang, Serena Wang, Taman Narayan, Seungil You, Karthik Sridharan, and
608 Maya R. Gupta. Optimization with non-differentiable constraints with applications to fairness,
609 recall, churn, and other goals. *Journal of Machine Learning Research*, 20(172):1–59, 2019.

611 Frank E. Curtis, Michael J. O’Neill, and Daniel P. Robinson. Worst-case complexity of an sqp
612 method for nonlinear equality constrained stochastic optimization. *Mathematical Programming*,
613 205(1):431–483, May 2024a. ISSN 1436-4646. doi: 10.1007/s10107-023-01981-1. URL
614 <https://doi.org/10.1007/s10107-023-01981-1>.

616 Frank E. Curtis, Daniel P. Robinson, and Baoyu Zhou. Sequential quadratic optimization for
617 stochastic optimization with deterministic nonlinear inequality and equality constraints. *SIAM
618 Journal on Optimization*, 34(4):3592–3622, 2024b. doi: 10.1137/23M1556149. URL <https://doi.org/10.1137/23M1556149>.

620 Damek Davis, Dmitriy Drusvyatskiy, Sham Kakade, and Jason D. Lee. Stochastic subgradient method
621 converges on tame functions, 2018. URL <https://arxiv.org/abs/1804.07795>.

623 MaryBeth Defrance, Maarten Buyl, and Tijl De Bie. Abcfair: an adaptable benchmark approach for
624 comparing fairness methods, 2024. URL <https://arxiv.org/abs/2409.16965>.

626 Eoin Delaney, Zihao Fu, Sandra Wachter, Brent Mittelstadt, and Chris Russell. OxonFair: A Flexible
627 Toolkit for Algorithmic Fairness, November 2024.

628 Frances Ding, Moritz Hardt, John Miller, and Ludwig Schmidt. Retiring adult: New datasets for fair
629 machine learning. *Advances in Neural Information Processing Systems*, 34, 2021.

631 Mengnan Du, Subhabrata Mukherjee, Guanchu Wang, Ruixiang Tang, Ahmed Hassan Awadallah,
632 and Xia Hu. Fairness via representation neutralization. In *Neurips*, 2021.

634 Harrison Edwards and Amos Storkey. Censoring representations with an adversary, 2016. URL
635 <https://arxiv.org/abs/1511.05897>.

637 Alessandro Fabris, Stefano Messina, Gianmaria Silvello, and Gian Antonio Susto. Algorithmic
638 fairness datasets: the story so far. *Data Mining and Knowledge Discovery*, 36(6):2074–2152, Nov
639 2022. ISSN 1573-756X. doi: 10.1007/s10618-022-00854-z. URL <https://doi.org/10.1007/s10618-022-00854-z>.

641 Francisco Facchinei and Vyacheslav Kungurtsev. Stochastic approximation for expectation objective
642 and expectation inequality-constrained nonconvex optimization, 2023. URL <https://arxiv.org/abs/2307.02943>.

645 Francisco Facchinei, Vyacheslav Kungurtsev, Lorenzo Lamariello, and Gesualdo Scutari. Ghost
646 penalties in nonconvex constrained optimization: Diminishing stepsizes and iteration complexity.
647 *Mathematics of Operations Research*, 46(2):595–627, 2021. doi: 10.1287/moor.2020.1079. URL
648 <https://doi.org/10.1287/moor.2020.1079>.

648 Yuchen Fang, Sen Na, Michael W. Mahoney, and Mladen Kolar. Fully stochastic trust-region
649 sequential quadratic programming for equality-constrained optimization problems. *SIAM Journal*
650 *on Optimization*, 34(2):2007–2037, 2024. doi: 10.1137/22M1537862. URL <https://doi.org/10.1137/22M1537862>.

652 Jose Gallego-Posada, Juan Ramirez, Meraj Hashemizadeh, and Simon Lacoste-Julien. Cooper: A
653 Library for Constrained Optimization in Deep Learning, April 2025.

654

655 Xiaotian Han, Jianfeng Chi, Yu Chen, Qifan Wang, Han Zhao, Na Zou, and Xia Hu. FFB: A Fair
656 Fairness Benchmark for In-Processing Group Fairness Methods. In *The Twelfth International*
657 *Conference on Learning Representations*, October 2023.

