

COMPOSITE DIFFERENCE-MAX PROGRAMS FOR MODERN
STATISTICAL ESTIMATION PROBLEMS*YING CUI[†], JONG-SHI PANG[†], AND BODHISATTVA SEN[‡]

Abstract. Many modern statistical estimation problems are defined by three major components: a statistical model that postulates the dependence of an output variable on the input features; a loss function measuring the error between the observed output and the model predicted output; and a regularizer that controls the overfitting and/or variable selection in the model. We study the sampling version of this generic statistical estimation problem where the model parameters are estimated by empirical risk minimization, which involves the minimization of the empirical average of the loss function at the data points weighted by the model regularizer. In our setup we allow all three component functions discussed above to be of the difference-of-convex (dc) type and illustrate them with a host of commonly used examples, including those in continuous piecewise affine regression and in deep learning (where the activation functions are piecewise affine). We describe a nonmonotone majorization-minimization (MM) algorithm for solving the unified nonconvex, nondifferentiable optimization problem which is formulated as a specially structured composite dc program of the pointwise max type, and present convergence results to a directional stationary solution. An efficient semismooth Newton method is proposed to solve the dual of the MM subproblems. Numerical results are presented to demonstrate the effectiveness of the proposed algorithm and the superiority of continuous piecewise affine regression over the standard linear model.

Key words. continuous piecewise affine regression, nonconvex optimization, nondifferentiable objective, ReLu activation function, semismooth Newton method

AMS subject classifications. 62J02, 90C26, 49J52

DOI. 10.1137/18M117337X

1. Introduction. Many modern statistical estimation problems are defined by three major components: a statistical model that postulates the dependence of an output quantity on the input features; a loss function measuring the error between the observed output and model predicted output; and a regularizer that controls the overfitting and/or variable selection in the model. The overall estimation problem is to determine certain unknown parameters in the statistical model. In practical computation, samples on the underlying covariates (inputs) and output are available and a data-based empirical objective combined with a regularizer is formulated as a minimization problem that constitutes the workhorse of the estimation process. This paper addresses the computational solution of the latter optimization problem which is challenged by its nonconvexity and nondifferentiability. These features immediately raise the question about the stationarity, let alone minimizing, properties of the computed solution by an iterative algorithm. Our goal is to compute a directional stationary solution [47], which is the sharpest kind among all stationary solutions.

Traditionally, a statistical model is usually described by a linear combination of the input features, resulting in a linear regression or classification model. Nevertheless,

*Received by the editors March 2, 2018; accepted for publication (in revised form) September 19, 2018; published electronically December 13, 2018.

<http://www.siam.org/journals/siop/28-4/M117337.html>

Funding: The work of the first two authors was based on research partially supported by the U.S. National Science Foundation grant IIS–1632971. The work of the third author was based on research partially supported by the U.S. National Science Foundation grant DMS–1712822.

[†]Department of Industrial and Systems Engineering, University of Southern California, Los Angeles, CA 90089 (yingcui@usc.edu, jongship@usc.edu).

[‡]Department of Statistics, Columbia University, New York, NY 10027 (bodhi@stat.columbia.edu).

in recent years, research into the use of piecewise affine, convex [19, 20, 39], and nonconvex [24] estimation functions has been on the rise. There are other piecewise affine/quadratic functions that arise from different applications, such as those using piecewise affine functions to approximate the signum function and multilayer neural networks in deep learning [59, Chapter 4]. These nontraditional, nondifferentiable estimation functions have provided an important impetus for our work.

Some classical convex loss measures are based on the absolute value or squared errors between the model predicted outputs and the observed outputs, leading to an ℓ_1 or ℓ_2 loss function for regression. Other loss functions include a logarithmic loss for logistic regression, and an exponential loss for boosting. In the area of support vector machines, a hinge loss is commonly used for binary classification. In recent years, a truncated hinge loss function has been proposed to reduce the effect of outliers [61]. A distinguished property of the latter loss function is that it is nonconvex and nondifferentiable. All these loss measures are composed with the statistical model involving the parameters to be estimated, leading to an overall composite loss function to be minimized, which is nonconvex and nondifferentiable, due either to a piecewise statistical model or a loss function that lacks convexity and/or differentiability. In the area of support vector machines and other applications, a norm function of the model variables is often used as the model regularizer to avoid overfitting. In recent years, sparsity constraints to avoid model overfitting [26] are an important consideration in high-dimensional statistical learning. Starting with the pioneering work of Fan and Li [22], various sparsity functions have been proposed as surrogates of the discontinuous counting step function; it was shown in [1, 34] that all these surrogate sparsity functions can be expressed as the difference of two convex functions of a particular type.

In this paper, we introduce a composite difference-of-convex-piecewise program of the pointwise max type as a unification of a host of statistical models, loss functions, and model regularizers. A highlight of this formulation is the emphasis on the separate roles of each component function that is key to the development of a nonmonotone majorization-minimization (MM) algorithm for solving the overall nonconvex, nondifferentiable optimization problem. While global optima of nonconvex problems cannot be computed, stationary solutions of various kinds are computable by iterative algorithms with guaranteed convergence. What is essential is that focus should be placed on computing sharp stationary solutions which distinguish themselves as being the ones that must satisfy all other relaxed definitions of stationarity. With guaranteed subsequential convergence, the developed MM algorithm aims to compute such a stationary point that is based on the elementary notion of directional derivatives of the objective function. By using the Kurdyka–Łojaziewicz (KL) theory of semianalytic functions [2, 3, 9], we also show the sequential convergence of the iterates produced by the algorithm. In order to handle the nondifferentiable pointwise max terms, auxiliary variables are employed to express these pointwise maximum functions as constraints. With the regularization of the added variables in defining the subproblems, the resulting MM algorithm no longer guarantees the monotone property of the original objective; this leads to our terminology of “nonmonotone MM algorithm.” Due to the modifications which are introduced to facilitate the fast and efficient solution of the subproblems in the iterative steps, a separate proof of convergence is needed for the algorithm. To ensure the rapid convergence and high accuracy required of the solution of the subproblems, we employ the semismooth Newton (SN) method [49, 46] for semismoothly differentiable (SC^1) functions applied to the dual of such subproblems. We present numerical results to demonstrate the effectiveness of the

overall algorithm and the superiority of a nontraditional continuous piecewise affine regression model over the traditional linear model using least-squares regression. In addition to this specific problem, the nonmonotone MM algorithm combined with the SN method (abbreviated as MM+SN) developed in this paper has wide applications to many related problems that include the multilayer neural network problems with piecewise activation functions in deep learning. Due to page limitation, we will report the details of the other applications in separate papers.

The major contributions of this paper are fourfold: (a) we identify a unified composite difference-max program for a host of modern statistical estimation problems and illustrate the formulation with a variety of special instances; (b) we develop an MM-based algorithm for computing, for the first time, a directional stationary solution of such a nonconvex, nonsmooth program; (c) we present an SN method for solving the MM subproblems and illustrate its effectiveness when applied to a piecewise affine regression problem; and (d) we report computational results on the latter application that demonstrate the superiority of this nontraditional statistical estimation approach over traditional linear regression.

2. A unified composite difference-convex-piecewise program. In the definition of a statistical estimation problem, we place particular emphasis and care on the mathematical properties of its defining functions; identifying these properties is needed to effectively deal with the joint feature of nonconvexity and nondifferentiability and to facilitate the design of the MM+SN algorithm and its analysis for solving the resulting optimization problem. Specifically, the optimization problem consists of the following four major components, each playing a separate role in the overall statistical estimation process. As such, distinguishing them offers flexibility to accommodate a variety of model specifications.

- A parametric statistical model: $y = \psi(x; \theta) + \varepsilon$, where y is the response (output) variable (taken to be a scalar) given the input $x \in \mathbb{R}^d$, ε is the unobserved random error assumed to have (conditional) mean zero, and $\theta \in \mathbb{R}^m$ is the model parameter to be estimated, for some positive integers d (number of input features) and m (number of parameters). In general, these two dimensions, d and m , may be different (see, e.g., (3)); however, they may be the same as in the case of linear regression.

- An output-dependent loss function $\varphi(y, \bullet)$, whose composition with the statistical model $\psi(x; \theta)$ leads to the composite function $(\theta; y, x) \mapsto \varphi(y, \psi(x; \theta))$ that provides a measure of the deviation between the model predicted output $\psi(x; \theta)$ and the observed response y ; we stress the importance of this composition as properties of φ and ψ may be very different and need to be understood separately for best results.

- A regularizing function $P(\theta)$ that is used either for the purpose of strongly convexifying the objective function or in high-dimensional problems to control the selection of the model parameter θ ; the latter control is particularly important to avoid overfitting when the number of features d is relatively large.

- An admissible set $\Theta \subseteq \mathbb{R}^m$ that further restricts the parameter space in the presence of domain knowledge on the parameter θ ; often this set Θ is the whole space \mathbb{R}^m , leading to an unconstrained selection of the parameter.

Together, the tuple (φ, ψ, Θ) defines the following constrained minimization problem that is central to many statistical estimation problems:

$$(1) \quad \underset{\theta \in \Theta}{\text{minimize}} \mathbb{E} [\varphi(y, \psi(x; \theta))],$$

where the expectation \mathbb{E} is taken over the joint distribution of the input x and the output y . The objective expresses the average loss of accuracy in the model output

versus the realized output taking into account the uncertainties in the random input-output pair (x, y) . This perspective of statistical estimation is in line with stochastic programming and is a departure from classical statistical estimation where there is always a postulate of a “ground truth” of the estimator. In this paper, we investigate the numerical solution of problem (1) via the application of the sample average approximation scheme. Specifically, taking N independent and identically distributed samples $\{(x^s, y_s)\}_{s=1}^N \subseteq \mathbb{R}^{d+1}$ and adding the regularizer $P(\theta)$ weighted by the scalar $\gamma_N \geq 0$, we consider the empirical optimization problem

$$(2) \quad \underset{\theta \in \Theta}{\text{minimize}} \quad \frac{1}{N} \sum_{s=1}^N [\varphi(y_s, \psi(x^s; \theta))] + \gamma_N P(\theta).$$

While this formulation (2) is rather classical, the detailed treatment of the distinguished properties of the functions (φ, ψ, P) constitutes the novelty and intellectual merits of the present paper. In what follows, we present some details of these functions and describe how they arise. An important aim of the discussion is to motivate a focused formulation of problem (2) that allows a unified treatment of these statistical estimation problems as a composite difference-max program by a common algorithmic procedure with guaranteed convergence properties; see formulation (6) and its setting (assumptions C1–C3) as well as developments in the subsequent sections.

Statistical models. We are interested in piecewise smooth models, including the following:

- A continuous piecewise affine parametric model; this is a (generally nonconvex) piecewise affine function expressed as the difference of two convex piecewise affine functions [51] each with its own parameters: for two positive integers k_a and k_b ,

$$(3) \quad \psi(x; \theta) \triangleq \max_{1 \leq i \leq k_a} \{ (a^i)^T x + \alpha_i \} - \max_{1 \leq i \leq k_b} \{ (b^i)^T x + \beta_i \}$$

with parameter $\theta \triangleq \{(a^i, \alpha_i)_{i=1}^{k_a}, (b^i, \beta_i)_{i=1}^{k_b}\} \in \mathbb{R}^{(k_a+k_b)(d+1)}$. The convex case $k_b = 1$ was studied in [19, 20, 39]; study of the nonconvex case (where $k_b > 1$) can be found in the unpublished manuscript [24] and also in [4] with a max-min representation of the piecewise affine function. Model (3) includes the 1-layer neural network by the rectified linear unit (ReLU) activation function [43, 23] that has the simple form $\psi(x; \theta) = \max(a^T x + \alpha, 0)$, where the parameter $\theta = (a, \alpha) \in \mathbb{R}^{d+1}$.

- A multilayer neural network with the ReLU activation function: for simplicity, we present a 2-layer model in deep learning [59],

$$(4) \quad \psi(x; \theta) \triangleq \max(b^T \max(Ax + a, 0) + \beta, 0),$$

where the parameter θ consists of the vectors b and a in \mathbb{R}^k , the matrix $A \in \mathbb{R}^{k \times d}$, and scalar $\beta \in \mathbb{R}$. The two occurrences of the max ReLU functions indicate the action of 2 hidden layers, where the max of $Ax + a$ and 0 is taken along each coordinate. Omitting the proof, we can show that this function can be written in the following difference-of-max form:

$$(5) \quad \psi(x; \theta) = \max_{1 \leq j \leq k_1} \psi_{1,j}(x; \theta) - \max_{1 \leq j \leq k_2} \psi_{2,j}(x; \theta)$$

for some positive integers k_1 and k_2 and convex functions $\psi_{i,j}(x; \bullet)$ that are all once but not twice continuously differentiable piecewise linear-quadratic. (A continuous function ψ is *piecewise linear-quadratic* [50, Definition 10.20] on a domain if the domain can be represented as the union of finitely many polyhedral sets on each of which

ψ is a quadratic function; see [50, 52, 53] for properties of piecewise functions of this type). The above difference-of-convex (dc) pointwise maximum representation (in short, difference-max representation) of $\psi(x; \theta)$ has several advantages over the original definition (4). Extending the piecewise affine model (3) to a piecewise quadratic model, the formulation (5) motivates a unified form of the function ψ in these two models. More importantly, one can take advantage of the piecewise linear-quadratic components for the computation of directional stationary solutions (to be defined later). Furthermore, while it is clear from (4) that $\psi(x; \bullet)$ is a piecewise quadratic function, definition (4) does not immediately reveal the linear-quadratic feature of this function.

We make two important remarks. One, it is possible to extend the above 2-layer treatment in two major directions: (i) to multilayers and (ii) to general piecewise affine activation functions of which the ReLu function is a special case. Due to their significance, the full treatment of these deep-learning problems is presented in a separate paper. Two, in the algorithmic development, we will take each $\psi_{i,j}(x; \theta)$ in (5) as a once continuously differentiable convex function.

Loss functions. These include both differentiable and piecewise affine (thus non-differentiable) functions. Of particular interest are the following convex loss functions: the classical quadratic function for least-squares regression; the Huber loss function in robust estimation; the quantile function [28] that includes the absolute deviation loss function; the one-parameter exponential family via log-likelihood maximization [7, 11]. We are also interested in the nonconvex yet piecewise affine truncated hinge loss function for binary and multiclass classification [61] with the form

$$\varphi(y, t) \triangleq \max(1 - (t - y), 0) - \max(\delta - (t - y), 0) \quad \text{for some parameter } \delta \leq 0.$$

With the proof omitted, we can show that a composition of this function with the difference-max function (5) is also a function of the same kind. Therefore, the composite function $\varphi(y, \psi(x; \theta))$ is itself a difference-max function, thus amenable to treatment by the methodology developed in the later sections.

Regularizers. Traditionally, strongly convex regularizers are very prominent; in sparse linear regression, convex and dc regularizers are becoming popular as they are employed as surrogate sparsity functions. Most of the latter regularizers are not differentiable. Frequently used convex regularizers include the ℓ_2 -norm used in support vector machines [14], the weighted ℓ_1 in sparsity representation [26], and the total variation norm in image processing [54]. Examples of dc surrogate sparsity functions include the smoothly clipped absolute deviation (SCAD) function [22], the minimax concave penalty function [60], and the truncated transformed ℓ_1 [58, 18], and the truncated logarithmic penalty [12, 18]. It can be shown that all of the above mentioned dc regularizers can be written in the unified form

$$P(\theta) = \sum_{i=1}^m c_i |\theta_i| - \sum_{i=1}^m p_i(\theta_i)$$

with each p_i being a univariate differentiable convex function; see, e.g., [1].

Putting together the above families of functions $\varphi(y, \bullet)$, $\psi(x; \bullet)$, and $P(\bullet)$, we arrive at the following detailed formulation of problem (2),

$$(6) \quad \underset{\theta \in \Theta}{\text{minimize}} \quad f_N(\theta) \triangleq \frac{1}{N} \sum_{s=1}^N \varphi_s \circ \psi_s(\theta) + \gamma_N [P_1(\theta) - P_2(\theta)]$$

with the composite difference-max structure summarized below:

C1: each φ_s is a univariate convex function;
 C2: each ψ_s is a dc function of the pointwise max type (5); specifically, for some nonnegative integers $k_{s;1}$ and $k_{s;2}$,

$$(7) \quad \psi_s(\theta) \triangleq \max_{1 \leq i \leq k_{s;1}} \psi_{s;1,i}(\theta) - \max_{1 \leq i \leq k_{s;2}} \psi_{s;2,i}(\theta)$$

with each $\psi_{s;1,i}(\theta)$ and $\psi_{s;2,i}(\theta)$ being convex and continuously differentiable;
 C3: both P_1 and P_2 are convex with P_2 being additionally the pointwise maximum of finitely many convex differentiable functions.

In this setting, the overall objective function f_N is not convex. Since the composition of a convex function with a dc function is of the dc type, by [25, Theorem II, p. 708], it follows that f_N is a dc function; thus the dc algorithm in dc programming [33, 32] is in principle applicable to compute a critical point of f_N on the feasible set Θ . In general, for a dc program, $\min_{z \in Z} [g(z) - h(z)]$, where Z is a closed convex set in \mathbb{R}^n and g and h are convex functions, a vector $\bar{z} \in Z$ is a critical point if $\partial h(\bar{z}) \cap [\partial g(\bar{z}) + \mathcal{N}(\bar{z}; Z)] \neq \emptyset$, where the notation $\partial \varphi(\bar{z})$ denotes the subdifferential of a convex function φ at a given vector \bar{z} and $\mathcal{N}(\bar{z}; Z)$ denotes the normal cone of the set Z at $\bar{z} \in Z$. Nevertheless, this criticality concept has several major drawbacks. One, it depends on the dc representation of the objective function f_N . Two, although f_N is known to be dc, a dc decomposition is not readily available when each of the composite functions $\varphi_s \circ \psi_s$ is derived from the statistical functions given above. Third, as will be shown in the next section, criticality is a very weak property and can have no bearing at all with a desirable minimizing property. For these reasons, our research goal is to seek an alternative definition of stationarity that is the “sharpest” of its kind and which is applicable to problem (5) without demanding a dc decomposition of the composite functions $\varphi_s \circ \psi_s$. One shall see from the subsequent discussion that the particular (difference of) pointwise maximum structures of ψ_s and P_2 are critical to achieve this goal.

3. A detour: Subdifferentials and stationarity. As a first step in studying the composite nonconvex, nondifferentiable problem (6), we take a detour to present some results from variational analysis; see [50] for details. These materials will prepare us for the introduction of two key notions of stationary solutions that provide the objects of convergence of the iterative algorithms for solving problem (6), to be presented in the following sections. Let $\Phi : \Omega \rightarrow \mathbb{R}^m$ be a locally Lipschitz continuous vector-valued function defined on an open set $\Omega \subseteq \mathbb{R}^n$. It follows that Φ is F(réchet)-differentiable almost everywhere on Ω (cf. [50, Theorem 9.60]). Denote by $\mathcal{D}_\Phi \subseteq \Omega$ the set of points where Φ is F-differentiable and by $J\Phi(x) \in \mathbb{R}^{m \times n}$ the Jacobian of Φ at $x \in \mathcal{D}_\Phi$. Let $\bar{x} \in \Omega$ be given. The B(ouligand)-subdifferential of Φ at \bar{x} is denoted by

$$\partial_B \Phi(\bar{x}) \triangleq \left\{ V \in \mathbb{R}^{m \times n} \mid \exists \{x^k\} \subseteq \mathcal{D}_\Phi \text{ with } \lim_{k \rightarrow \infty} x^k = \bar{x} \text{ and } \lim_{k \rightarrow \infty} J\Phi(x^k) = V \right\}.$$

The Clarke subdifferential (also called the Clarke generalized Jacobian) of Φ at \bar{x} is defined as $\partial_C \Phi(\bar{x}) \triangleq \text{conv}(\partial_B \Phi(\bar{x}))$, where “conv” stands for the convex hull of a given set. For a real-valued function $\phi : \Omega \rightarrow \mathbb{R}$, the Clarke subdifferential of ϕ at \bar{x} can also be characterized by

$$\partial_C \phi(\bar{x}) = \left\{ v \in \mathbb{R}^n \mid \limsup_{x \rightarrow \bar{x}, t \downarrow 0} \frac{\phi(x + tw) - \phi(x) - t v^T w}{t} \geq 0 \quad \forall w \in \mathbb{R}^n \right\}.$$

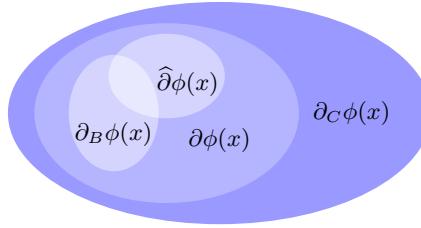


FIG. 1. Relationship between the subdifferentials.

The regular subdifferential of ϕ at \bar{x} is defined as

$$\widehat{\partial}\phi(\bar{x}) \triangleq \left\{ v \in \mathbb{R}^n \mid \liminf_{\bar{x} \neq x \rightarrow \bar{x}} \frac{\phi(x) - \phi(\bar{x}) - v^T(x - \bar{x})}{\|x - \bar{x}\|} \geq 0 \right\}.$$

The limiting subdifferential of ϕ at \bar{x} is defined as

$$\partial\phi(\bar{x}) \triangleq \left\{ v \in \mathbb{R}^n \mid \exists \{x^k\} \rightarrow \bar{x} \text{ and } \{v^k\} \rightarrow v \text{ such that } v^k \in \widehat{\partial}\phi(x^k) \text{ for all } k \right\}.$$

When ϕ is convex, its Clarke subdifferential, regular subdifferential, and limiting subdifferential coincide with the set of all subgradients of ϕ . In general, the relationship between the above four definitions is demonstrated in Figure 1. The detailed explanations are given in the proposition below.

PROPOSITION 1. *Let Ω be an open set in \mathbb{R}^n . For any locally Lipschitz continuous function $\phi : \Omega \rightarrow \mathbb{R}$ and any $\bar{x} \in \Omega$, it holds that*

- (i) $(\partial_B\phi(\bar{x}) \cup \widehat{\partial}\phi(\bar{x})) \subseteq \partial\phi(\bar{x}) \subseteq \partial_C\phi(\bar{x})$;
- (ii) $\text{conv}(\widehat{\partial}\phi(\bar{x})) \subseteq \text{conv}(\partial_B\phi(\bar{x})) = \text{conv}(\partial\phi(\bar{x})) = \partial_C\phi(\bar{x})$.

Moreover, counterexamples exist for the omitted inclusions.

Proof. It is known from [50, Theorem 8.6] that $\widehat{\partial}\phi(\bar{x}) \subseteq \partial\phi(\bar{x})$ and from [42, Theorem 3.57] that $\text{conv}(\partial\phi(\bar{x})) = \partial_C\phi(\bar{x})$. Thus, it remains to show that $\partial_B\phi(\bar{x}) \subseteq \partial\phi(\bar{x})$. Let $\bar{v} \in \partial_B\phi(\bar{x})$. Then there exists $\{x^k\} \subseteq \Omega$ converging to \bar{x} such that ϕ is F-differentiable at x^k and $\nabla\phi(x^k) \rightarrow \bar{v}$. Since $\widehat{\partial}\phi(x^k) = \{\nabla\phi(x^k)\}$ (cf. [50, Exercise 8.8]), we have $\bar{v} \in \partial\phi(\bar{x})$ by the definition of $\partial\phi$.

We provide two examples to show the possible invalidity of the omitted inclusions. For $\phi(x) = |x|$, we have $\partial\phi(0) = \widehat{\partial}\phi(0) = [-1, 1] \not\subseteq \partial_B\phi(0) = \{-1, 1\}$. For $\phi(x) = -|x|$, we have $\partial_C\phi(0) = [-1, 1] \not\subseteq \partial_B\phi(0) = \partial\phi(0) = \{-1, 1\} \not\subseteq \widehat{\partial}\phi(0) = \emptyset$ and $\text{conv}(\partial_B\phi(0)) = [-1, 1] \not\subseteq \text{conv}(\widehat{\partial}\phi(0)) = \emptyset$. \square

Let X be a closed convex set in Ω . It is known that a necessary condition for $\bar{x} \in X$ to be a local minimum of ϕ is $0 \in \widehat{\partial}(\phi(\bar{x}) + \delta_X(\bar{x}))$ [50, Theorem 10.1], where $\delta_X(x) \triangleq \begin{cases} 0 & \text{if } x \in X \\ +\infty & \text{otherwise} \end{cases}$ is the indicator function of X at x . If ϕ is locally Lipschitz continuous near \bar{x} and directionally differentiable (dd) at \bar{x} , the latter condition is equivalent to the d(irectional)-stationarity of \bar{x} , i.e.,

$$\phi'(\bar{x}; v) \triangleq \lim_{\delta \downarrow 0} \frac{\phi(\bar{x} + \delta v) - \phi(\bar{x})}{\delta} \geq 0 \quad \forall v \in X - \bar{x}.$$

Using the limiting subdifferential and Clarke subdifferential, respectively, we say that a point $\bar{x} \in X$ is a l(imiting)-stationary point of ϕ on X if $0 \in \partial(\phi(\bar{x}) + \delta_X(\bar{x}))$, and

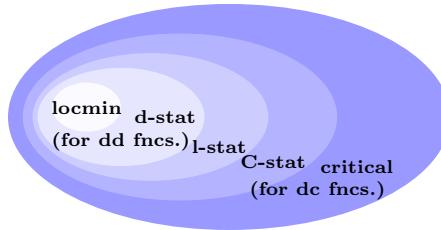


FIG. 2. Relationship between the stationary points.

a C(larke)-stationary point if $0 \in \partial_C \phi(\bar{x}) + \mathcal{N}(\bar{x}; X)$, which implies

$$\phi^\circ(\bar{x}; v) \triangleq \limsup_{x \rightarrow \bar{x}, \delta \downarrow 0} \frac{\phi(x + \delta v) - \phi(x)}{\delta} \geq 0 \quad \forall v \in X - \bar{x}.$$

Based on Proposition 1, l-stationarity implies C-stationarity; and if ϕ is locally Lipschitz near \bar{x} and dd at this point, d-stationarity implies l-stationarity. The following examples show that the reverse implications do not hold without additional assumptions.

Example 1. Consider the univariate function $\phi(x) = \max(-|x|, x - 1)$ for $x \in \mathbb{R}$. Since $\partial_C \phi(0) = [-1, 1]$ and $\partial \phi(0) = \{-1, 1\}$, it holds that $x = 0$ is a C-stationary point of ϕ , but fails to be an l-stationary point. The unique l-stationary point of ϕ on \mathbb{R} is $x = \frac{1}{2}$. For the univariate function $\phi(x) = \max(-x - 1, \min(-x, 0))$, since $\partial \phi(0) = \{-1, 0\}$ and $\widehat{\partial} \phi(0) = \emptyset$, it holds that $x = 0$ is an l-stationary point of ϕ , but not a d-stationary point. The unique d-stationary point of ϕ on \mathbb{R} is $x = -1$. \square

If $\phi = \phi_1 - \phi_2$ is a difference of two convex functions ϕ_1 and ϕ_2 , it holds that $\partial_C \phi(x) \subseteq \partial_C \phi_1(x) - \partial_C \phi_2(x)$ for any $x \in \Omega$ (cf. [13, Corollary 2]). A point $\bar{x} \in \Omega$ is said to be a critical point of ϕ on X if

$$\partial_C \phi_2(\bar{x}) \cap [\partial_C \phi_1(\bar{x}) + \mathcal{N}(\bar{x}; X)] = \partial \phi_2(\bar{x}) \cap [\partial \phi_1(\bar{x}) + \mathcal{N}(\bar{x}; X)] \neq \emptyset.$$

Different from the above mentioned concepts of stationary points, a critical point depends on the dc decomposition. Since there are infinitely many dc decompositions of a given dc function, it is likely that a critical point provides no information on the local minima of that function. This can be seen from the following example.

Example 2. Consider the univariate function $\phi(x) = \max\{x, -x - 4\}$ on \mathbb{R} , for which $x = -2$ is the unique C-stationary point (and the global minimizer). For the dc decomposition $\phi = \phi_1 - \phi_2$ with $\phi_1(x) = \max\{2x, 0, -2x - 4\}$ and $\phi_2(x) = |x|$, we have $\partial \phi_1(0) \cap \partial \phi_2(0) = [0, 1]$, implying that $x = 0$ is a critical point of ϕ on \mathbb{R} . \square

The relationship between the various kinds of stationary points and a critical point (for dc problems) is summarized in Figure 2. As a caution to the reader, we note that $0 \in \partial_B \phi(\bar{x})$ is not a necessary condition for \bar{x} being a local minima of a locally Lipschitz continuous function ϕ ; this can be seen from $\phi(x) = |x|$ at $x = 0$.

We close this section by mentioning three properties of d-stationarity that further highlight the fundamental importance of this stationarity concept. The first property asserts that a d-stationary point must be “locally ε -first-order minimizing”; the second and third property are applicable to the least-squares piecewise affine regression problem. We recall that a function ψ is B(ouligand)-differentiable at a point \bar{x} [21, Definition 3.1.2] if it is both locally Lipschitz continuous near \bar{x} and directionally

differentiable at \bar{x} . Proof of the proposition below is omitted as it is not difficult; for related results, see [15, 16].

PROPOSITION 2. *Let X be a closed convex set contained in the open set $\Omega \subseteq \mathbb{R}^n$. The following three statements hold:*

- (i) *Let $\psi : \Omega \rightarrow \mathbb{R}$ be B -differentiable at a d -stationary point $\bar{x} \in X$. It holds that for every $\varepsilon > 0$, there exists an open neighborhood \mathcal{N} of \bar{x} such that $\psi(x) \geq \psi(\bar{x}) - \varepsilon \|x - \bar{x}\|$ for all $x \in X \cap \mathcal{N}$.*
- (ii) *Let $\psi : \Omega \rightarrow \mathbb{R}^m$ be piecewise affine and $\Phi : \mathbb{R}^m \rightarrow \mathbb{R}$ be convex. It holds that every d -stationary point of the composite function $\Phi \circ \psi$ on the set X is a local minimizer.*
- (iii) *If X is polyhedral and ψ is piecewise linear-quadratic on X , then the set of values of ψ on the set of d -stationary points of ψ on X is finite.*

4. The nonmonotone MM algorithm. We aim to compute a d -stationary solution of problem (6). The approach we propose is to apply the MM algorithm with suitable modifications. Originally described in [44, section 9.3(d)], the idea of the basic MM algorithm is to solve a sequence of convex minimization subproblems by creating surrogate functions that majorize the original objective function. It unifies various optimization methods, including the projected/proximal gradient method [6] and the dc algorithm [33, 32]. It is a generalization of the expectation-maximization algorithm, as known in the statistics community, for finding maximum likelihood estimators of parameters in statistical models [17, 56]. See [27, 31] for comprehensive discussions of the MM algorithm (without modifications) and [37, 38, 9] for recent developments.

Given a locally Lipschitz continuous function $\phi : \Omega \rightarrow \mathbb{R}$ defined on an open set $\Omega \subseteq \mathbb{R}^n$ containing the closed convex set X and a point $y \in X$, the continuous function $\hat{\phi}(\bullet, y)$ is said to be a majorizing function of ϕ if (i) $\hat{\phi}(\bullet, y)$ is convex on X ; (ii) $\phi(x) \leq \hat{\phi}(x, y)$ for all $x \in X$; and (iii) $\hat{\phi}(y) = \hat{\phi}(y, y)$. The key to a successful application of the MM algorithm to the optimization problem (6) hinges on a readily available convex majorization of the objective function f_N at a given iterate θ^ν . In turn, it suffices to derive such a majorizing function for each summand $\varphi_s \circ \psi_s$; for the dc regularizer $P(\theta) = P_1(\theta) - P_2(\theta)$ a convex majorizing function is easily obtained as $\hat{P}(\theta, \theta^\nu) = P_1(\theta) - [P_2(\theta^\nu) + (a^{2:\nu})^T(\theta - \theta^\nu)]$, where $a^{2:\nu} \in \partial P_2(\theta^\nu)$. The following fact combined with the difference-max structure (7) of each function ψ_s easily yields a convex majorant of each composite function $\varphi_s \circ \psi_s$.

LEMMA 3. *A univariate convex function f can be written as the sum of a convex nondecreasing function f^\uparrow and a convex nonincreasing function f^\downarrow .*

Proof. Note that if there exists $t_0 \in \mathbb{R}$ such that $f'(t_0; \pm 1) \geq 0$, then f achieves its minimum value at t_0 . By the convexity of f , one can easily check that f is nonincreasing on $(-\infty, t_0]$ and nondecreasing on $[t_0, \infty)$, leading to a choice of

$$f^\uparrow(t) \triangleq \begin{cases} f(t) - f(t_0) & \text{if } t \geq t_0 \\ 0 & \text{if } t < t_0 \end{cases} \quad \text{and} \quad f^\downarrow(t) \triangleq \begin{cases} f(t_0) & \text{if } t \geq t_0 \\ f(t) & \text{if } t < t_0. \end{cases}$$

To see the convexity of f^\uparrow , it suffices to observe that for any t_1 and t_2 in \mathbb{R} satisfying $t_1 < t_0 < t_2$ and for any $\lambda \in (0, 1)$,

$$\begin{aligned} f^\uparrow((1 - \lambda)t_1 + \lambda t_2) &\leq f^\uparrow((1 - \lambda)t_0 + \lambda t_2) = f((1 - \lambda)t_0 + \lambda t_2) - f(t_0) \\ &\leq (1 - \lambda)f(t_0) + \lambda f(t_2) - f(t_0) = (1 - \lambda)f^\uparrow(t_1) + \lambda f^\uparrow(t_2). \end{aligned}$$

By a similar argument, we can show the convexity of f^\downarrow .

If such t_0 does not exist, then for all $t \in \mathbb{R}$, either $f'(t; 1) \geq 0 \geq f'(t; -1)$ or $f'(t; 1) \leq 0 \leq f'(t; -1)$. The former situation implies that f is nondecreasing on \mathbb{R} ; hence we may take $f^\uparrow = f$ and $f^\downarrow = 0$ in this case; while the latter situation implies that f is nonincreasing on \mathbb{R} ; hence we may choose $f^\downarrow = f$ and $f^\uparrow = 0$. \square

Since a convex nondecreasing (nonincreasing) function composed with a convex (concave) function is convex, the following corollary is easy to obtain. The proof is omitted for brevity.