658 Moritz Hardt, Eric Price, and Nathan Srebro. Equality of opportunity in supervised learning. In *Pro-*
659 *ceedings of the 30th International Conference on Neural Information Processing Systems*, NIPS’16,
660 pp. 3323–3331, Red Hook, NY, USA, 2016. Curran Associates Inc. ISBN 9781510838819.

661 Ruichuan Huang, Jiawei Zhang, and Ahmet Alacaoglu. Stochastic smoothed primal-dual algorithms
662 for nonconvex optimization with linear inequality constraints, 2025. URL <https://arxiv.org/abs/2504.07607>.

663

664 Yankun Huang and Qihang Lin. Oracle complexity of single-loop switching subgradient methods for
665 non-smooth weakly convex functional constrained optimization, 2023. URL <https://arxiv.org/abs/2301.13314>.

666

667 Toshihiro Kamishima, Shotaro Akaho, Hideki Asoh, and Jun Sakuma. Fairness-Aware Classifier
668 with Prejudice Remover Regularizer. In Peter A. Flach, Tijl De Bie, and Nello Cristianini (eds.),
669 *Machine Learning and Knowledge Discovery in Databases*, pp. 35–50, Berlin, Heidelberg, 2012.
670 Springer. ISBN 978-3-642-33486-3. doi: 10.1007/978-3-642-33486-3_3.

671

672 Michael P. Kim, Amirata Ghorbani, and James Zou. Multiaccuracy: Black-box post-processing
673 for fairness in classification. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics,*
674 *and Society*, AIES ’19, pp. 247–254, New York, NY, USA, 2019. Association for Computing
675 Machinery. ISBN 9781450363242. doi: 10.1145/3306618.3314287. URL <https://doi.org/10.1145/3306618.3314287>.

676

677 Tai Le Quy, Arjun Roy, Vasileios Iosifidis, Wenbin Zhang, and Eirini Ntoutsi. A survey on datasets
678 for fairness-aware machine learning. *WIREs Data Mining and Knowledge Discovery*, 12(3), 03
679 2022. ISSN 1942-4795. doi: 10.1002/widm.1452. URL <https://doi.org/10.1002/widm.1452>.

680

681 Zhu Li, Adrián Pérez-Suay, Gustau Camps-Valls, and Dino Sejdinovic. Kernel dependence regulariz-
682 ers and Gaussian processes with applications to algorithmic fairness. *Pattern Recognition*, 132:
683 108922, December 2022. ISSN 0031-3203. doi: 10.1016/j.patcog.2022.108922.

684

685 Gilles Louppe, Michael Kagan, and Kyle Cranmer. Learning to pivot with adversarial networks.
686 *Advances in neural information processing systems*, 30, 2017.

687

688 David Madras, Elliot Creager, Toniann Pitassi, and Richard Zemel. Learning adversarially fair and
689 transferable representations. In *International Conference on Machine Learning*, pp. 3384–3393.
690 PMLR, 2018.

691

692 Sen Na, Mihai Anitescu, and Mladen Kolar. An adaptive stochastic sequential quadratic program-
693 ming with differentiable exact augmented lagrangians. *Mathematical Programming*, 199(1):
694 721–791, May 2023a. doi: 10.1007/s10107-022-01846-z. URL <https://doi.org/10.1007/s10107-022-01846-z>.

695

696 Sen Na, Mihai Anitescu, and Mladen Kolar. Inequality constrained stochastic nonlinear optimization
697 via active-set sequential quadratic programming, 2023b. URL <https://arxiv.org/abs/2109.11502>.

698

699 Yatin Nandwani, Abhishek Pathak, Mausam, and Parag Singla. A primal dual formulation for
700 deep learning with constraints. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc,
701 E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Cur-
ran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/cf708fc1decf0337aded484f8f4519ae-Paper.pdf.

702 Figen Oztoprak, Richard Byrd, and Jorge Nocedal. Constrained optimization in the presence of
703 noise. *SIAM Journal on Optimization*, 33(3):2118–2136, 2023. doi: 10.1137/21M1450999. URL
704 <https://doi.org/10.1137/21M1450999>.

705 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan,
706 Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas
707 Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy,
708 Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-
709 performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pp.
710 8024–8035. Curran Associates, Inc., 2019. URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.