COROLLARY 4. *Let φ be a univariate convex function with the decomposition $\varphi = \varphi^\uparrow + \varphi^\downarrow$ as in Lemma 3. For any $\theta^0 \in \Theta$, if $\widehat{\psi}(\bullet, \theta^0)$ is a concave minorizing function and $\check{\psi}(\bullet, \theta^0)$ a convex majoring function of ψ on Θ , i.e.,*

$$\widehat{\psi}(\theta, \theta^0) \leq \psi(\theta) \leq \check{\psi}(\theta, \theta^0) \quad \forall \theta \in \Theta,$$

then $\varphi^\uparrow \circ \check{\psi}(\bullet, \theta^0) + \varphi^\downarrow \circ \widehat{\psi}(\bullet, \theta^0)$ is a convex majorant of the composite $\varphi \circ \psi$ on Θ .

To simplify the notation, we present the MM algorithm and its convergence property for solving the following single-summand formulation of (6):

$$(8) \quad \underset{\theta \in \Theta}{\text{minimize}} \quad \Psi(\theta) \triangleq \varphi \circ \psi(\theta),$$

where φ is a univariate convex function and $\psi(\theta) = g(\theta) - h(\theta)$, where

$$g(\theta) \triangleq \max_{1 \leq i \leq k_1} \psi_{1,i}(\theta) \quad \text{and} \quad h(\theta) \triangleq \max_{1 \leq i \leq k_2} \psi_{2,i}(\theta)$$

with each $\psi_{1,i}$ and $\psi_{2,i}$ being a differentiable convex function. The treatment is clearly extendable to $\frac{1}{N} \sum_{s=1}^N \varphi_s \circ \psi_s(\theta) + \gamma_N P(\theta)$ with each pair (φ_s, ψ_s) as above and P satisfying assumption C3. We omit the details of the extended treatment but will employ it in the computational experiments; see section 6.

Denote the index sets of maximizing functions in g and h at $\theta \in \Theta$ as

$$\mathcal{A}_1(\theta) \triangleq \underset{1 \leq i \leq k_1}{\text{argmax}} \{ \psi_{1,i}(\theta) \} \quad \text{and} \quad \mathcal{A}_2(\theta) \triangleq \underset{1 \leq i \leq k_2}{\text{argmax}} \{ \psi_{2,i}(\theta) \}.$$

For any $(\theta, \bar{\theta}) \in \Theta \times \Theta$ and any $(i_1, i_2) \in \mathcal{A}_1(\bar{\theta}) \times \mathcal{A}_2(\bar{\theta})$, we have

$$\underbrace{[g(\bar{\theta}) + \nabla \psi_{1,i_1}(\bar{\theta})^T (\theta - \bar{\theta})]}_{\text{linearization of } \psi_{1,i_1} \text{ at } \bar{\theta}} - h(\theta) \leq \psi(\theta) \leq g(\theta) - \underbrace{[h(\bar{\theta}) + \nabla \psi_{2,i_2}(\bar{\theta})^T (\theta - \bar{\theta})]}_{\text{linearization of } \psi_{2,i_2} \text{ at } \bar{\theta}},$$

which, by Corollary 4, leads to the following convex majorant of Ψ in (8):

$$\begin{aligned} \mathcal{M}\Psi_{(i_1, i_2)}(\theta, \bar{\theta}) &\triangleq \varphi^\uparrow (g(\theta) - [h(\bar{\theta}) + \nabla \psi_{2,i_2}(\bar{\theta})^T (\theta - \bar{\theta})]) \\ &\quad + \varphi^\downarrow ([g(\bar{\theta}) + \nabla \psi_{1,i_1}(\bar{\theta})^T (\theta - \bar{\theta})] - h(\theta)). \end{aligned}$$

The above function is nonsmooth because of the “max” in the functions g and h . Notice that we essentially leave φ unchanged instead of approximating it. In this way, we may presumably obtain a tighter approximation of the original composite objective function $\varphi \circ \psi$.

4.1. From pointwise max to constraints. Before describing the MM-based algorithm that can be shown to converge to a d-stationary solution of problem (8), we first present the following lemma that characterizes a d-stationary solution of this problem as a solution of finitely many convex programs.

LEMMA 5. *The point $\bar{\theta} \in \Theta$ is a d-stationary point of (8) if and only if for all $(i_1, i_2) \in \mathcal{A}_1(\bar{\theta}) \times \mathcal{A}_2(\bar{\theta})$, $\bar{\theta} \in \operatorname{argmin}_{\theta \in \Theta} \mathcal{M}\Psi_{(i_1, i_2)}(\theta, \bar{\theta})$.*

Proof. *Only if:* If $\bar{\theta} \in \Theta$ is a d-stationary point of (6), then for all $\theta \in \Theta$,

$$\begin{aligned} 0 \leq \Psi'(\bar{\theta}; \theta - \bar{\theta}) &= \varphi'(\psi(\bar{\theta}); \psi'(\bar{\theta}; \theta - \bar{\theta})) \\ &= (\varphi^\uparrow)'(\psi(\bar{\theta}); \psi'(\bar{\theta}; \theta - \bar{\theta})) + (\varphi^\downarrow)'(\psi(\bar{\theta}); \psi'(\bar{\theta}; \theta - \bar{\theta})). \end{aligned}$$

We have $\psi(\bar{\theta}) = g(\bar{\theta}) - h(\bar{\theta})$ and

$$\begin{aligned} \psi'(\bar{\theta}; \theta - \bar{\theta}) &= \underbrace{\max_{i \in \mathcal{A}_1(\bar{\theta})} \nabla \psi_{1,i}(\bar{\theta})^T (\theta - \bar{\theta})}_{= g'(\bar{\theta}; \theta - \bar{\theta})} - \underbrace{\max_{i \in \mathcal{A}_2(\bar{\theta})} \nabla \psi_{2,i}(\bar{\theta})^T (\theta - \bar{\theta})}_{= h'(\bar{\theta}; \theta - \bar{\theta})}. \end{aligned}$$

Since φ^\uparrow is a univariate nondecreasing convex function, it follows that the directional derivative $(\varphi^\uparrow)'(t; \bullet)$ is a nondecreasing function of the direction for every $t \in \mathbb{R}$. Hence, for every $i_2 \in \mathcal{A}_2(\bar{\theta})$,

$$\begin{aligned} (\varphi^\uparrow)'(\psi(\bar{\theta}); \psi'(\bar{\theta}; \theta - \bar{\theta})) &= (\varphi^\uparrow)'(g(\bar{\theta}) - h(\bar{\theta}); g'(\bar{\theta}; \theta - \bar{\theta}) - h'(\bar{\theta}; \theta - \bar{\theta})) \\ &\leq (\varphi^\uparrow)'(g(\bar{\theta}) - h(\bar{\theta}); g'(\bar{\theta}; \theta - \bar{\theta}) - \nabla \psi_{2,i_2}(\bar{\theta})^T (\theta - \bar{\theta})). \end{aligned}$$

Similarly, we also have, for every $i_1 \in \mathcal{A}_1(\bar{\theta})$,

$$(\varphi^\downarrow)'(\psi(\bar{\theta}); \psi'(\bar{\theta}; \theta - \bar{\theta})) \leq (\varphi^\downarrow)'(g(\bar{\theta}) - h(\bar{\theta}); \nabla \psi_{1,i_1}(\bar{\theta})^T (\theta - \bar{\theta}) - h'(\bar{\theta}; \theta - \bar{\theta})),$$

because $(\varphi^\downarrow)'(t; \bullet)$ is a nonincreasing function for every $t \in \mathbb{R}$. It therefore follows that $\mathcal{M}\Psi_{(i_1, i_2)}(\bullet, \bar{\theta})'(\bar{\theta}; \theta - \bar{\theta}) \geq 0$ for all $\theta \in \Theta$. Since $\mathcal{M}\Psi_{(i_1, i_2)}(\bullet, \bar{\theta})$ is a convex program, it follows that $\bar{\theta}$ is a minimizer of this function over Θ .

If: Conversely, suppose that for every $(i_1, i_2) \in \mathcal{A}_1(\bar{\theta}) \times \mathcal{A}_2(\bar{\theta})$, it holds that $\bar{\theta} \in \operatorname{argmin}_{\theta \in \Theta} \mathcal{M}\Psi_{(i_1, i_2)}(\theta, \bar{\theta})$. Let $\theta \in \Theta$ be arbitrary. Pick

$$i_1 \in \operatorname{argmax}_{i \in \mathcal{A}_1(\bar{\theta})} \nabla \psi_{1,i}(\bar{\theta})^T (\theta - \bar{\theta})$$

and $i_2 \in \operatorname{argmax}_{i \in \mathcal{A}_2(\bar{\theta})} \nabla \psi_{2,i}(\bar{\theta})^T (\theta - \bar{\theta})$. We then have

$$\Psi'(\bar{\theta}; \theta - \bar{\theta}) = \mathcal{M}\Psi_{(i_1, i_2)}(\bullet, \bar{\theta})'(\bar{\theta}; \theta - \bar{\theta}) \geq 0.$$

Since $\theta \in \Theta$ is arbitrary, it follows that $\bar{\theta}$ is a d-stationary point of Ψ on Θ . \square

Lemma 5 indicates that there is a “for all index pairs” condition in the requirement of d-stationarity. Though $\mathcal{M}\Psi_{(i_1, i_2)}(\bullet, \theta^\nu)$ is a convex majorant of Ψ for any pair $(i_1, i_2) \in \mathcal{A}_1(\theta^\nu) \times \mathcal{A}_2(\theta^\nu)$, by arbitrarily picking a single pair of indices (i_1, i_2) at each iteration of the MM algorithm, one may not obtain an algorithm that converges to a d-stationary point. This nonconvergence has been observed in [47] when the dc algorithm (a special MM algorithm) is applied to solve dc programs (a special case of problem (8) with φ being the identity function); see specifically Example 4 in this reference. In order for the MM algorithm to converge to a d-stationary point, we

employ the ε -technique to expand the argmax index sets of the functions g and h . Specifically, given $\varepsilon > 0$ and $\theta \in \Theta$, let

$$\begin{cases} \mathcal{A}_{1;\varepsilon}(\theta) \triangleq \underset{1 \leq i \leq k_1}{\varepsilon\text{-argmax}} g(\theta) = \{1 \leq i \leq k_1 \mid \psi_{1,i}(\theta) \geq g(\theta) - \varepsilon\}, \\ \mathcal{A}_{2;\varepsilon}(\theta) \triangleq \underset{1 \leq i \leq k_2}{\varepsilon\text{-argmax}} h(\theta) = \{1 \leq i \leq k_2 \mid \psi_{2,i}(\theta) \geq h(\theta) - \varepsilon\}. \end{cases}$$

To facilitate the solution of the MM subproblem, we further introduce auxiliary variables r and s in \mathbb{R} to write the max functions in g and h as constraints, thus smoothing out the arguments in the functions φ^\uparrow and φ^\downarrow , and then regularize the added variables in the objective function. Specifically, for any $\bar{\theta} \in \Theta$ and any $(i_1, i_2) \in \mathcal{A}_1(\bar{\theta}) \times \mathcal{A}_2(\bar{\theta})$, define the convex set

$$\begin{aligned} & \mathcal{Z}_{(i_1, i_2)}(\bar{\theta}) \\ & \triangleq \left\{ \begin{array}{l} z \triangleq (\theta, r, s) \mid \psi_{1,j}(\theta) - h(\bar{\theta}) - \nabla \psi_{2,i_2}(\bar{\theta})^T (\theta - \bar{\theta}) \leq r, 1 \leq j \leq k_1 \\ \in \Theta \times \mathbb{R} \times \mathbb{R} \mid g(\bar{\theta}) + \nabla \psi_{1,i_1}(\bar{\theta})^T (\theta - \bar{\theta}) - \psi_{2,j}(\theta) \geq s, 1 \leq j \leq k_2 \end{array} \right\}, \end{aligned}$$

which must contain the triple $(\bar{\theta}, r, s)$ for any pair of scalars (r, s) satisfying $r \geq \psi(\bar{\theta}) \geq s$. Since, φ^\uparrow is nondecreasing and φ^\downarrow is nonincreasing, we have

$$(9) \quad \varphi^\uparrow(r) + \varphi^\downarrow(s) \geq \mathcal{M}\Psi_{(i_1, i_2)}(\theta, \bar{\theta}) \geq \Psi(\theta) \quad \forall (\theta, r, s) \in \mathcal{Z}_{(i_1, i_2)}(\bar{\theta}).$$

Moreover, equalities hold throughout (9) if $r = s = \psi(\theta)$.

With the above preparations, the MM-based algorithm is given below.

The Nonmonotone MM algorithm for solving (8).

Given are positive scalars c and ε and an initial point $\theta^0 \in \Theta$. Let $z^0 \triangleq (\theta^0, r_0, s_0)$, where $r_0 \triangleq \psi(\theta^0) \triangleq s_0$. Set $\nu = 0$.

Step 1. For every pair $(i_1^\nu, i_2^\nu) \in \mathcal{A}_{1;\varepsilon}(\theta^\nu) \times \mathcal{A}_{2;\varepsilon}(\theta^\nu)$, compute

$$(10) \quad z^{\nu+\frac{1}{2};i_1^\nu, i_2^\nu} \triangleq \operatorname{argmin}_{z \in \mathcal{Z}_{(i_1^\nu, i_2^\nu)}(\theta^\nu)} \left\{ \widehat{\Psi}_c(z, z^\nu) \triangleq \varphi^\uparrow(r) + \varphi^\downarrow(s) + \frac{c}{2} \|z - z^\nu\|^2 \right\}.$$

Step 2. Set $z^{\nu+1} \triangleq z^{\nu+\frac{1}{2};\widehat{i}_1^\nu, \widehat{i}_2^\nu}$, where $(\widehat{i}_1^\nu, \widehat{i}_2^\nu)$ is a minimizing index in

$$\operatorname{argmin} \left\{ \widehat{\Psi}_c(z^{\nu+\frac{1}{2};i_1^\nu, i_2^\nu}, z^\nu) \mid (i_1^\nu, i_2^\nu) \in \mathcal{A}_{1;\varepsilon}(\theta^\nu) \times \mathcal{A}_{2;\varepsilon}(\theta^\nu) \right\}.$$

Step 3. If $z^{\nu+1}$ satisfies a prescribed stopping rule, terminate; otherwise, return to Step 1 with ν replaced by $\nu + 1$. \square

Because of the regularization of the variables r and s , subproblem (10) is a slight modification of the MM subproblem which would be

$$(11) \quad \underset{z \in \mathcal{Z}_{(i_1^\nu, i_2^\nu)}(\theta^\nu)}{\operatorname{minimize}} \quad \widetilde{\Psi}_c(z, \theta^\nu) \triangleq \varphi^\uparrow(r) + \varphi^\downarrow(s) + \frac{c}{2} \|\theta - \theta^\nu\|^2.$$

The modification is needed when we discuss the solution of (10) that is based on a smoothness property of its Lagrangian dual. Unlike (11), a minimizing triple (θ, r, s) of (10) does not necessarily satisfy $r = g(\theta) - [h(\theta^\nu) + \nabla \psi_{2,i_2^\nu}(\theta^\nu)^T (\theta - \theta^\nu)]$ or $s =$

$[g(\theta^\nu) + \nabla\psi_{1,i_1^\nu}(\theta^\nu)^T(\theta - \theta^\nu)] - h(\theta)$ because of the regularization of these variables. In addition, the feasible set $\mathcal{Z}_{(i_1^\nu, i_2^\nu)}(\theta^\nu)$ changes with the iterate θ^ν . Due to all these anomalies, care is needed in the proof of convergence of the algorithm. In particular, the sequence $\{\Psi(\theta^\nu)\}$ of objective values of the original function Ψ to be minimized is not shown to be decreasing; instead it is the substituted sequence $\{\varphi^\uparrow(r_{\nu+1}) + \varphi^\downarrow(s_{\nu+1})\}$ that is decreasing. The term “nonmonotone” is employed to highlight this nonstandard feature of the algorithm. We first give a necessary and sufficient condition for a vector $\bar{\theta}$ to be an optimal solution of the problem $\text{minimize}_{\theta \in \Theta} \mathcal{M}\Psi_{(i_1, i_2)}(\theta, \bar{\theta})$, this being a key requirement in the d-stationarity of $\bar{\theta}$.