711 Patricia Pauli, Anne Koch, Julian Berberich, Paul Kohler, and Frank Allgöwer. Training robust neural
712 networks using lipschitz bounds. *IEEE Control Systems Letters*, 6:121–126, 2021.

713 Dana Pessach and Erez Shmueli. A review on fairness in machine learning. *ACM Comput. Surv.*, 55
714 (3), February 2022. ISSN 0360-0300. doi: 10.1145/3494672. URL <https://doi.org/10.1145/3494672>.

715 Evaggelia Pitoura, Kostas Stefanidis, and Georgia Koutrika. Fairness in rankings and recommendations:
716 an overview. *The VLDB Journal*, 31(3):431–458, May 2022. ISSN 0949-877X. doi: 10.1007/s00778-021-00697-y. URL <https://doi.org/10.1007/s00778-021-00697-y>.

717 Juan Ramirez, Meraj Hashemizadeh, and Simon Lacoste-Julien. Position: Adopt Constraints Over
718 Penalties in Deep Learning, July 2025.

719 Sathya N. Ravi, Tuan Dinh, Vishnu Suresh Lokhande, and Vikas Singh. Explicitly imposing
720 constraints in deep networks via conditional gradients gives improved generalization and faster
721 convergence. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):4772–4779, Jul.
722 2019. doi: 10.1609/aaai.v33i01.33014772. URL <https://ojs.aaai.org/index.php/AAAI/article/view/4404>.

723 Qiankun Shi, Xiao Wang, and Hao Wang. A momentum-based linearized augmented lagrangian
724 method for nonconvex constrained stochastic optimization. *Optimization Online*, 2022. URL
725 <https://optimization-online.org/?p=19870>.

726 Amal Tawakuli and Thomas Engel. Make your data fair: A survey of data preprocessing techniques
727 that address biases in data towards fair ai. *Journal of Engineering Research*, 2024. ISSN 2307-1877.
728 doi: <https://doi.org/10.1016/j.jer.2024.06.016>. URL <https://www.sciencedirect.com/science/article/pii/S2307187724001871>.

729 Sahil Verma and Julia Rubin. Fairness definitions explained. In *Proceedings of the International
730 Workshop on Software Fairness*, FairWare ’18, pp. 1–7, New York, NY, USA, 2018. Association
731 for Computing Machinery. ISBN 9781450357463. doi: 10.1145/3194770.3194776. URL
732 <https://doi.org/10.1145/3194770.3194776>.

733 Sandra Wachter, Brent Mittelstadt, and Chris Russell. Bias Preservation in Machine Learning: The
734 Legality of Fairness Metrics Under EU Non-Discrimination Law, January 2021.

735 Mingyang Wan, Daochen Zha, Ninghao Liu, and Na Zou. In-processing modeling techniques for
736 machine learning fairness: A survey. *ACM Trans. Knowl. Discov. Data*, 17(3), March 2023. ISSN
737 1556-4681. doi: 10.1145/3551390. URL <https://doi.org/10.1145/3551390>.

738 Quan Zhou, Jakub Mareček, and Robert Shorten. Subgroup fairness in two-sided markets. *Plos one*,
739 18(2):e0281443, 2023a.

740 Quan Zhou, Jakub Mareček, and Robert Shorten. Fairness in forecasting of observations of linear
741 dynamical systems. *Journal of Artificial Intelligence Research*, 76:1247–1280, April 2023b.
742 ISSN 1076-9757. doi: 10.1613/jair.1.14050. URL <http://dx.doi.org/10.1613/jair.1.14050>.

743


756 A ADDITIONAL DETAILS ON THE FAIRNESS OF THE BINARY EXPERIMENT.

757
 758 Figure 3 presents the distribution of predictions over both groups. The distribution of prediction
 759 without fairness guarantees (col. 1) clearly does not meet the group fairness standard. Indeed, the
 760 “non-white” group has a significantly higher likelihood than the “white” group of receiving small
 761 predicted values, and the converse holds for large predicted values. **Among the fairness-constrained**
 762 **models, the ALM and SSL-ALM distributions are the closest to the distributions of SGD without**
 763 **fairness, which is consistent with retaining good prediction information. The Fairret penalized**
 764 **formulation (col. 2) and SSw (col. 5) have a center-heavy distribution, which, in this case, is evidence**
 765 **of poor objective minimization.**


766 The SSL-ALM, ALM, and Stochastic Ghost (col. 3, 4, and 6) have closer distributions across groups
 767 than the unconstrained formulation.