LEMMA 6. *A vector $\bar{\theta} \in \text{argmin}_{\theta \in \Theta} \mathcal{M}\Psi_{(i_1, i_2)}(\theta, \bar{\theta})$, where $(i_1, i_2) \in \mathcal{A}_1(\bar{\theta}) \times \mathcal{A}_2(\bar{\theta})$, if and only if there exist a scalar $c > 0$ and a pair of scalars (\bar{r}, \bar{s}) such that the triple $\bar{z} \triangleq (\bar{\theta}, \bar{r}, \bar{s}) \in \text{argmin}_{z \in \mathcal{Z}_{(i_1, i_2)}(\bar{\theta})} \hat{\Psi}_c(z, \bar{z})$.*

Proof. Only if. Suppose $\bar{\theta} \in \text{argmin}_{\theta \in \Theta} \mathcal{M}\Psi_{(i_1, i_2)}(\theta, \bar{\theta})$. Let $c > 0$ be arbitrary. It then follows that $\bar{\theta} \in \text{argmin}_{\theta \in \Theta} [\mathcal{M}\Psi_{(i_1, i_2)}(\theta, \bar{\theta}) + \frac{c}{2} \|\theta - \bar{\theta}\|^2]$. Define $\bar{r} \triangleq \psi(\bar{\theta}) \triangleq \bar{s}$. Then $\bar{z} \triangleq (\bar{\theta}, \bar{r}, \bar{s}) \in \mathcal{Z}_{(i_1, i_2)}(\bar{\theta})$. Letting $z \triangleq (\theta, r, s) \in \mathcal{Z}_{(i_1, i_2)}(\bar{\theta})$ be arbitrary, we have

$$\begin{aligned} \hat{\Psi}_c(z, \bar{z}) &= \varphi^\uparrow(r) + \varphi^\downarrow(s) + \frac{c}{2} \|z - \bar{z}\|^2 \\ &\geq \mathcal{M}\Psi_{(i_1, i_2)}(\theta, \bar{\theta}) + \frac{c}{2} \|\theta - \bar{\theta}\|^2 \quad \text{by (9)} \\ &\geq \mathcal{M}\Psi_{(i_1, i_2)}(\bar{\theta}, \bar{\theta}) \quad \text{by assumption on } \bar{\theta} \\ &= \varphi^\uparrow(\bar{r}) + \varphi^\downarrow(\bar{s}) = \hat{\Psi}_c(\bar{z}, \bar{z}) \quad \text{since } \bar{r} = \psi(\bar{\theta}) = \bar{s}. \end{aligned}$$

If. Conversely, suppose $\bar{z} \triangleq (\bar{\theta}, \bar{r}, \bar{s}) \in \text{argmin}_{z \in \mathcal{Z}_{(i_1, i_2)}(\bar{\theta})} \hat{\Psi}_c(z, \bar{z})$ for some scalar $c > 0$. It follows that $\bar{z} \triangleq (\bar{\theta}, \bar{r}, \bar{s}) \in \text{argmin}_{z \in \mathcal{Z}_{(i_1, i_2)}(\bar{\theta})} [\varphi^\uparrow(r) + \varphi^\downarrow(s)]$, which implies in particular that $\bar{r} \geq \psi(\bar{\theta}) \geq \bar{s}$. Hence $(\theta, \psi(\bar{\theta}), \psi(\bar{\theta})) \in \text{argmin}_{z \in \mathcal{Z}_{(i_1, i_2)}(\bar{\theta})} [\varphi^\uparrow(r) + \varphi^\downarrow(s)]$ and $\varphi^\uparrow(\bar{r}) + \varphi^\downarrow(\bar{s}) = \Psi(\bar{\theta}) = \mathcal{M}\Psi_{(i_1, i_2)}(\bar{\theta}, \bar{\theta})$. Let $\theta \in \Theta$ be arbitrary. Define

$$r \triangleq g(\theta) - [h(\bar{\theta}) + \nabla\psi_{2, i_2}(\bar{\theta})^T(\theta - \bar{\theta})], \quad s \triangleq [g(\bar{\theta}) + \nabla\psi_{1, i_1}(\bar{\theta})^T(\theta - \bar{\theta})] - h(\theta).$$

Then $(\theta, r, s) \in \mathcal{Z}_{(i_1, i_2)}(\bar{\theta})$ and $\varphi^\uparrow(r) + \varphi^\downarrow(s) = \mathcal{M}\Psi_{(i_1, i_2)}(\theta, \bar{\theta})$. It then follows that $\mathcal{M}\Psi_{(i_1, i_2)}(\theta, \bar{\theta}) \geq \mathcal{M}\Psi_{(i_1, i_2)}(\bar{\theta}, \bar{\theta})$, establishing that $\theta \in \text{argmin}_{\theta \in \Theta} \mathcal{M}\Psi_{(i_1, i_2)}(\theta, \bar{\theta})$. \square

Now we present the main result of this section on the subsequential convergence of $\{\theta^\nu\}$ to a d-stationary point of (8), which generalizes the result in [47, Proposition 6] for the dc algorithm to solve dc programs.

PROPOSITION 7. *Suppose that Ψ is bounded below on the closed convex set Θ . Let $\{z^\nu = (\theta^\nu, r_\nu, s_\nu)\}$ be a sequence generated by the MM algorithm. The following four statements hold:*

(a) *For any $(i_1^\nu, i_2^\nu) \in \mathcal{A}_{1; \varepsilon}(\theta^\nu) \times \mathcal{A}_{2; \varepsilon}(\theta^\nu)$ and $z \in \mathcal{Z}_{(i_1^\nu, i_2^\nu)}(\theta^\nu)$, it holds that*

$$\varphi^\uparrow(r_{\nu+1}) + \varphi^\downarrow(s_{\nu+1}) + \frac{c}{2} \|z^{\nu+1} - z^\nu\|^2 \leq \varphi^\uparrow(r) + \varphi^\downarrow(s) + \frac{c}{2} \|z - z^\nu\|^2.$$

(b) $\lim_{\nu \rightarrow \infty} \|z^{\nu+1} - z^\nu\| = 0$.

(c) *For any accumulation point $(\theta^\infty, r_\infty, s_\infty)$ of $\{z^\nu\}$, if it exists, θ^∞ is a d-stationary point of (8).*

(d) If the set $\{\theta \in \Theta \mid \Psi(\theta) \leq \Psi(\theta^0)\}$ is bounded, then the sequence $\{\theta^\nu\}$ is bounded. If in addition, the set $\{t \mid \varphi(t) \leq \Psi(\theta^0)\}$ is bounded, then the sequence $\{(r_\nu, s_\nu)\}$ is also bounded.

Proof. (a) For any $(i_1^\nu, i_2^\nu) \in \mathcal{A}_{1;\varepsilon}(\theta^\nu) \times \mathcal{A}_{2;\varepsilon}(\theta^\nu)$ and $z \in \mathcal{Z}_{(i_1^\nu, i_2^\nu)}(\theta^\nu)$, we derive

$$\begin{aligned} & \varphi^\uparrow(r) + \varphi^\downarrow(s) + \frac{c}{2} \|z - z^\nu\|^2 \\ & \geq \varphi^\uparrow(r_{\nu+\frac{1}{2};i_1^\nu, i_2^\nu}) + \varphi^\downarrow(s_{\nu+\frac{1}{2};i_1^\nu, i_2^\nu}) + \frac{c}{2} \|z^{\nu+\frac{1}{2};i_1^\nu, i_2^\nu} - z^\nu\|^2 \\ & \geq \varphi^\uparrow(r_{\nu+1}) + \varphi^\downarrow(s_{\nu+1}) + \frac{c}{2} \|z^{\nu+1} - z^\nu\|^2, \end{aligned}$$

where the two inequalities are due to the definitions of $z^{\nu+\frac{1}{2};i_1^\nu, i_2^\nu}$ and $z^{\nu+1}$.

(b) By the strong convexity of $\widehat{\Psi}_c(\bullet, z^\nu)$, it follows that $z^{\nu+\frac{1}{2};i_1^\nu, i_2^\nu}$ is the unique minimizer of (10). We claim that $z^\nu \in \mathcal{Z}_{(i_1, i_2)}(\theta^\nu)$ for all ν and all $(i_1, i_2) \in \{1, \dots, k_1\} \times \{1, \dots, k_2\}$. Indeed, this is clearly true for $\nu = 0$ by the choice of (r_0, s_0) . For any $\nu \geq 0$, since $z^{\nu+1} \in \mathcal{Z}_{(i_1^\nu, i_2^\nu)}(\theta^\nu)$ for some $(i_1^\nu, i_2^\nu) \in \mathcal{A}_{1;\varepsilon}(\theta^\nu) \times \mathcal{A}_{2;\varepsilon}(\theta^\nu)$, we have

$$\begin{aligned} r_{\nu+1} & \geq g(\theta^{\nu+1}) - [h(\theta^\nu) + \nabla \psi_{2,i_2^\nu}(\theta^\nu)^T (\theta^{\nu+1} - \theta^\nu)] \\ & \geq g(\theta^{\nu+1}) - \psi_{2,i_2^\nu}(\theta^{\nu+1}) \geq g(\theta^{\nu+1}) - h(\theta^{\nu+1}) = \psi(\theta^{\nu+1}). \end{aligned}$$

Similarly, we can deduce $s_{\nu+1} \leq \psi(\theta^{\nu+1})$. Consequently, the claim holds. It then follows from (a) that

$$(12) \quad \varphi^\uparrow(r_{\nu+1}) + \varphi^\downarrow(s_{\nu+1}) + \frac{c}{2} \|z^{\nu+1} - z^\nu\|^2 \leq \varphi^\uparrow(r_\nu) + \varphi^\downarrow(s_\nu) \quad \forall \nu.$$

Thus the sequence $\{\varphi^\uparrow(r_\nu) + \varphi^\downarrow(s_\nu)\}$ is nonincreasing. Since $\varphi^\uparrow(r_\nu) + \varphi^\downarrow(s_\nu) \geq \Psi(\theta^\nu)$ and Ψ is bounded below on Θ , it follows that the sequence $\{\varphi^\uparrow(r_\nu) + \varphi^\downarrow(s_\nu)\}$ converges. Hence by (12), the sequence $\{\|z^{\nu+1} - z^\nu\|\}$ converges to zero.

(c) Let $z^\infty \triangleq (\theta^\infty, r_\infty, s_\infty)$ be the limit of a convergent subsequence $\{z^\nu\}_{\nu \in \kappa}$. We must have $\theta^\infty \in \Theta$. To prove that θ^∞ is a d-stationary point, it suffices to show, by Lemmas 5 and 6, that z^∞ belongs to $\operatorname{argmin}_{z \in \mathcal{Z}_{(i_1, i_2)}(\theta^\infty)} \widehat{\Psi}_c(z, z^\infty)$ for all $(i_1, i_2) \in \mathcal{A}_1(\theta^\infty) \times \mathcal{A}_2(\theta^\infty)$. Consider any $(i_1, i_2) \in \mathcal{A}_1(\theta^\infty) \times \mathcal{A}_2(\theta^\infty)$ and any $z \in \mathcal{Z}_{(i_1, i_2)}(\theta^\infty)$. Define

$$\begin{aligned} \widehat{r}_\nu & \triangleq r + |[h(\theta^\nu) + \nabla \psi_{2,i_2}(\theta^\nu)^T (\theta - \theta^\nu)] - [h(\theta^\infty) + \nabla \psi_{2,i_2}(\theta^\infty)^T (\theta - \theta^\infty)]|, \\ \widehat{s}_\nu & \triangleq s + |[g(\theta^\nu) + \nabla \psi_{1,i_1}(\theta^\nu)^T (\theta - \theta^\nu)] - [g(\theta^\infty) + \nabla \psi_{1,i_1}(\theta^\infty)^T (\theta - \theta^\infty)]|. \end{aligned}$$

We then have $\widehat{z}^\nu \triangleq (\theta, \widehat{r}_\nu, \widehat{s}_\nu) \in \mathcal{Z}_{(i_1, i_2)}(\theta^\nu)$ and $\lim_{\nu \in \kappa \rightarrow \infty} (\widehat{r}_\nu, \widehat{s}_\nu) = (r, s)$. One can easily show that for all $\nu \in \kappa$ sufficiently large, $\mathcal{A}_1(\theta^\infty) \times \mathcal{A}_2(\theta^\infty) \subseteq \mathcal{A}_{1;\varepsilon}(\theta^\nu) \times \mathcal{A}_{2;\varepsilon}(\theta^\nu)$, which in turn implies $(i_1, i_2) \in \mathcal{A}_1(\theta^\nu) \times \mathcal{A}_2(\theta^\nu)$ for all such ν . It therefore follows from (a) that

$$\varphi^\uparrow(r_{\nu+1}) + \varphi^\downarrow(s_{\nu+1}) + \frac{c}{2} \|z^{\nu+1} - z^\nu\|^2 \leq \varphi^\uparrow(\widehat{r}_\nu) + \varphi^\downarrow(\widehat{s}_\nu) + \frac{c}{2} \|\widehat{z}^\nu - z^\nu\|^2.$$

Passing to the limit $\nu \in \kappa \rightarrow \infty$, we obtain $z^\infty \in \operatorname{argmin}_{z \in \mathcal{Z}_{(i_1, i_2)}(\theta^\infty)} \widehat{\Psi}_c(z, z^\infty)$. This completes the proof of statement (c).

(d) Since $\Psi(\theta^\nu) \leq \varphi^\uparrow(r_\nu) + \varphi^\downarrow(s_\nu) \leq \varphi^\uparrow(r_0) + \varphi^\downarrow(s_0) = \Psi(\theta^0)$ for all ν , it follows that the sequence $\{\theta^\nu\}$ is bounded by assumption. Since $r_\nu \geq s_\nu$, it follows that

$$\varphi^\uparrow(r_\nu) + \varphi^\downarrow(s_\nu) \geq \max(\varphi(s_\nu), \varphi(r_\nu));$$

the boundedness of $\{(r_\nu, s_\nu)\}$ follows similarly. \square

One may further establish the convergence of the full sequence $\{z^\nu\}$ under an isolatedness assumption of an accumulation point as in the general theory of sequential convergence [21, Proposition 8.3.10] or by invoking the KL theory of semianalytic functions as in [2, 3, 9, 10]. Recall that a proper lower semicontinuous function $f : \mathbb{R}^n \rightarrow (-\infty, +\infty]$ is said to have the KL property at $\bar{x} \in \text{dom}(\partial f) \triangleq \{x \in \text{dom } f \mid \partial f(x) \neq \emptyset\}$ if there exist $\alpha > 0$, a neighborhood \mathcal{N} of \bar{x} , and a continuous concave function $\phi : [0, \alpha) \rightarrow \mathbb{R}_+$ such that

- $\phi(0) = 0$ and ϕ is continuously differentiable on $(0, \alpha)$ with $\phi' > 0$; and
- $\phi'(f(x) - f(\bar{x})) \text{dist}(0, \partial f(x)) \geq 1$ for any $x \in \mathcal{N}$ with $f(\bar{x}) < f(x) < f(\bar{x}) + \alpha$, where $\text{dist}(x, C) \triangleq \min_{y \in C} \|x - y\|$ is the distance from a point $x \in \mathbb{R}^n$ to a nonempty closed set $C \subseteq \mathbb{R}^n$. If the function ϕ in the above definition can be chosen as $\phi(s) = \gamma s^{1-\beta}$ for some scalars $\gamma > 0$ and $\beta \in [0, 1)$, we say that f has the KL property at \bar{x} with an exponent β . It is known that a proper closed semialgebraic function is a KL function [8].

To proceed, we introduce the following function of $z \triangleq (\theta, r, s)$:

$$\bar{\Psi}(z) \triangleq \begin{cases} \varphi^\uparrow(r) + \varphi^\downarrow(s) & \text{if } z \in \bar{Z} \triangleq \{(\theta, r, s) \in \Theta \times \mathbb{R} \times \mathbb{R} \mid s \leq \psi(\theta) \leq r\}, \\ +\infty & \text{otherwise} \end{cases}$$

and also define the set $Z^\infty \triangleq \Theta^\infty \times R_\infty \times S_\infty$ of all accumulation points of the sequence $\{z^\nu\}$ produced by the nonmonotone MM algorithm. This set is nonempty under condition (d) of Proposition 7; moreover, every one of its elements is a d-stationary solution of problem (8).

PROPOSITION 8. *Suppose that Ψ is bounded below on the closed convex set Θ . Then the whole sequence $\{z^\nu\}$ converges to an unique element of the set Z^∞ under either one of the following two conditions (a) or (b):*

- (a) Z^∞ contains an isolated element;
- (b) $\{z^\nu\}$ is bounded; the index sets $\mathcal{A}_{1;2\varepsilon}(\theta)$ and $\mathcal{A}_{2;2\varepsilon}(\theta)$ are singletons for all $\theta \in \Theta^\infty$; $\nabla\psi_{1,i_1}$ and $\nabla\psi_{2,i_2}$ are locally Lipschitz continuous near all $\theta \in \Theta^\infty$ for all $(i_1, i_2) \in \mathcal{A}_{1;2\varepsilon}(\theta) \times \mathcal{A}_{2;2\varepsilon}(\theta)$, respectively; the functions φ^\uparrow and φ^\downarrow are differentiable with Lipschitz continuous gradients near all $r \in R_\infty$ and $s \in S_\infty$, respectively; moreover, the function $\bar{\Psi}$ satisfies the KL property at all $z^\infty \in Z^\infty$.

If $\{z^k\}$ converges to $z^\infty \in Z^\infty$ under condition (b) and the function $\bar{\Psi}$ has the KL exponent $\beta \in [0, 1)$ at z^∞ , it holds that

- if $\beta = 0$, the sequence $\{z^\nu\}$ converges in a finite number of steps;
- if $\beta \in (0, \frac{1}{2}]$, the sequence $\{z^\nu\}$ converges R-linearly, i.e., there exist $\eta > 0$ and $r \in [0, 1)$ such that $\|z^\nu - z^\infty\| \leq \eta r^\nu$ for all ν sufficiently large;
- if $\beta \in (\frac{1}{2}, 1)$, the sequence $\{z^\nu\}$ converges R-sublinearly, specifically, there exists $\eta > 0$ such that $\|z^\nu - z^\infty\| \leq \eta \nu^{-\frac{1-\beta}{2\beta-1}}$ for all ν sufficiently large.

Proof. It follows from Proposition 7(b) that $\lim_{\nu \rightarrow \infty} \|z^{\nu+1} - z^\nu\| = 0$. Under assumption (a) here, the convergence of the whole sequence $\{z^\nu\}$ to the isolated element of Z^∞ follows from [21, Proposition 8.3.10]. To prove the sequential convergence of $\{z^\nu\}$ under the conditions in assumption (b), it suffices to show the following three properties of $\{z^\nu\}$ and then apply the convergence results in [3, Theorem 2.9] to the function $\bar{\Psi}$:

- (i) $\bar{\Psi}(z^{\nu+1}) \leq \bar{\Psi}(z^\nu) - \frac{c}{2} \|z^{\nu+1} - z^\nu\|^2$ for all $\nu \geq 1$;
- (ii) there exists a subsequence $\{z^\nu\}_{\nu \in \kappa}$ of $\{z^\nu\}$ such that $\lim_{\nu \in \kappa \rightarrow \infty} z^\nu = z^\infty$ and $\lim_{\nu \in \kappa \rightarrow \infty} \bar{\Psi}(z^\nu) = \bar{\Psi}(z^\infty)$;

(iii) there exists a scalar $\eta > 0$ such that for all ν sufficiently large, $a^{\nu+1} \in \partial \bar{\Psi}(z^{\nu+1})$ exists satisfying $\|a^{\nu+1}\| \leq \eta \|z^{\nu+1} - z^\nu\|$.

Properties (i) and (ii) readily follow from Proposition 7. To show (iii), we first note that Z^∞ is a nonempty, compact, and connected set [21, Proposition 8.3.9]. By assumption (b), we further derive the existence of $\bar{i}_j \in \{1, \dots, k_j\}$ such that $\mathcal{A}_{j;2\varepsilon}(\theta) = \{\bar{i}_j\}$ for all $\theta \in \Theta^\infty$, $j = 1, 2$. Since the scalar sequence $\{\text{dist}(\theta^\nu, \Theta^\infty)\}$ is bounded with a unique accumulation point 0, we have $\lim_{\nu \rightarrow \infty} \text{dist}(\theta^\nu, \Theta^\infty) = 0$. Hence, for some constant $\kappa > 0$ it holds that for all ν sufficiently large, some $\theta^{\nu;\infty} \in \Theta^\infty$ exists for which $\|\theta^\nu - \theta^{\nu;\infty}\| \leq \frac{\varepsilon}{2\kappa}$ and $|\psi_{1,i}(\theta^\nu) - \psi_{1,i}(\theta^{\nu;\infty})| \leq \kappa \|\theta^\nu - \theta^{\nu;\infty}\| \leq \frac{1}{2}\varepsilon$ for any $i = 1, \dots, k_1$. Then for any $i \in \mathcal{A}_{1;\varepsilon}(\theta^\nu)$,

$$\begin{aligned} \psi_{1,i}(\theta^{\nu;\infty}) &\geq \psi_{1,i}(\theta^\nu) - \frac{1}{2}\varepsilon \geq \max_{1 \leq i \leq k_1} \psi_{1,i}(\theta^\nu) - \frac{3}{2}\varepsilon \quad \text{since } i \in \mathcal{A}_{1;\varepsilon}(\theta^\nu) \\ &\geq \psi_{1,\bar{i}_1}(\theta^\nu) - \frac{3}{2}\varepsilon \geq \psi_{1,\bar{i}_1}(\theta^{\nu;\infty}) - 2\varepsilon \\ &= \max_{1 \leq i \leq k_1} \psi_{1,i}(\theta^{\nu;\infty}) - 2\varepsilon \quad \text{since } \mathcal{A}_{1;2\varepsilon}(\theta^{\nu;\infty}) = \{\bar{i}_1\} = \mathcal{A}_1(\theta^{\nu;\infty}), \end{aligned}$$

i.e., $i \in \mathcal{A}_{1;2\varepsilon}(\theta^{\nu;\infty}) = \{\bar{i}_1\}$. Therefore, for all ν sufficiently large, $\mathcal{A}_{1;\varepsilon}(\theta^\nu) = \{\bar{i}_1\}$, and, similarly, $\mathcal{A}_{2;\varepsilon}(\theta^\nu) = \{\bar{i}_2\}$. By assumption (b) and [50, Exercise 10.10], we further derive

$$(13) \quad \begin{cases} \nabla \psi(\theta^{\nu+1}) = \nabla \psi_{1,\bar{i}_1}(\theta^{\nu+1}) - \nabla \psi_{2,\bar{i}_2}(\theta^{\nu+1}), \\ \partial \bar{\Psi}(z^{\nu+1}) = \begin{pmatrix} 0 \\ \nabla \varphi^{\uparrow}(r_{\nu+1}) \\ \nabla \varphi^{\downarrow}(s_{\nu+1}) \end{pmatrix} + \mathcal{N}(z^{\nu+1}; \bar{Z}) \end{cases} \quad \forall \nu \text{ sufficiently large.}$$

Note that \bar{Z} may not be a convex set; we write $\mathcal{N}(z; \bar{Z})$ as the normal cone of \bar{Z} at $z \in \bar{Z}$ based on the definition in [50, Definition 6.3]. One can easily check that

$$\left\{ \begin{array}{l} \alpha_1 \in \mathcal{N}(\psi(\theta^{\nu+1}) - r_{\nu+1}; \mathbb{R}_-), \alpha_2 \in \mathcal{N}(-\psi(\theta^{\nu+1}) + s_{\nu+1}; \mathbb{R}_-), \\ 0 \in \begin{pmatrix} (\alpha_1 - \alpha_2) \nabla \psi(\theta^{\nu+1}) + \mathcal{N}(\theta^{\nu+1}; \Theta) \\ -\alpha_1 \\ \alpha_2 \end{pmatrix} \end{array} \right\} \implies \alpha_1 = \alpha_2 = 0,$$

where we denote $\mathbb{R}_- \triangleq \{t \in \mathbb{R} \mid t \leq 0\}$. It then follows from [50, Theorem 6.14] that

$$\mathcal{N}(z^{\nu+1}; \bar{Z}) = \left\{ \begin{pmatrix} (\alpha_1 - \alpha_2) \nabla \psi(\theta^{\nu+1}) + \mathcal{N}(\theta^{\nu+1}; \Theta) \\ -\alpha_1 \\ \alpha_2 \end{pmatrix} \middle| \begin{array}{l} \alpha_1 \geq 0, \alpha_2 \geq 0, \\ \alpha_1 = 0 \text{ if } \psi(\theta^{\nu+1}) < r_{\nu+1}, \\ \alpha_2 = 0 \text{ if } \psi(\theta^{\nu+1}) > s_{\nu+1} \end{array} \right\}.$$

For all ν sufficiently large, since $\mathcal{A}_{1;\varepsilon}(\theta^\nu) \times \mathcal{A}_{2;\varepsilon}(\theta^\nu) = \{(\bar{i}_1, \bar{i}_2)\}$, we deduce $z^{\nu+1} = z^{\nu+\frac{1}{2}, \bar{i}_1, \bar{i}_2} = \text{argmin}_{z \in \mathcal{Z}(\bar{i}_1, \bar{i}_2)(\theta^\nu)} \hat{\Psi}_c(z, z^\nu)$. To proceed, we denote

$$\begin{aligned} I_1(z^{\nu+1}) &\triangleq \{i : \psi_{1,i}(\theta^{\nu+1}) - \psi_{2,\bar{i}_2}(\theta^\nu) - \nabla \psi_{2,\bar{i}_2}(\theta^\nu)^T (\theta^{\nu+1} - \theta^\nu) = r_{\nu+1}\}, \\ I_2(z^{\nu+1}) &\triangleq \{i : \psi_{1,\bar{i}_1}(\theta^\nu) + \nabla \psi_{1,\bar{i}_1}(\theta^\nu)^T (\theta^{\nu+1} - \theta^\nu) - \psi_{2,i}(\theta^{\nu+1}) = s_{\nu+1}\}. \end{aligned}$$

It then follows from $z^{\nu+1} \in \mathcal{Z}_{(\bar{i}_1, \bar{i}_2)}(\theta^\nu)$ that one of the following four cases must hold for any ν sufficiently large: (1) $I_1(z^{\nu+1}) = I_2(z^{\nu+1}) = \emptyset$; (2) $I_1(z^{\nu+1}) = \{\bar{i}_1\}$ and $I_2(z^{\nu+1}) = \emptyset$; (3) $I_1(z^{\nu+1}) = \emptyset$ and $I_2(z^{\nu+1}) = \{\bar{i}_2\}$; (4) $I_1(z^{\nu+1}) = \{\bar{i}_1\}$ and $I_2(z^{\nu+1}) = \{\bar{i}_2\}$. For those $z^{\nu+1}$ satisfying case (1), we derive

$$\begin{aligned}\psi(\theta^{\nu+1}) &= \psi_{1, \bar{i}_1}(\theta^{\nu+1}) - \psi_{2, \bar{i}_2}(\theta^{\nu+1}) \\ &\leq \psi_{1, \bar{i}_1}(\theta^{\nu+1}) - \psi_{2, \bar{i}_2}(\theta^\nu) - \nabla \psi_{2, \bar{i}_2}(\theta^\nu)^T (\theta^{\nu+1} - \theta^\nu) < r_{\nu+1}.\end{aligned}$$

Similarly one has $\psi(\theta^{\nu+1}) > s_{\nu+1}$. Then $\mathcal{N}(z^{\nu+1}; \bar{Z}) = \begin{pmatrix} \mathcal{N}(\theta^{\nu+1}; \Theta) \\ 0_{\mathbb{R}^2} \end{pmatrix}$. Since all the inequality constraints in $z \in \mathcal{Z}_{(\bar{i}_1, \bar{i}_2)}(\theta^\nu)$ are inactive at $z = z^{\nu+1}$, by applying the optimality condition of the problem: $\min_{z \in \mathcal{Z}_{(\bar{i}_1, \bar{i}_2)}(\theta^\nu)} \widehat{\Psi}_c(z, z^\nu)$, we deduce

$$\begin{aligned}0 &\in \mathcal{N}(\theta^{\nu+1}; \Theta) + c(\theta^{\nu+1} - \theta^\nu), \quad \text{and} \\ \nabla \varphi^\uparrow(r_{\nu+1}) + c(r_{\nu+1} - r_\nu) &= 0, \quad \nabla \varphi^\downarrow(s_{\nu+1}) + c(s_{\nu+1} - s_\nu) = 0,\end{aligned}$$

which, together with (13), yield $-c(z^{\nu+1} - z^\nu) \in \partial \bar{\Psi}(z^{\nu+1})$. Hence, property (iii) holds with the constant $\eta_1 = c$. For those $z^{\nu+1}$ satisfying case (2), there exists a multiplier $\lambda_{\nu+1} \in \mathbb{R}$ corresponding to the active constraint $\psi_{1, \bar{i}_1}(\theta^{\nu+1}) - \psi_{2, \bar{i}_2}(\theta^\nu) - \nabla \psi_{2, \bar{i}_2}(\theta^\nu)^T (\theta^{\nu+1} - \theta^\nu) = r_{\nu+1}$ such that

$$0 \in \begin{pmatrix} \lambda_{\nu+1} [\nabla \psi_{1, \bar{i}_1}(\theta^{\nu+1}) - \nabla \psi_{2, \bar{i}_2}(\theta^\nu)] + \mathcal{N}(\theta^{\nu+1}; \Theta) \\ \nabla \varphi^\uparrow(r_{\nu+1}) - \lambda_{\nu+1} \\ \nabla \varphi^\downarrow(s_{\nu+1}) \end{pmatrix} + c(z^{\nu+1} - z^\nu).$$

Since φ^\uparrow is nondecreasing, we have $\nabla \varphi^\uparrow(r_{\nu+1}) = \lambda_{\nu+1} - c(r_{\nu+1} - r_\nu) \geq 0$, which, together with (13) and (14), yields

$$\begin{pmatrix} -b^{\nu+1} \\ c(r_{\nu+1} - r_\nu) \\ 0 \end{pmatrix} - c(z^{\nu+1} - z^\nu) \in \partial \bar{\Psi}(z^{\nu+1}), \quad \text{where}$$

$$\begin{aligned}b^{\nu+1} &\triangleq c(r_{\nu+1} - r_\nu) [\nabla \psi_{1, \bar{i}_1}(\theta^{\nu+1}) - \nabla \psi_{2, \bar{i}_2}(\theta^\nu)] \\ &\quad + \nabla \varphi^\uparrow(r_{\nu+1}) [\nabla \psi_{2, \bar{i}_2}(\theta^{\nu+1}) - \nabla \psi_{2, \bar{i}_2}(\theta^\nu)].\end{aligned}$$

In addition, $\{\nabla \psi_{1, \bar{i}_1}(\theta^{\nu+1}) - \nabla \psi_{2, \bar{i}_2}(\theta^\nu)\}$ and $\{\nabla \varphi^\uparrow(r_{\nu+1})\}$ are bounded for all ν sufficiently large under the assumptions in (b). Hence, property (iii) holds for some constant $\eta_2 > 0$. Similarly, those points $z^{\nu+1}$ satisfying cases (3) and (4) can be proved to have property (iii) for some positive constants η_3 and η_4 , respectively. Consequently, property (iii) holds for all ν sufficiently large with $\eta = \max\{\eta_1, \dots, \eta_4\}$. The desired global convergence of $\{z^\nu\}$ is thus established based on the above arguments. Finally, with the proven properties (i)–(iii), the stated convergence rate can be established by arguments similar to [2, Theorem 2]; see also [9, Proposition 4]. \square

The sequential convergence of the standard MM algorithm (without regularizing the slack variables) has been studied in [10] under a condition that is in the same spirit as property (iii) in the above proof. In order to verify this property, we essentially assume that φ^\uparrow , φ^\downarrow , ψ_1 , and ψ_2 are differentiable with locally Lipschitz continuous gradients near Z^∞ . These types of assumptions have also been made in [32, 55, 36] to prove the sequential convergence of the dc algorithm (with extrapolation) for solving the dc program $\min_{x} f(x) - g(x)$, where the function g is assumed to be differentiable with a locally Lipschitz gradient near all accumulation points. Unlike the dc algorithm that only employs a majorization of the function g , we majorize both ψ_1 and ψ_2 and do not rely on a dc decomposition of the overall objective function Ψ . Admittedly, this differentiability assumption is rather restrictive for a nondifferentiable problem. The difficulty in trying to relax this assumption is the evaluation of the subdifferential $\partial\bar{\Psi}(z^{\nu+1})$ of the extended-value function $\bar{\Psi}$ that involves the nonconvex set \bar{Z} defined by the nondifferentiable function ψ . It is our interest to pursue a thorough investigation of this issue that links the theory of error bounds to the sequential convergence of iterative methods for specially structured nondifferentiable optimization problems, without relaying the hard-to-evaluate subdifferential as needed here. Regrettably, such a study is beyond the scope of this paper.

4.2. Two variants of the MM algorithm. Though the MM algorithm discussed above converges to a d-stationary point, it potentially requires high computational cost per iteration because multiple subproblems need to be solved. To reduce the computational cost, we consider two variants of the algorithm. The first one requires solving only one subproblem associated with an arbitrary pair of indices $(i_1, i_2) \in \mathcal{A}_1(\theta^\nu) \times \mathcal{A}_2(\theta^\nu)$ in each iteration.

MM-1: A simplified version of the MM algorithm for solving (8).

Initialization. Given are a scalar $c > 0$, an initial point $\theta^0 \in \Theta$, and an arbitrary pair in indices $(i_1^0, i_2^0) \in \mathcal{A}_1(\bar{\theta}) \times \mathcal{A}_2(\bar{\theta})$. Let $r_0 \triangleq \psi(\theta^0) \triangleq s_0$; thus $z^0 \triangleq (\theta^0, r_0, s_0) \in \mathcal{Z}_{(i_1^0, i_2^0)}(\bar{\theta})$. Set $\nu = 0$.

Step 1. Let $z^{\nu+1} \triangleq \operatorname{argmin}_{z \in \mathcal{Z}_{(i_1^\nu, i_2^\nu)}(\theta^\nu)} \hat{\Psi}_c(z, z^\nu)$ for some $(i_1^\nu, i_2^\nu) \in \mathcal{A}_1(\theta^\nu) \times \mathcal{A}_2(\theta^\nu)$.