768 Numerically, this is expressed in Table 4 (col. Wd), which reports the value of the Wasserstein
 769 distance between group distributions for each model.

770 Table 4 displays the fairness metrics presented in Section 2: independence (Ind), separation (Sp), and
 771 sufficiency (Sf), along with inaccuracy (Ina). The mean value and standard deviation over 10 runs are
 772 presented for the four fairness-constrained models and the two baselines, both on train and test sets.
 773 Figure 4 presents the mean values as spider plots. For all metrics, smaller is better.



774
 775
 776
 777
 778
 779
 780 Figure 3: Distribution of predictions for each algorithm. Left to right: SGD, SGD-Fairret, SSL-ALM,
 781 ALM, SSw, StGh, Blue and red denote “white” and “non-white” groups.



781
 782
 783
 784
 785
 786
 787
 788
 789
 790
 791
 792
 793
 794
 795
 796
 797 Figure 4: Average value of the three fairness metrics (independence (Ind), separation (Sp), and
 798 sufficiency (Sf)), along with mean inaccuracy (Ina), and difference in accuracy between the two
 799 groups (DifAcc). For all metrics, smaller values are better.

800 B HYPERPARAMETER SENSITIVITY ANALYSIS

801
 802 In this section, we present the performance of algorithms on the validation set, justifying our
 803 hyperparameter choices in the experimental section. For each algorithm, we provide a table of loss
 804 and constraint violation values under different hyperparameter choices after a training run.

805 For each algorithm, we perform 5 runs with each hyperparameter combination. Every 100 iterations,
 806 we save the model state (except for Stochastic Ghost, where we save every 10 iterations due to higher
 807 iteration cost). We then pick our hyperparameters based on the average loss value and constraint

810 violation over the last 20 states over 5 runs. For algorithms that exhibit exceptionally noisy trajectories,
 811 we also provide convergence plots.
 812

813 **B.1 EXPERIMENT 1 - BINARY ATTRIBUTE**
 814

815 We recall that the constraint bound δ was set to 0.005. With random parameter initialization, initial
 816 loss value was ≈ 0.7 .
 817

818 **B.1.1 STOCHASTIC SWITCHING SUBGRADIENT**
 819

820 Stochastic Switching Subgradient admits a choice of constant, diminishing, or adaptive stepsizes for
 821 both objective and constraint updates (η^f and η^c , respectively). In this experiment, we were unable
 822 to find a combination of the above that would enable the algorithm to both minimize the objective on
 823 par with other methods and maintain the feasibility of the solution. The run results are presented in
 824 table 5.
 825

826 As suggested in the original paper, the diminishing stepsize follows the rule
 827

$$\eta_k = \frac{\eta_0}{\sqrt{k+1}}$$

828 and the adaptive stepsize follows the rule
 829

$$\eta_k = \frac{c(x_k, \zeta)}{\|S^c(x_k, \zeta)\|^2}$$

830 We use diminishing infeasibility tolerance, with $\epsilon_0 = 0.01, \epsilon_k = \frac{\epsilon_0}{\sqrt{k+1}}$
 831

Hyperparameters	Loss		Constraint violation	
	mean	std	mean	std
η^f adaptive, η^c adaptive	0.681	0.006	0.000	0.000
$\eta^f = 0.05, \eta^c$ adaptive	0.679	0.006	0.000	0.000
$\eta^f = 0.05, \eta^c = 0.05$	0.626	0.006	0.000	0.000
$\eta^f = 0.05, \eta^c = 0.045$	0.589	0.01	0.0012	0.002
$\eta^f = 0.05, \eta^c = 0.04$	0.567	0.014	0.003	0.003
$\eta^f = 0.05, \eta^c = 0.025$	0.469	0.012	0.018	0.004
$\eta^f = 0.05, \eta^c$ dimin with $\eta_0^c = 0.05$	0.409	0.001	0.025	0.004
$\eta^f = 0.05, \eta^c$ dimin with $\eta_0^c = 0.25$	0.409	0.001	0.025	0.004
$\eta^f = 0.05, \eta^c$ dimin with $\eta_0^c = 0.5$	0.410	0.002	0.027	0.005
$\eta^f = 0.05, \eta^c$ dimin with $\eta_0^c = 3.0$	0.476	0.016	0.018	0.004
η^f dimin with $\eta_0^f = 0.05, \eta^c$ dimin with $\eta_0^c = 0.05$	0.660	0.002	0.000	0.000
η^f dimin with $\eta_0^f = 0.5, \eta^c$ dimin with $\eta_0^c = 0.5$	0.633	0.004	0.000	0.000