Step 2. If $z^{\nu+1}$ satisfies a prescribed stopping rule, terminate; otherwise, return to Step 1 with ν replaced by $\nu + 1$. \square

Unfortunately, the above MM algorithm is not guaranteed to converge to a d-stationary solution of (8). The reason is easy to understand; namely, there is a for all index pairs condition in the requirement of d-stationarity, whereas each iteration of the MM-1 algorithm solves only one subprogram corresponding to a single index pair. Relaxing the “for all” requirement in d-stationarity, we say that a point $\bar{\theta} \in \Theta$ is a *weak \mathcal{M} -stationary point* of (8) if there exists $(\bar{i}_1, \bar{i}_2) \in \mathcal{A}_1(\bar{\theta}) \times \mathcal{A}_2(\bar{\theta})$ such that $\bar{\theta} \in \operatorname{argmin}_{\theta \in \Theta} \mathcal{M}\Psi_{(\bar{i}_1, \bar{i}_2)}(\theta, \bar{\theta})$. The prefix \mathcal{M} is used to highlight that the concept employs the majorization of the function Ψ at $\bar{\theta}$ by the family of convex functions $\{\mathcal{M}\Psi_{(\bar{i}_1, \bar{i}_2)}(\bullet, \bar{\theta})\}$. In the case where φ is the identity function, a weak \mathcal{M} -stationary point of the function Ψ on Θ coincides with a “weak d-stationary point” as defined in [47, subsection 3.3], thus a weak \mathcal{M} -stationary point must be a critical point (in the sense of dc programming) of a uncomposed difference-max program.

The convergence of the MM-1 algorithm is stated in the proposition below, whose proof can be derived similarly to the proof of Proposition 7. We omit the details here.

PROPOSITION 9. *Suppose that Ψ is bounded below on the closed convex set Θ . For any accumulation point $(\theta^\infty, r_\infty, s_\infty)$ of the sequence $\{z^\nu = (\theta^\nu, r_\nu, s_\nu)\}$ generated by the MM-1 algorithm, if it exists, θ^∞ is a weak \mathcal{M} -stationary point of (8).*

It is worth mentioning that unlike a d-stationary point of (8) which is independent of the representation of Ψ by its components φ and ψ , the concept of a weak \mathcal{M} -stationary point not only depends on the decomposition of $\varphi = \varphi^\uparrow + \varphi^\downarrow$, but also depends on the “max-max” representation of ψ . To see the dependence on the former, consider univariate functions

$$\varphi(t) = \underbrace{2t}_{\varphi_1^\uparrow(t)} + \underbrace{-t}_{\varphi_1^\downarrow(t)} = \underbrace{t}_{\varphi_2^\uparrow(t)} + \underbrace{0}_{\varphi_2^\downarrow(t)}, \quad \psi(t) = \underbrace{\max(2t, 1.5t)}_{g(t)} - \underbrace{\max(t, 0.5t)}_{h(t)}.$$

One can easily check that $t = 0$ is an \mathcal{M} -stationary point of $\varphi \circ \psi$ under the decomposition $\varphi = \varphi_1^\uparrow + \varphi_1^\downarrow$ (with $(i_1, i_2) = (1, 1) \in \mathcal{A}_1(0) \times \mathcal{A}_2(0)$), but fails to be that under the decomposition $\varphi = \varphi_2^\uparrow + \varphi_2^\downarrow$. To see the dependence on the latter, we may consider the example $\varphi(t) = t$ and $\psi(t) = \max(2t, 0) - \max(t, -t) = t - 0$ for a scalar t , and verify $t = 0$ is a weak \mathcal{M} -stationary point of $\varphi \circ \psi$ under the former representation, but fails to be weakly \mathcal{M} -stationary under the latter representation. In addition, a weak \mathcal{M} -stationary point of (8) may not be its C-stationary point as indicated by $\varphi(t) = t$ and $\psi(t) = \max(2t, 0) - \max(t, -t)$ at $t = 0$, and vice versa as shown by $\varphi(t) = t$ and $\psi(t) = t - \max(0, 2t)$ at $t = 0$.

The second variant of the MM algorithm is a probabilistic version that requires solving only one single convex subprogram corresponding to a randomly selected pair $(i_1, i_2) \in \mathcal{A}_{1;\varepsilon}(\theta^\nu) \times \mathcal{A}_{2;\varepsilon}(\theta^\nu)$, each with positive probability of being selected.

MM-2: A randomized version of the MM algorithm for solving (8).

Initialization. Given are a scalar $c > 0$, an initial point $\theta^0 \in \Theta$, and an arbitrary pair in indices $(i_1^0, i_2^0) \in \mathcal{A}_1(\bar{\theta}) \times \mathcal{A}_2(\bar{\theta})$. Let $r_0 \triangleq \psi(\theta^0) \triangleq s_0$; thus $z^0 \triangleq (\theta^0, r_0, s_0) \in \mathcal{Z}_{(i_1^0, i_2^0)}(\bar{\theta})$. Set $\nu = 0$.

Step 1. Choose $(i_1^\nu, i_2^\nu) \in \mathcal{A}_{1;\varepsilon}(\theta^\nu) \times \mathcal{A}_{2;\varepsilon}(\theta^\nu)$ randomly and independently from the previous iterations so that

$$p^{\nu, i_1, i_2} \triangleq \text{Prob}\{ (i_1, i_2) \text{ is chosen } \mid \text{given } z^\nu \} \geq p_{\min} > 0.$$

Let $z^{\nu+\frac{1}{2}; i_1^\nu, i_2^\nu} \triangleq \text{argmin}_{z \in \mathcal{Z}_{(i_1^\nu, i_2^\nu)}(\theta^\nu)} \widehat{\Psi}_c(z, z^\nu)$.

Step 2. Set $z^{\nu+1} \triangleq \begin{cases} z^{\nu+\frac{1}{2}; i_1^\nu, i_2^\nu} & \text{if } \widehat{\Psi}_c(z^{\nu+\frac{1}{2}; i_1^\nu, i_2^\nu}, z^\nu) < \varphi^\uparrow(r_\nu) + \varphi^\downarrow(s_\nu), \\ z^\nu & \text{otherwise.} \end{cases}$

Step 3. If $z^{\nu+1}$ satisfies a prescribed stopping rule, terminate; otherwise, return to Step 1 with μ replaced by $\mu + 1$. \square

The following proposition extends the convergence result of [47, Proposition 7] for dc programming. We omit the proof again since it can be derived with little difficulty by combining the arguments of Proposition 7 and [47, Proposition 7].

PROPOSITION 10. Suppose that Ψ is bounded below on the closed convex set Θ . Let $\{z^\nu = (\theta^\nu, r_\nu, s_\nu)\}$ be a sequence generated by the MM-2 algorithm. For any accumulation point $(\theta^\infty, r_\infty, s_\infty)$ of this sequence, if it exists, θ^∞ is a d -stationary point of (8) with probability one.

5. An SN method for the subproblems. The convex subproblems in Step 1 (see (10)) are the workhorse of the MM algorithms. In the present section, we discuss an SN method for solving them. There are two important motivations for us to adopt this Newton-type method instead of various first-order methods that are quite popular in recent years. One, although the computational cost for each step of the SN method may be more than that of first-order methods, this cost is compensated for by many fewer iterations, especially if the SN method is warm started by the solution of the preceding subproblem. The overall computational time of the SN method is thus potentially less than that of first-order methods. The recent papers [48, 62, 41, 57, 35] provide computational evidence to support the practical advantages of the SN method. Two, the MM algorithm itself is a first-order method, whose objective value can easily be stagnant after several iterations; see, e.g., [31, Chapter 1]. First-order methods may perform fairly well if they are directly employed to solve stand-alone problems to low or moderate accuracy. But unlike the SN method, which is able to compute very accurate solutions of the subproblems reasonably fast, first-order methods may not provide sufficiently accurate solutions when embedded in an iterative process such as the MM algorithm as the methods of choice for solving the subproblems.

Recall that a locally Lipschitz continuous function $\Phi : \Omega \rightarrow \mathbb{R}^m$ defined on the open set Ω is said to be *semismooth* at $\bar{x} \in \Omega$ if Φ is directionally differentiable at \bar{x} and for any $V \in \partial_C \Phi(x)$ with x sufficiently near \bar{x} , $\Phi(x) - \Phi(\bar{x}) - V(x - \bar{x}) = o(\|x - \bar{x}\|)$; Φ is said to be *strongly semismooth* at \bar{x} if $o(\|x - \bar{x}\|)$ is strengthened to $O(\|x - \bar{x}\|^2)$ [40, 49]. The function Φ is said to be a (strongly) semismooth function on Ω if it is (strongly) semismooth everywhere in Ω . The class of semismooth functions is broad, including piecewise semismooth functions, and thus piecewise smooth functions [21, Proposition 7.4.6]. It is also known that every piecewise affine map from \mathbb{R}^n into \mathbb{R}^m is strongly semismooth [21, Proposition 7.4.7]. A real-valued function $\phi : \Omega \rightarrow \mathbb{R}$ is said to be SC^1 at $\bar{x} \in \Omega$ [21, subsections 7.4.1 and 8.3.3] if it is differentiable near \bar{x} and its gradient is a semismooth function at \bar{x} .

With pioneering work by Kojima and Shindo [29], Kummer [30], Pang [45], Qi and Sun [49], among others, the SN method is a generalization of Newton's method for solving the semismooth equation $\Phi(x) = 0$, which has also been extended to a convex constrained SC^1 minimization problem in [46]; this is the method that we propose to apply to solve the dual problem of (10).

Letting $\mathbf{1}$ denote the vector of all ones of appropriate dimensions, the subproblem (10) can be written as, for the given pair (i_1^ν, i_2^ν) ,

$$(15) \quad \begin{aligned} & \underset{\theta \in \Theta, r, s \in \mathbb{R}}{\text{minimize}} && \varphi^\uparrow(r) + \varphi^\downarrow(s) + \frac{c}{2} [\|\theta - \theta^\nu\|^2 + (r - r_\nu)^2 + (s - s_\nu)^2] \\ & \text{subject to} && b^{1;i_2^\nu}(\theta; \theta^\nu) - \mathbf{1}r \leq 0, \text{ and } b^{2;i_1^\nu}(\theta; \theta^\nu) + \mathbf{1}s \leq 0 \end{aligned}$$

with $b^{1;i_2^\nu}(\bullet; \theta^\nu)$ and $b^{2;i_1^\nu}(\bullet; \theta^\nu)$ denoting two vector functions with convex differentiable components given by

$$\begin{aligned} b_i^{1;i_2^\nu}(\theta; \theta^\nu) &\triangleq \psi_{1,i}(\theta) - [h(\theta^\nu) + \nabla \psi_{2,i_2^\nu}(\theta^\nu)^T(\theta - \theta^\nu)] \quad \forall i = 1, \dots, k_1, \quad \text{and} \\ b_i^{2;i_1^\nu}(\theta; \theta^\nu) &\triangleq g(\theta^\nu) + \nabla \psi_{1,i_1^\nu}(\theta^\nu)^T(\theta - \theta^\nu) - \psi_{2,i}(\theta) \quad \forall i = 1, \dots, k_2. \end{aligned}$$

The Lagrangian dual program is

$$\begin{aligned}
 & \underset{\lambda \geq 0, \mu \geq 0}{\text{maximize}} \quad \zeta_{(i_1^\nu, i_2^\nu)}(\lambda, \mu; z^\nu), \quad \text{where} \\
 \zeta_{(i_1^\nu, i_2^\nu)}(\lambda, \mu; z^\nu) & \triangleq \underset{\theta \in \Theta, r, s \in \mathbb{R}}{\text{minimize}} \left\{ \begin{array}{l} \varphi^\uparrow(r) + \varphi^\downarrow(s) + \frac{c}{2} \|z - z^\nu\|^2 \\ + \lambda^T (b^{1;i_2^\nu}(\theta; \theta^\nu) - \mathbf{1} r) + \mu^T (b^{2;i_1^\nu}(\theta; \theta^\nu) + \mathbf{1} s) \end{array} \right\} \\
 & = \underbrace{\underset{\theta \in \Theta}{\text{minimize}} \left\{ \lambda^T b^{1;i_2^\nu}(\theta; \theta^\nu) + \mu^T b^{2;i_1^\nu}(\theta; \theta^\nu) + \frac{c}{2} \|\theta - \theta^\nu\|^2 \right\}}_{\text{subprob}_\theta^\nu} \\
 & \quad + \underbrace{\underset{r \in \mathbb{R}}{\text{minimize}} \left\{ \varphi^\uparrow(r) - (\mathbf{1}^T \lambda)r + \frac{c}{2} (r - r_\nu)^2 \right\}}_{\text{subprob}_r^\nu} \\
 & \quad + \underbrace{\underset{s \in \mathbb{R}}{\text{minimize}} \left\{ \varphi^\downarrow(s) + (\mathbf{1}^T \mu)s + \frac{c}{2} (s - s_\nu)^2 \right\}}_{\text{subprob}_s^\nu}.
 \end{aligned}$$

By Danskin's theorem [21, Theorem 10.2.1], the concave dual objective function $\zeta_{(i_1^\nu, i_2^\nu)}(\bullet, \bullet; z^\nu)$ is differentiable with the gradient given by

$$\nabla_\lambda \zeta_{(i_1^\nu, i_2^\nu)}(\lambda, \mu; z^\nu) = b^{1;i_2^\nu}(\bar{\theta}^\nu; \theta^\nu) - \mathbf{1} \bar{r}_\nu, \quad \nabla_\mu \zeta_{(i_1^\nu, i_2^\nu)}(\lambda, \mu; z^\nu) = b^{2;i_1^\nu}(\bar{\theta}^\nu; \theta^\nu) + \mathbf{1} \bar{s}_\nu,$$

where $\bar{\theta}^\nu$, \bar{r}_ν , and \bar{s}_ν are, respectively, the unique minimizers of $\text{subprob}_\theta^\nu$, subprob_r^ν , and subprob_s^ν corresponding to the pair (λ, μ) . Here, we see the benefit of regularizing the scalar variables r and s that ensures the uniqueness of the minimizers \bar{r}_ν and \bar{s}_ν . We can use the concept of proximal mappings to characterize these three minimizers. For a proper closed convex function ϕ defined on \mathbb{R}^n and a scalar $c > 0$, the proximal mapping P_ϕ^c is defined by

$$P_\phi^c(u) \triangleq \underset{v \in \mathbb{R}^n}{\text{argmin}} \left\{ \phi(v) + \frac{c}{2} \|v - u\|^2 \right\}, \quad u \in \mathbb{R}^n,$$

which is globally Lipschitz continuous [5, Proposition 12.27]. One property of this mapping is that a function ϕ is convex piecewise linear-quadratic if and only if P_ϕ^c is piecewise linear [50, Proposition 12.30]. (A function ϕ is called piecewise linear-quadratic if $\text{dom } \phi$ can be represented as the union of finitely many polyhedral sets, relative to each of which $\phi(x)$ is given by a quadratic function; see, e.g., [50, Definition 10.20].) Given that $b^{1;i_2^\nu}(\bullet; \theta^\nu)$ and $b^{2;i_1^\nu}(\bullet; \theta^\nu)$ are smooth mappings, the function $\zeta_{(i_1^\nu, i_2^\nu)}$ is an SC^1 function for many interesting instances, in particular, if Θ is a polyhedral set and the functions $\lambda^T b^{1;i_2^\nu}(\bullet; \theta^\nu) + \mu^T b^{2;i_1^\nu}(\bullet; \theta^\nu)$, φ^\uparrow , and φ^\downarrow are all convex piecewise linear-quadratic, or the proximal mappings of these functions are piecewise smooth.

The following is the globally convergent and locally superlinearly convergent SN method to minimize a convex SC^1 function ϕ on a closed convex set X ; see [46] for details. In the case of the dual program in question, the set X is the nonnegative orthant in the (λ, μ) -space. Specializations of the method are possible and will be illustrated with a least-squares piecewise affine regression problem.

An SN method for minimize _{$x \in X$} $\phi(x)$ with an SC¹ function ϕ .

Initialization. Given an initial point $x^0 \in X$ and two scalars $\rho, \sigma \in (0, 1)$. Set $k = 0$.

Step 1. (Find the search direction.) Given $x^k \in X$, pick $V^k \in \partial_C \nabla \phi(x^k)$ and a scalar $\varepsilon_k > 0$. Solve the following strictly convex quadratic program in the variable d to obtain d^k :

$$\underset{d \in X - x^k}{\text{minimize}} \quad \nabla \phi(x^k)^T d + \frac{1}{2} d^T (V^k + \varepsilon_k I) d$$

Step 2. (Line search by the Armijo rule.) Set $x^{k+1} \triangleq x^k + \rho^{m_k} d^k$, where m_k is the smallest nonnegative integer m for which

$$\phi(x^k + \rho^m d^k) \leq \phi(x^k) + \sigma \rho^m \nabla \phi(x^k)^T d^k.$$

Step 3. If x^{k+1} satisfies a prescribed stopping rule, terminate; otherwise, return to Step 1 with k replaced by $k + 1$. \square

The major computational cost of the above SN method is to find the search direction in Step 1. To accomplish this task, one may apply any quadratic programming solver. For the special case when ψ is a dc piecewise affine function as in (3), we may instead find the search direction by solving a system of linear equations, leading to a very effective, provably convergent, overall algorithm for solving problem (8) that includes the problem of piecewise affine regression.