832 Table 5: Loss and constraint violation on the validation set after 5 30-second runs of the Stochastic
 833 Switching Subgradient in the setup of Exp. 1, rounded to 3 digits.
 834

835 We pick $\eta^f = 0.05, \eta^c = 0.04$, as it results in the best objective value with relatively low infeasibility;
 836 we then conduct
 837

838 **B.1.2 SSL-ALM**
 839

840 SSL-ALM allows the user to tune the primal (τ) and dual (η) parameter update rate, the penalty
 841 multiplier ρ , as well as the smoothing update rate β and multiplier μ . We consider only constant
 842 stepsizes τ and η , and fix $\rho = 1$. The run results are presented in table 6.
 843

844 The lowest loss value with the least constraint violation is offered by $\tau = 0.01, \eta = 0.05, \mu = 0$; the
 845 second best combination is $\tau = 0.01, \eta = 0.05, \mu = 2$. As the original paper assumes $\mu \geq 2$, we test
 846 both options as "ALM" and "SSL-ALM", respectively.
 847

	Hyperparameters			Loss		Constraint violation	
	τ	η	μ	mean	std	mean	std
867	0.05	0.05	2	0.4376	0.0283	0.0194	0.0294
868	0.05	0.01	2	0.4240	0.0096	0.0112	0.0097
869	0.01	0.05	2	0.4465	0.0122	0.0074	0.0085
870	0.01	0.01	2	0.4416	0.0138	0.0133	0.0145
871	0.05	0.05	0	0.4560	0.0962	0.0282	0.0620
872	0.05	0.01	0	0.4222	0.0099	0.0111	0.0100
873	0.01	0.05	0	0.4448	0.0114	0.0070	0.0078
874	0.01	0.01	0	0.4407	0.0155	0.0157	0.0154
875	0.05	0.05	4	0.4377	0.0242	0.0180	0.0271
876	0.05	0.01	4	0.4252	0.0098	0.0121	0.0095
877	0.01	0.05	4	0.4471	0.0112	0.0073	0.0077
878	0.01	0.01	4	0.4387	0.0132	0.0143	0.0099

Table 6: Loss and constraint violation on the validation set after 5 30-second runs of the SSL-ALM variants in the setup of Exp. 1, rounded to 4 digits.




Figure 5: Loss (top row) and constraint (bottom row) evolution of the Stochastic Ghost on the validation dataset with different hyperparameter choices. The line corresponds to the mean value over 5 runs, the shaded region - to the area between the 1st and 3rd quartiles over 5 runs.

B.1.3 STOCHASTIC GHOST

Among the algorithms discussed, Stochastic Ghost features the largest number of hyperparameters, which makes the algorithm difficult to tune. We set p_0 , the parameter of the geometric distribution that controls the number of samples taken at each iteration, to 0.4, as, in expectation, this matches the batch size used for the other algorithms. As in the original paper Facchinei & Kungurtsev (2023), we use the OSQP solver to solve the subproblems.

We attempt to tune the stepsize α_0 , the stepsize decay rate $\hat{\alpha}$, and two subproblem parameters β and ρ .

The run results are presented in table 7. Due to the level of noise in the algorithm's trajectories, we add the plots for some of the more stable hyperparameter configurations in fig. 5.

We pick the values $\alpha_0 = 0.05$, $\hat{\alpha} = 0.05$, $\rho = 0.1$, $\beta = 20$.

B.1.4 SGD+FAIRRET

For regularized SGD, we tune the regularization penalty multiplier; once we find a value that leads to minimizing both loss and constraint violation, we try different stepsize values.

The run results are presented in table 8.

We pick a multiplier equal to 3 and LR of 0.05.

	Hyperparameters				Loss		Constraint violation	
	α_0	$\hat{\alpha}$	ρ	β	mean	std	mean	std
921	0.0100	0.0500	0.1000	10	0.6144	0.0287	0.0135	0.0115
922	0.0100	0.0500	0.1000	20	0.6390	0.0461	0.0287	0.0297
923	0.0100	0.0500	0.1000	5	0.6007	0.0448	0.0256	0.0170
924	0.0250	0.0250	0.1000	10	0.4781	0.0245	0.0504	0.0183
925	0.0250	0.0250	0.1000	20	0.4674	0.0169	0.0381	0.0175
926	0.0250	0.0250	0.1000	5	0.4896	0.0371	0.0367	0.0238
927	0.0250	0.0500	0.1000	10	0.5041	0.0402	0.0318	0.0249
928	0.0250	0.0500	0.1000	20	0.5246	0.0566	0.0326	0.0236
929	0.0250	0.0500	0.1000	5	0.5218	0.0301	0.0343	0.0229
930	0.0500	0.0500	0.1000	10	0.5054	0.0295	0.0280	0.0196
931	0.0500	0.0500	0.1000	20	0.4745	0.0266	0.0152	0.0154
932	0.0500	0.0500	0.1000	5	0.4800	0.0338	0.0442	0.0243
933	0.0500	0.0500	0.8000	10	0.5700	0.0432	0.0441	0.0375
934	0.0500	0.0500	0.8000	20	0.6124	0.1448	0.0851	0.0744
935	0.0500	0.0500	0.8000	5	0.6139	0.0783	0.0647	0.0618
936	0.0500	0.1000	0.1000	10	0.5136	0.0413	0.0486	0.0297
937	0.0500	0.1000	0.1000	20	0.5124	0.0620	0.0379	0.0384
938	0.0500	0.1000	0.1000	5	0.5131	0.0505	0.0356	0.0197

Table 7: Loss and constraint violation on the validation set after 5 30-second runs of the Stochastic Ghost in the setup of Exp. 1, rounded to 4 digits.

Multiplier	LR	Hyperparameters		Loss		Constraint violation	
		mean	std	mean	std	mean	std
0	0.05	0.406	0.003	0.020	0.004		
1	0.05	0.406	0.004	0.010	0.007		
2	0.05	0.421	0.008	0.012	0.006		
3	0.05	0.539	0.026	0.004	0.005		
3	0.07	0.523	0.029	0.006	0.008		
3	0.1	0.523	0.033	0.007	0.008		
4	0.05	0.676	0.009	0.001	0.002		

Table 8: Loss and constraint violation on the validation set after 5 30-second runs of fairret-regularized SGD in the setup of Exp. 1, rounded to 3 digits.

B.2 EXPERIMENT 2 - MULTI-VALUED ATTRIBUTE

We recall that in this experiment, we set the constraint bound to be 0.05. We test Regularized SGD, SSL-ALM, Stochastic Ghost, and Switching Subgradient.

B.2.1 STOCHASTIC SWITCHING SUBGRADIENT

The run results are presented in table 9. We pick $\eta^f = 0.05$ constant, η_k^c diminishing with $\eta_0^c = 0.25$ as it offers the lowest loss value while being close to feasibility.

B.2.2 SSL-ALM

Unlike Exp. 1, we do not tune the μ hyperparameter, setting $\mu = 2$. The run results are presented in appendix B.2.2. We pick $\tau = 0.05$, $\eta = 0.01$.

B.2.3 STOCHASTIC GHOST

The run results are presented in appendix B.2.3. We pick $\alpha_0 = 0.01$, $\hat{\alpha} = 0.2$, $\rho = 0.1$, $\beta = 5$.

	Hyperparameters				Loss		Max constraint violation	
	η_0^f	η^f rule	η_0^c	η^c rule	mean	std	mean	std
0.05	const	0.05	const	0.598	0.002	0.000	0.000	
0.5	const	0.05	const	0.537	0.007	0.000	0.000	
0.05	const	0.5	const	0.631	0.004	0.000	0.000	
0.5	const	0.5	const	0.590	0.006	0.000	0.000	
0.05	const	0.05	dimin	0.543	0.018	0.031	0.001	
0.005	const	0.05	dimin	0.551	0.002	0.000	0.000	
0.05	const	0.05	dimin	0.499	0.003	0.023	0.003	
0.05	const	0.1	dimin	0.506	0.006	0.011	0.003	
0.05	const	0.25	dimin	0.519	0.004	0.001	0.002	

Table 9: Loss and constraint violation on the validation set after 5 60-second runs of the Switching Subgradient in the setup of Exp. 2, rounded to 3 digits.