5.1. A special case where ψ is difference-convex-piecewise-affine. Let ψ be a dc piecewise affine functions as in (3). Then the functions $b^{1;i_2^\nu}(\bullet; \theta^\nu)$ and $b^{2;i_1^\nu}(\bullet; \theta^\nu)$ in (15) are affine, for which we may write (dropping the indices i_1^ν and i_2^ν and the dependence on θ^ν) $b^1(\theta) \triangleq B^1 \theta - \beta^1$ and $b^2(\theta) \triangleq B^2 \theta - \beta^2$ for some matrices $B^1 \in \mathbb{R}^{k_1 \times m}$ and $B^2 \in \mathbb{R}^{k_2 \times m}$ and vectors $\beta^1 \in \mathbb{R}^{k_1}$ and $\beta^2 \in \mathbb{R}^{k_2}$. Instead of (10), we may consider the following alternative subproblems in the MM algorithm with two additional auxiliary variables $\hat{r} \in \mathbb{R}^{k_1}$ and $\hat{s} \in \mathbb{R}^{k_2}$:

$$\begin{aligned} \underset{\theta \in \Theta, r, s, \hat{r}, \hat{s}}{\text{minimize}} \quad & \varphi^\uparrow(r) + \varphi^\downarrow(s) \\ & + \frac{c}{2} [\|\theta - \theta^\nu\|^2 + (r - r_\nu)^2 + (s - s_\nu)^2 + \|\hat{r} - \hat{r}^\nu\|^2 + \|\hat{s} - \hat{s}^\nu\|^2] \\ \text{subject to} \quad & B^1 \theta - \mathbf{1} r + \hat{r} = \beta^1, \quad B^2 \theta + \mathbf{1} s + \hat{s} = \beta^2, \quad \hat{r} \geq 0, \quad \hat{s} \geq 0. \end{aligned}$$

The (subsequential) convergence of the MM algorithms discussed in section 4 can be easily extended to the above formulation with the additional regularization term $\frac{c}{2} [\|\hat{r} - \hat{r}^\nu\|^2 + \|\hat{s} - \hat{s}^\nu\|^2]$. The Lagrangian dual program is

$$\underset{\lambda \in \mathbb{R}^{k_1}, \mu \in \mathbb{R}^{k_2}}{\text{maximize}} \quad \xi(\lambda, \mu; \hat{z}^\nu), \quad \text{where} \quad \hat{z}^\nu \triangleq (\theta^\nu, r_\nu, s_\nu, \hat{r}^\nu, \hat{s}^\nu) \quad \text{and}$$

$$\begin{aligned}
\xi(\lambda, \mu; \hat{z}^\nu) &\triangleq \underset{\substack{\theta \in \Theta, r, s \in \mathbb{R}, \\ \hat{r} \geq 0, \hat{s} \geq 0}}{\text{minimize}} \left\{ \begin{array}{l} \varphi^\uparrow(r) + \varphi^\downarrow(s) + \frac{c}{2} [\|\theta - \theta^\nu\|^2 + (r - r_\nu)^2] \\ + (s - s_\nu)^2 + \|\hat{r} - \hat{r}^\nu\|^2 + \|\hat{s} - \hat{s}^\nu\|^2 \\ + \lambda^T(B^1\theta - \mathbf{1}r + \hat{r} - \beta^1) + \mu^T(B^2\theta + \mathbf{1}s + \hat{s} - \beta^2) \end{array} \right\} \\
&= -\lambda^T\beta^1 - \mu^T\beta^2 + \underset{\theta \in \Theta}{\text{minimize}} \left\{ \lambda^T B^1 \theta + \mu^T B^2 \theta + \frac{c}{2} \|\theta - \theta^\nu\|^2 \right\} \\
&\quad + \underset{r \in \mathbb{R}}{\text{minimize}} \left\{ \varphi^\uparrow(r) - (\mathbf{1}^T \lambda) r + \frac{c}{2} (r - r_\nu)^2 \right\} \\
&\quad + \underset{s \in \mathbb{R}}{\text{minimize}} \left\{ \varphi^\downarrow(s) + (\mathbf{1}^T \mu) s + \frac{c}{2} (s - s_\nu)^2 \right\} \\
&\quad + \underset{\hat{r} \geq 0}{\text{minimize}} \left\{ \lambda^T \hat{r} + \frac{c}{2} \|\hat{r} - \hat{r}^\nu\|^2 \right\} + \underset{\hat{s} \geq 0}{\text{minimize}} \left\{ \mu^T \hat{s} + \frac{c}{2} \|\hat{s} - \hat{s}^\nu\|^2 \right\}.
\end{aligned}$$

As in the previous discussions, the concave function $\xi(\bullet, \bullet; \hat{z}^\nu)$ is continuously differentiable. The proximal mappings associated with \hat{r} and \hat{s} are essentially the projections onto the nonnegative orthant, which must be semismooth [21, Theorem 4.5.2 and Proposition 7.4.6]. If Θ is a polyhedral set and the proximal mappings of φ^\uparrow and φ^\downarrow are also semismooth, the dual program is an unconstrained SC¹ problem. Step 1 in the SN method is thus to solve an unconstrained strictly convex quadratic program or, equivalently, to solve a linear equation. Furthermore, the minimization in θ within the dual function $\xi(\lambda, \mu; \hat{z}^\nu)$ is the projection onto the set Θ , which is easy if Θ is a simple set, e.g., containing upper and lower bounds only. In this case, the bulk of the computational effort per iteration of the combined MM and SN methods essentially consists of solving systems of linear equations in the (λ, μ) -space within the fast convergent SN method.

6. Numerical experiments. In order to demonstrate the effectiveness of the proposed nonsmooth MM+SN method for solving some nonstandard statistical estimation problems, we conduct various numerical experiments on an unconstrained least-squares continuous piecewise affine regression problem by taking the error function $\varphi(y, t) = \frac{1}{2}(t - y)^2$ and the piecewise affine model ψ as in (3). Besides providing evidence of the promise of the MM+SN method itself (see Table 1 and Figures 12 and 13 for summaries of its performance statistics), our experiments also aim to establish the superiority of the piecewise affine statistical model over the standard linear model.

Specifically, the optimization problem we consider in this section is the following:

$$\underset{\theta}{\text{minimize}} \left\{ f_N(\theta) \triangleq \frac{1}{N} \sum_{s=1}^N \left[y_s - \max_{1 \leq i \leq k_1} \left\{ (a^i)^T x^s + \alpha_i \right\} - \max_{1 \leq i \leq k_2} \left\{ (b^i)^T x^s + \beta_i \right\} \right]^2 + \gamma_N \sum_{i=1}^m [|\theta_i| - p_i(\theta_i)] \right\},$$

where $\theta = \{(a^i, \alpha_i)_{i=1}^{k_1}, (b^i, \beta_i)_{i=1}^{k_2}\} \in \mathbb{R}^{(k_1+k_2)(d+1)}$ and each p_i is a differentiable convex function. For this optimization problem which must have a global minimizer, Proposition 2 yields some interesting properties of its d-stationary points, provided that each univariate $p_i(\theta_i)$ is piecewise linear-quadratic. All our computations are done in MATLAB on a Mac OS X with 1.7 GHz Intel Core i7 and 8 GB RAM. We employ the randomized version of the MM algorithm to solve the problem with $\varepsilon = 10^{-4}$ for the “ ε -argmax” expansion. The iterations of the MM algorithm are terminated if $\frac{|f_N(\theta^{\nu+1}) - f_N(\theta^\nu)|}{\max(1, |f_N(\theta^\nu)|)} \leq 10^{-4}$; the iterations of the inner SN algorithm are

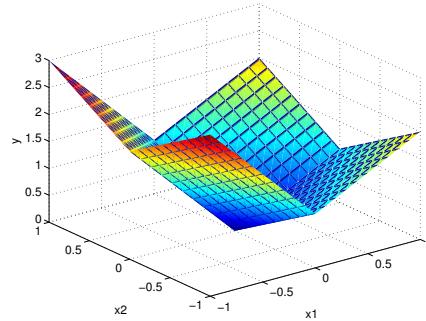


FIG. 3. The 3-dimensional plot of Example 1.

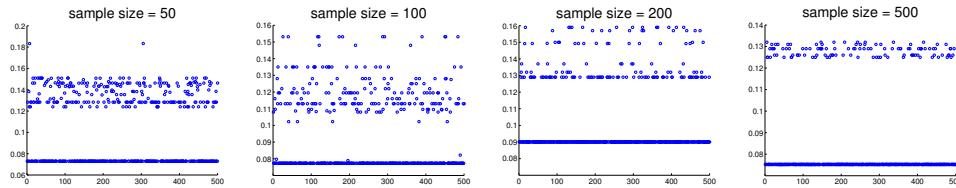


FIG. 4. The objective values computed by the MM+SN algorithm of Example 1.

terminated if $\|\nabla \xi(\lambda^k, \mu^k; \hat{z}^\nu)\| \leq \max(10^{-6}, 10^{-2} |f_N(\theta^{\nu+1}) - f_N(\theta^\nu)|)$, where ξ is the objective of the ν th dual subproblem.

6.1. Synthetic data. In this set of runs we consider two simulation examples. In both examples the penalty parameter γ_N is taken to be 0.

Example 3. We consider a 2-dimensional convex piecewise linear model

$$y = \max \{x_1 + x_2, x_1 - x_2, -2x_1 + x_2, -2x_1 - x_2\} + \varepsilon,$$

whose 3-dimensional plot is given in Figure 3. Although this is a convex model, the resulting least-squares problem is still nonconvex. Random samples of size $N = 50, 100, 200, 500$ are generated uniformly from $[-1, 1] \times [-1, 1]$, respectively, and ε is drawn from a uniform distribution in $[-0.5, 0.5]$. The objective values of the computed solutions for $k_1 = 4$ and $k_2 = 0$ over 500 runs (with each run corresponding to one initial point) by the MM+SN algorithm are shown in Figure 4. Two observations are made from this figure. One, the d-stationary values (which are also locally minimum values) are clustered into a few distinguished groups. In a recent paper [15], we provide theoretical justification of this observation by showing that there are only finitely many d-stationary values for several classes of piecewise programs, of which the least-squares piece affine regression problem is a special instance. Two, with the increase of the sample size, more and more initial points lead to the solutions that appear to yield the globally minimum value. The total numbers of initial points that lead to the smallest objective values with respect to different sample sizes are listed in Figure 5. In Figure 6, we plot the solutions given by the smallest objective values of different sample sizes in Figure 4. These solutions almost recover the original convex regression model, thus (empirically) validating the MM+SN methodology for a nonconvex minimization problem in the recovery of a convex statistical model.

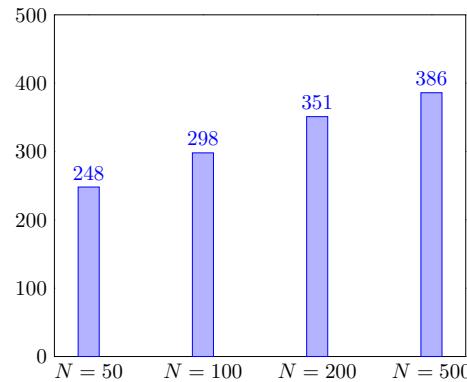


FIG. 5. The number of initial points that lead to the smallest objective values.

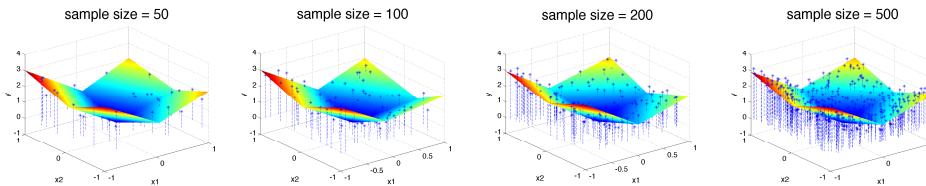


FIG. 6. The solutions of Example 1 with the smallest objective values.

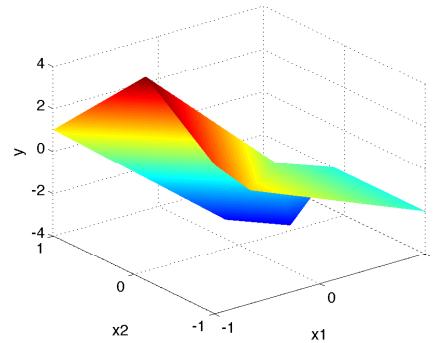


FIG. 7. The 3-dimensional plot of Example 2.

Example 4. Next we consider a 2-dimensional nonconvex piecewise affine model

$$y = \max \{x_1 - 2x_2, -2x_1 + x_2 + 1\} - \max \{3x_1 - 2x_2, 2x_1 + 5x_2\} + \varepsilon,$$

whose 3-dimensional plot is given in Figure 7. The projection of this figure onto $y = 0$ is demonstrated in Figure 8. As in Example 3, random samples of size $N = 50, 100, 200, 500$ are generated uniformly from $[-1, 1] \times [-1, 1]$, and ε is drawn from a uniform distribution in $[-0.5, 0.5]$. The objective values of the computed solutions over 500 runs are shown in Figure 9 summarized in Figure 10. Similar observations to Example 3 can be made with regards to the concentration of the (computed) stationary values. We also plot the solutions corresponding to the smallest objective values for different sample sizes in Figure 11.

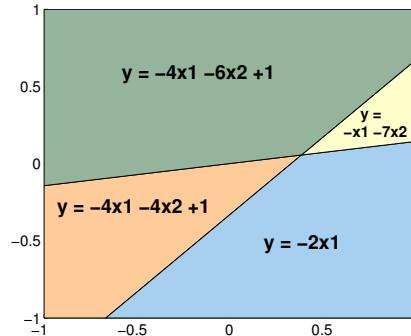
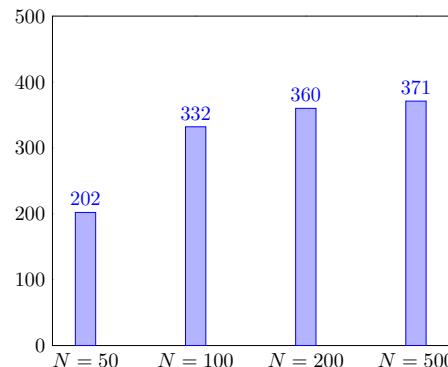
FIG. 8. The projection of Example 2 onto $y = 0$.

FIG. 9. The number of initial points that lead to the smallest objective values.

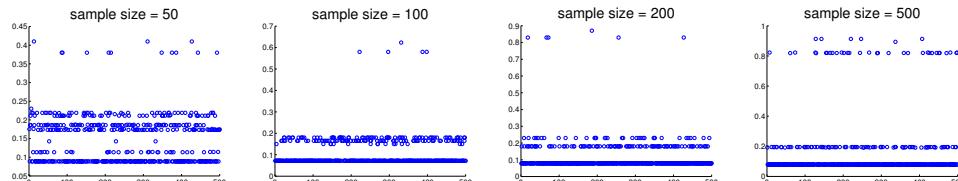


FIG. 10. The objective values computed by the MM+SN algorithm of Example 2.

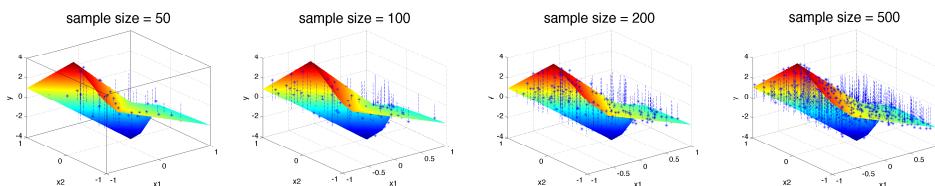


FIG. 11. The solutions of Example 2 with the smallest objective values.

6.2. Real data. We also conduct the experimental evaluation of the proposed algorithm on the real datasets obtained from the UCI Repository of machine learning databases <https://archive.ics.uci.edu/ml/index.php>. We first list the problem characteristics and report the performance of the MM+SN method averaged over the pairs

TABLE 1
The performance of the MM+SN algorithm.

Problem name	<i>N</i>	<i>d</i>	Iterations		Time
			MM	SN	
Banknote authentication	1372	4	9	72	3.5 s
Concrete compressive strength	1030	8	23	197	10.8 s
Auto MPG	392	7	25	185	3.9 s
Airfoil self-noise	1503	5	18	156	12.3 s
Libras movement	360	91	22	187	14.7 s
Communities and crime	2215	147	31	251	43.7 s

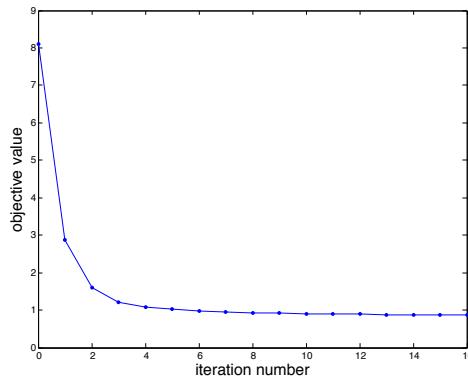


FIG. 12. *The performance of the MM algorithm.*

(k_1, k_2) employed in the experiments, i.e., satisfying $1 \leq k_1, k_2 \leq 4$ in Table 1, where the two columns under “iterations” present, respectively, the number of iterations of the MM algorithm and the total numbers of iterations of the SN algorithm. For the first four problems where the sample sizes N are much larger than the dimensions d , no sparse regularization term is added to the model. For the other two problems where d is relatively large, we take the SCAD function as the sparsity inducing penalty function. The sparse penalty parameter is estimated by fivefold cross validation. One can see the efficiency of the proposed MM+SN method, where it usually takes no more than 10 SN iterations for each MM subproblem. As an illustration, Figures 12 and 13 exhibit the performance of the MM algorithm (objective values versus iteration number) and the SN algorithm ($\log_{10} \|\nabla \xi(\lambda^k, \mu^k; \hat{z}^\nu)\|$ versus iteration number) for solving the “Auto MPG” problem with $(k_1, k_2) = (2, 2)$.