	Hyperparameters		Loss		Max constraint violation	
	τ	η	mean	std	mean	std
0.05	0.05	0.536	0.007	0.000	0.001	
0.05	0.01	0.530	0.006	0.000	0.001	
0.01	0.05	0.542	0.003	0.000	0.000	
0.01	0.01	0.538	0.003	0.000	0.000	

Table 10: Loss and constraint violation on the validation set after 5 60-second runs of the SSL-ALM in the setup of Exp. 2, rounded to 4 digits.

	Hyperparameters				Loss		Constraint violation	
	α_0	$\hat{\alpha}$	ρ	β	mean	std	mean	std
0.0050	0.0100	0.1000	5	0.5948	0.0218	0.0016	0.0031	
0.0050	0.0100	0.1000	10	0.6149	0.0193	0.0005	0.0012	
0.0050	0.0500	0.1000	5	0.6161	0.0142	0.0001	0.0006	
0.0050	0.0500	0.1000	10	0.6187	0.0191	0.0016	0.0033	
0.0100	0.0100	0.1000	5	0.5373	0.0243	0.0275	0.0296	
0.0100	0.0100	0.1000	10	0.5321	0.0280	0.0290	0.0259	
0.0100	0.0500	0.1000	5	0.5442	0.0222	0.0180	0.0196	
0.0100	0.0500	0.1000	10	0.5583	0.0449	0.0099	0.0166	
0.0100	0.0500	0.1000	20	0.5498	0.0311	0.0106	0.0176	
0.0100	0.0500	0.8000	10	0.5716	0.0235	0.0096	0.0123	
0.0100	0.0500	0.8000	20	0.5571	0.0409	0.0181	0.0184	
0.0100	0.1000	0.1000	5	0.5448	0.0297	0.0197	0.0198	
0.0100	0.1000	0.1000	10	0.5522	0.0213	0.0146	0.0108	
0.0100	0.2000	0.1000	5	0.5868	0.0181	0.0070	0.0071	
0.0100	0.2000	0.1000	10	0.5898	0.0198	0.0016	0.0035	
0.0100	0.2000	0.8000	5	0.5616	0.0330	0.0121	0.0119	
0.0500	0.2000	0.1000	5	0.5144	0.0296	0.0199	0.0255	

Table 11: Loss and constraint violation on the validation set after 5 60-second runs of the Stochastic Ghost in the setup of Exp. 2, rounded to 4 digits.

B.2.4 REGULARIZED SGD

The run results are presented in appendix B.2.4. We pick $\lambda = 0.4$.

Penalty multiplier	Loss		Constraint violation	
	mean	std	mean	std
0	0.432	0.001	0.075	0.004
0.3	0.475	0.002	0.030	0.003
0.4	0.512	0.002	0.001	0.002
0.5	0.550	0.002	0.000	0.000

Table 12: Loss and constraint violation on the validation set after 5 60-second runs of the Regularized SGD in the setup of Exp. 2, rounded to 4 digits.

C ALGORITHMS IN MORE DETAIL

In this section, we provide the pseudocodes of algorithms presented in Section 3 as Algorithms 1 to 3. Recall that we denote by $X_k^J = \{X_{k,j}\}_{j=1}^J$ a mini-batch of size J with the j -th element

$$X_{k,j} = (\nabla f(x_k, \xi_{k,j}), c(x_k, \zeta_{k,j}), \nabla c(x_k, \zeta_{k,j})). \quad (14)$$

1080

1081

Algorithm 1 Stochastic Ghost algorithm

1083 **Require:** Training dataset \mathcal{D} , constraint dataset \mathcal{C} , initial neural network weights x_0 1084 **Require:** Parameters $p_0 \in (0, 1)$, $\alpha_0, \hat{\alpha}, \rho, \tau, \beta$ 1085 1: **for** Iteration $k = 0$ **to** $K - 1$ **do**1086 2: Sample $\xi \stackrel{iid}{\sim} \mathcal{P}_\xi$ and $\zeta \stackrel{iid}{\sim} \mathcal{P}_\zeta$ 1087 3: Sample $N \sim \mathcal{G}(p_0)$ 1088 4: Set $J = 2^{N+1}$ 1089 5: Sample a mini-batch $\{\zeta_j\}_{j=1}^J$ so that $\zeta_1, \dots, \zeta_J \stackrel{iid}{\sim} \mathcal{P}_\zeta$ 1090 6: Sample a mini-batch $\{\xi_j\}_{j=1}^J$ so that $\xi_1, \dots, \xi_J \stackrel{iid}{\sim} \mathcal{P}_\xi$ 1091 7: Set X_k^1 and $X_k^{2^{N+1}}$ using (14)1092 8: Compute $d(x_k)$ from (9)1093 9: Set $\alpha_k = \alpha_{k-1}(1 - \hat{\alpha}\alpha_{k-1})$ 1094 10: Update $x_{k+1} = x_k + \alpha_k d(x_k)$ 1095 11: **end for**

1096

1097

1098

1099

Algorithm 2 Stochastic Smoothed and Linearized AL Method for solving (1)

1100 **Require:** Training dataset \mathcal{D} , constraint dataset \mathcal{C} , initial neural network weights x_0 1101 **Require:** Parameters $\mu, \eta, M_y > 0, \tau, \beta, \rho \geq 0$ 1102 1: **for** Iteration $k = 0$ **to** $K - 1$ **do**1103 2: Sample $\xi \stackrel{iid}{\sim} \mathcal{P}_\xi$ and $\zeta_1, \zeta_2 \stackrel{iid}{\sim} \mathcal{P}_\zeta$ 1104 3: $y_{k+1} = y_k + \eta c(x, \zeta_1)$ 1105 4: **if** $\|y_{k+1}\| \geq M_y$ **then**1106 5: $y_{k+1} = 0$ 1107 6: **end if**1108 7: $x_{k+1} = \text{proj}_{\mathcal{X}}(x_k - \tau G(x_k, y_{k+1}, z_k; \xi, \zeta_1, \zeta_2))$, where G is defined in (10)1109 8: $z_{k+1} = z_k + \beta(x_k - z_k)$ 1110 9: **end for**

1111

1112

1113

1114

1115

Algorithm 3 Stochastic Switching Subgradient Method

1116 **Require:** Training dataset \mathcal{D} , constraint dataset \mathcal{C} , initial neural network weights $x_0 \in \mathcal{X}$ 1117 **Require:** Total number of iterations K , sequence of tolerances of infeasibility $\epsilon_k \geq 0$, sequences of
1118 stepsizes η_k^f and η_k^c , mini-batch size J , starting index k_0 for recording outputs, $I = \emptyset$ 1119 1: **for** Iteration $k = 0$ **to** $K - 1$ **do**1120 2: Sample a mini-batch $\{\zeta_j\}_{j=1}^J$ so that $\zeta_1, \dots, \zeta_J \stackrel{iid}{\sim} \mathcal{P}_\zeta$ 1121 3: Set $\bar{c}^J(x_k) = \frac{1}{J} \sum_{j=1}^J c(x_k, \zeta_j)$ 1122 4: **if** $\bar{c}^J(x_k) \leq \epsilon_k$ **then**1123 5: Sample $\xi \stackrel{iid}{\sim} \mathcal{P}_\xi$ and generate $S^f(x_k, \xi)$ 1124 6: Set $x_{k+1} = \text{proj}_{\mathcal{X}}(x_k - \eta_k^f S^f(x_k, \xi))$ and, if $k \geq k_0$, $I = I \cup \{k\}$ 1125 7: **else**1126 8: Sample $\zeta \stackrel{iid}{\sim} \mathcal{P}_\zeta$ and generate $S^c(x_k, \zeta)$ 1127 9: Set $x_{k+1} = \text{proj}_{\mathcal{X}}(x_k - \eta_k^c S^c(x_k, \zeta))$ and, if $k \geq k_0$, $I = I \cup \{k\}$ 1128 10: **end if**1129 11: **end for**1130 12: **Output:** x_τ with τ randomly sampled from I using $P(\tau = k) = \frac{\eta_k}{\sum_{s \in I} \eta_s}$.

1131

1132

1133