To demonstrate the advantage of using the piecewise affine model over the classical linear regression model, we compare the solution of the least-squares piecewise affine regression θ^{PA} with the ordinary least-squares solution

$$\theta^{\text{LS}} \in \operatorname{argmin}_{\theta} \frac{1}{2N} \sum_{s=1}^N \|y^s - (x^s, 1)^T \theta\|^2$$

in the following way. Let the full dataset of each instance obtained from the UCI Repository be T , which is further divided into a training set T_{training}^f and a test set T_{test}^f for $f = 1, \dots, 5$ by fivefold cross validation. For each fold, we find the estimator $\theta^{\text{PA}, f}$ based on the training set T_{training}^f by choosing the solution with the smallest objective value over 20 runs of the MM+SN algorithm (with each run corresponding

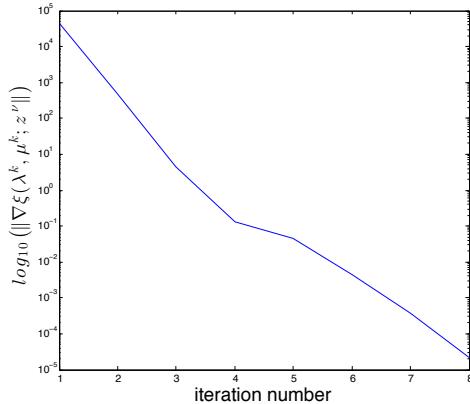


FIG. 13. The performance of the SN algorithm.

TABLE 2
Ratio of the prediction error for different choices of (k_1, k_2) .

		Banknote authentication						Concrete compressive strength				
		k_1	1	2	3	4	k_1	1	2	3		
k_2		1	0.74	0.68	0.67	1	1.00	0.74	0.47	0.38	4	
1		1.00	0.74	0.68	0.67	2	0.84	0.73	0.65	0.63		
2		0.84	0.73	0.65	0.63	3	0.81	0.72	0.74	0.74		
3		0.81	0.72	0.74	0.74	4	0.85	0.71	0.69	0.74		
Auto MPG												
		k_1	1	2	3	4			Airfoil self-noise			
		1	1.00	0.77	0.72	0.77			k_1	1	2	3
k_2		1	1.00	0.77	0.72	0.77	1	1.00	0.83	0.77	0.74	4
1		1.00	0.77	0.72	0.77	2	0.97	0.74	0.79	0.73		
2		0.97	0.74	0.79	0.73	3	0.83	0.78	0.76	0.76		
3		0.83	0.78	0.76	0.76	4	0.85	0.84	0.73	0.73		
Libras movement												
		k_1	1	2	3	4			Communities and crime			
		1	1.00	0.93	0.82	0.80			k_1	1	2	3
k_2		1	1.00	0.93	0.82	0.80	1	1.00	0.88	0.85	0.83	4
1		1.00	0.93	0.82	0.80	2	0.83	0.86	0.77	0.77		
2		0.83	0.86	0.77	0.77	3	0.78	0.72	0.72	0.76		
3		0.78	0.72	0.72	0.76	4	0.90	0.73	0.78	0.80		

to one initial point). Such a $\theta^{\text{PA},f}$ is likely to be an (unproven) global minimizer of the nonconvex least-squares piecewise affine regression program. We compute the prediction errors of the two estimators, respectively, by

$$\begin{aligned} E_{\text{PA}} &\triangleq \sum_{f=1}^5 \frac{1}{|T_{\text{test}}^f|} \sum_{s \in T_{\text{test}}^f} (y_s - \psi(x^s; \theta^{\text{PA},f}))^2, \\ E_{\text{LS}} &\triangleq \sum_{f=1}^5 \frac{1}{|T_{\text{test}}^f|} \sum_{s \in T_{\text{test}}^f} (y_s - (x^s, 1)^T \theta^{\text{LS},f})^2. \end{aligned}$$

We take 100 simulations and report the average ratio $E_{\text{PA}}/E_{\text{LS}}$ for different choices of (k_1, k_2) in Table 2. It can be seen that all piecewise affine results are better than the

linear regression results ($k_1 = k_2 = 1$); moreover, with proper choices of (k_1, k_2) , the prediction error given by the piecewise affine regression model is significantly smaller than that given by the linear regression model as highlighted by the best entry in each table.

As a final remark, we have observed in our simulation studies that even when the true model is linear, the continuous piecewise affine regression fits (for different choices of k_1 and k_2) are very close to that of linear regression, i.e., the ratios of the prediction errors (in fivefold cross validation as given in Table 2) are all very close to 1. This provides further evidence that the continuous piecewise affine regression model could be a worthwhile and useful extension of linear regression.

In conclusion, we have demonstrated the practical worthiness of the least-squares piecewise affine regression model formulated as a nonconvex nondifferentiable program and solved by the nonmonotone MM+SN method. Further applications of this method to other statistical estimation problems will be reported separately.

REFERENCES

- [1] M. AHN, J.-S. PANG, AND J. XIN, *Difference-of-convex learning: Directional stationarity, optimality, and sparsity*, SIAM J. Optim., 27 (2017), pp. 1637–1665.
- [2] H. ATTOUCH AND J. BOLTE, *On the convergence of the proximal algorithm for nonsmooth functions involving analytic features*, Math. Program., 116 (2009), pp. 5–16.
- [3] H. ATTOUCH, J. BOLTE, AND B.F. SVAITER, *Convergence of descent methods for semi-algebraic and tame problems: Proximal algorithms, forward-backward splitting, and regularized Gauss-Seidel methods*, Math. Program., 137 (2013), pp. 91–129.
- [4] A.M. BAGIROV, C. CLAUSEN, AND M. KOHLER, *An algorithm for the estimation of a regression function by continuous piecewise linear functions*, Comput. Optim. Appl., 45 (2010), pp. 159–179.
- [5] H.H. BAUSCHKE AND P.L. COMBETTES, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*, LMS Books Math., Springer, New York, 2011.
- [6] D.P. BERTSEKAS, *Nonlinear Programming*, 3rd ed., Athena Scientific, Belmont, MA, 2016.
- [7] P.J. BICKEL AND K. DOKSUM, *Mathematical Statistics, Basic Ideas and Selected Topics*, Prentice Hall, Upper Saddle River, NJ, 2006.
- [8] J. BOLTE, A. DANILIDIS, AND A. LEWIS, *The Lojasiewicz inequality for nonsmooth subanalytic functions with applications to subgradient dynamical systems*, SIAM J. Optim., 17 (2007), pp. 1205–1223.
- [9] J. BOLTE AND E. PAUWELS, *Majorization-minimization procedures and convergence of SQP methods for semi-algebraic and tame programs*, Math. Oper. Res., 41 (2016), pp. 442–465.
- [10] J. BOLTE, S. SABACH, AND M. TEBOULLE, *Proximal alternating linearized minimization for nonconvex and nonsmooth problems*, Math. Program., 146 (2014), pp. 459–494.
- [11] L.D. BROWN, *Fundamentals of Statistical Exponential Families*, IMS Lect. Notes Monogr. Ser., Institute of Mathematical Statistics, Hayward, CA, 1986.
- [12] E.J. CANDÈS, M. WAKIN, AND S. BOYD, *Enhancing sparsity by reweighted ℓ_1 minimization*, J. Fourier Anal. Appl., 14 (2008), pp. 877–905.
- [13] F.H. CLARKE, *Optimization and Nonsmooth Analysis*, Classics Appl. Math. 5, SIAM, Philadelphia, 1990.
- [14] C. CORTES AND V. VAPNIK, *Support-vector networks*, Mach. Learn., 20 (1995), pp. 273–297.
- [15] Y. CUI AND J.-S. PANG, *On the Finite Number of Directional Stationary Values of Piecewise Programs*, <https://arxiv.org/abs/1803.00190> (2018).
- [16] Y. CUI, T.H. CHANG, M. HONG, AND J.-S. PANG, *A Study of Piecewise Linear-Quadratic Programs*, <https://arxiv.org/abs/1709.05758> (2018).
- [17] A.P. DEMPSTER, N.M. LAIRD, AND D.B. RUBIN, *Maximum likelihood from incomplete data via the EM algorithm*, J. Roy. Statist. Soc. B, 39 (1977), pp. 1–38.
- [18] H. DONG, M. AHN, AND J.-S. PANG, *Structural properties of affine sparsity constraints*, Math. Program., to appear.
- [19] D. DUNSON AND L.A. HANNAH, *Bayesian Nonparametric Multivariate Convex Regression*, <https://arxiv.org/abs/1109.0322> (2011).
- [20] D. DUNSON AND L.A. HANNAH, *Multivariate convex regression with adaptive partitioning*, J. Mach. Learn. Res., 14 (2013), pp. 3261–3294.

- [21] F. FACCHINEI AND J.-S. PANG, *Finite-Dimensional Variational Inequalities and Complementarity Problems*, Springer, New York, 2003.
- [22] J. FAN AND R. LI, *Variable selection via nonconcave penalized likelihood and its oracle properties*, J. Amer. Statist. Assoc., 96 (2001), pp. 1348–1360.
- [23] X. GLOROT, A. BORDES, AND Y. BENGIO, *Deep sparse rectifier neural networks*, in Proceedings of the 14th International Conference on Artificial Intelligence and Statistics, JMLR Workshop Conf. Proc. 15, Journal of Machine Learning Research, Cambridge, MA, 2011, pp. 315–323.
- [24] G. HAHN, M. BANERGJEE, AND B. SEN, *Parameter estimation and inference in a continuous piecewise linear regression model*, Technical report, Columbia University, New York, 2017.
- [25] P. HARTMAN, *On functions representable as a difference of convex functions*, Pac. J. Optim., 9 (1959), pp. 707–713.
- [26] T. HASTIE, R. TIBSHIRANI, AND M. WEINWRIGHT, *Statistical Learning with Sparsity: The Lasso and Generalizations*, CRC Press, Boca Raton, FL, 2015.
- [27] D.R. HUNTER AND K. LANGE, *A tutorial on MM algorithms*, Amer. Statist., 58 (2004), pp. 30–37.
- [28] R. KOENKER AND G. BASSETT, JR., *Regression quantiles*, Econometrica, 46 (1978), pp. 33–50.
- [29] M. KOJIMA AND S. SHINDO, *Extensions of Newton and quasi-Newton methods to systems of PC^1 equations*, J. Oper. Res. Soc. Japan, 29 (1986), pp. 352–374.
- [30] B. KUMMER, *Newton's method for nondifferentiable functions*, in Advances in Mathematical Optimization, Akademie, Berlin, 1988, pp. 114–125.
- [31] K. LANGE, *MM Optimization Algorithms*, Other Titles Appl. Math. 147, SIAM, Philadelphia, 2016.
- [32] H.A. LE THI, V.N. HUYNH, AND D.T. PHAM, *Convergence analysis of difference-of-convex algorithm with subanalytic data*, J. Optim. Theory Appl., 179 (2018), pp. 103–126.
- [33] H.A. LE THI AND D.T. PHAM, *The DC programming and DCA revised with DC models of real world nonconvex optimization problems*, Ann. Oper. Res., 133 (2005), pp. 25–46.
- [34] H.A. LE THI, D.T. PHAM, AND X.T. VO, *DC approximation approaches for sparse optimization*, European J. Oper. Res., 244 (2015), pp. 26–46.
- [35] X. LI, D.F. SUN, AND K.-C. TOH, *A highly efficient semismooth Newton augmented Lagrangian method for solving Lasso problems*, SIAM J. Optim., 28 (2018), pp. 433–458.
- [36] Z. LU, Z. ZHOU, AND Z. SUN, *Enhanced proximal DC algorithms with extrapolation for a class of structured nonsmooth DC minimization*, Math. Program., to appear.
- [37] J. MAIRAL, *Optimization with first-order surrogate functions*, in Proceedings of the 30th International Conference on Machine Learning, Curran, Red Hook, NY, 2013, pp. 783–791.
- [38] J. MAIRAL, *Incremental majorization-minimization optimization with application to large-scale machine learning*, SIAM J. Optim., 25 (2015), pp. 829–855.
- [39] R. MAZUMDER, A. CHOUDHURY, G. IYENGAR, AND B. SEN, *A computational framework for multivariate convex regression and its variants*, J. Amer. Statist. Assoc., to appear.
- [40] R. MIFFLIN, *Semismooth and semiconvex functions in constrained optimization*, SIAM J. Control Optim., 15 (1977), pp. 959–972.
- [41] A. MILZAREK AND M. ULRICH, *A semismooth Newton method with multidimensional filter globalization for ℓ_1 optimization*, SIAM J. Optim., 24 (2014), pp. 298–333.
- [42] B.S. MORDUKHOVICH, *Variational Analysis and Generalized Differentiation I*, Springer, New York, 2006.
- [43] V. NAIR AND G.E. HINTON, *Rectified linear units improve restricted Boltzmann machines*, in Proceedings of the 27th International Conference on Machine Learning, Omnipress, Madison, WI, 2010, pp. 807–814.
- [44] J.M. ORTEGA AND W.C. RHEINBOLDT, *Iterative Solution of Nonlinear Equations in Several Variables*, Classics Appl. Math. 30, SIAM, Philadelphia, 2000.
- [45] J.-S. PANG, *Newton's method for B -differentiable equations*, Math. Oper. Res., 15 (1990), pp. 311–341.
- [46] J.-S. PANG AND L. QI, *A globally convergent Newton method for convex SC^1 minimization problems*, J. Optim. Theory Appl., 85 (1995), pp. 633–648.
- [47] J.-S. PANG, M. RAZAVIYAYN, AND A. ALVARADO, *Computing B -stationary points of nonsmooth DC programs*, Math. Oper. Res., 42 (2016), pp. 95–118.
- [48] H. QI AND D. SUN, *A quadratically convergent Newton method for computing the nearest correlation matrix*, SIAM J. Matrix Anal. Appl., 28 (2006), pp. 360–385.
- [49] L. QI AND J. SUN, *A nonsmooth version of Newton's method*, Math. Program., 58 (1993), pp. 353–367.
- [50] R.T. ROCKAFELLAR AND R.J.-B. WETS, *Variational Analysis*, Springer, New York, 1998.
- [51] S. SCHOLTES, *Introduction to Piecewise Differentiable Equations*, SpringerBriefs Optim., Springer, New York, 2002.

- [52] J. SUN, *On Monotropic Piecewise Quadratic Programming*, Ph.D. dissertation, Department of Applied Mathematics, University of Washington, Seattle, WA, 1986.
- [53] J. SUN, *On the structure of convex piecewise quadratic functions*, J. Optim. Theory Appl., 72 (1992), pp. 499–510.
- [54] R. TIBSHIRANI, M. SAUNDERS, S. ROSSET, J. ZHU, AND K. KNIGHT, *Sparsity and smoothness via the fused lasso*, J. Roy. Statist. Soc. B, 67 (2005), pp. 91–108.
- [55] B. WEN, X. CHEN, AND T.K. PONG, *A proximal difference-of-convex algorithm with extrapolation*, Comput. Optim. Appl., 2 (2018), pp. 297–324.
- [56] C.J. WU, *On the convergence properties of the EM algorithm*, Ann. Statist., 11 (1983), pp. 95–103.
- [57] L.Q. YANG, D.F. SUN, AND K.C. TOH, *SDPNAL+: A majorized semismooth Newton-CG augmented Lagrangian method for semidefinite programming with nonnegative constraints*, Math. Program. Comput., 7 (2015), pp. 331–366.
- [58] P. YIN, Y. LOU, Q. HE, AND J. XIN, *Minimization of $\ell_1 - \ell_2$ for compressed sensing*, SIAM J. Sci. Comput., 37 (2015), pp. A536–A563.
- [59] D. YU AND L. DENG, *Automatic Speech Recognition: A Deep Learning Approach*, Signals Commun. Technol., Springer, London, 2015.
- [60] C. ZHANG, *Nearly unbiased variable selection under minimax concave penalty*, Ann. Statist., 38 (2010), pp. 894–942.
- [61] C. ZHANG, M. PHAM, S. FU, AND Y. LIU, *Robust multiclass support vector machines using difference convex algorithm*, Math. Program., 169 (2018), pp. 277–305.
- [62] X.-Y. ZHAO, D. SUN, AND K.-C. TOH, *A Newton-CG augmented Lagrangian method for semidefinite programming*, SIAM J. Optim., 20 (2010), pp. 1737–1765.