

Sparse Linear Regression with Constraints: A Flexible Entropy-based Framework

Amber Srivastava, Alisina Bayati, Srinivasa M. Salapaka

Abstract—This work presents a new approach to solve the sparse linear regression problem, i.e., to determine a k -sparse vector $\mathbf{w} \in \mathbb{R}^d$ that minimizes the cost $\|\mathbf{y} - A\mathbf{w}\|_2^2$. In contrast to the existing methods, our proposed approach splits this k -sparse vector into two parts — (a) a column stochastic binary matrix V , and (b) a vector $\mathbf{x} \in \mathbb{R}^k$. Here, the binary matrix V encodes the location of the k non-zero entries in \mathbf{w} . Equivalently, it encodes the subset of k columns in the matrix A that map \mathbf{w} to \mathbf{y} . We demonstrate that this enables modeling several non-trivial application specific structural constraints on \mathbf{w} as constraints on V . The vector \mathbf{x} comprises of the actual non-zero values in \mathbf{w} . We use Maximum Entropy Principle (MEP) to solve the resulting optimization problem. In particular, we ascribe a probability distribution to the set of all feasible binary matrices V , and iteratively determine this distribution and the vector \mathbf{x} such that the associated Shannon entropy gets minimized, and the regression cost attains a pre-specified value. The resulting algorithm employs homotopy from the convex entropy function to the non-convex cost function to avoid poor local minimum. We demonstrate the efficacy and flexibility of our proposed approach in incorporating a variety of practical constraints, that are otherwise difficult to model using the existing benchmark methods.

I. INTRODUCTION

Sparse solutions to the linear regression problems have been of interest to multiple fields such as signal and image processing, genomics, economics and finance, flight load prediction, machine learning, and remote sensing [1]. One of the fundamental ways to formulate this problem is the best subset selection problem, where given a matrix $A = [a_1 \ a_2 \ \dots \ a_d] \in \mathbb{R}^{n \times d}$, a measurement vector $\mathbf{y} \in \mathbb{R}^n$, and a sparsity level $k (\ll d)$, we solve

$$\min_{\mathbf{w} \in \mathbb{R}^d} \|\mathbf{y} - A\mathbf{w}\|_2^2, \quad \text{subject to } \|\mathbf{w}\|_0 \leq k, \quad (1)$$

where $\|\mathbf{w}\|_0$ is the number of non-zero entries in \mathbf{w} . In other words, the optimization problem (1) determines (a) the best subset of k -columns (features) out of the d columns $\{a_1, \dots, a_d\}$ in A , and (b) their corresponding coefficients such that the vector \mathbf{w} linearly maps to the measurement vector \mathbf{y} with minimum squared euclidean loss in (1).

Note that the $\|\cdot\|_0$ norm is non-convex and the optimization problem (1) is NP-hard [2]. A large number of work done in this area design approximate solutions to (1), where they iteratively add or remove the non-zero coefficients in \mathbf{w} to minimize (1); for instance matching pursuit [3] and

forward-backward approaches [4]. See [5] for a detailed survey. Other relevant approaches involve replacing the non-convex constraint in (1) with a sparsity promoting term $\mathcal{T}(\mathbf{w})$ in the objective function, and solving the problem

$$\min_{\mathbf{w} \in \mathbb{R}^d} \|\mathbf{y} - A\mathbf{w}\|_2^2 + \lambda \mathcal{T}(\mathbf{w}), \quad (2)$$

where λ is a regularization parameter. A popular choice for $\mathcal{T}(\mathbf{w})$ is the l_1 norm $\|\mathbf{w}\|_1$, which results into a convex optimization problem (2). Algorithms such as gradient projection [6], iterative shrinkage-thresholding [7], and (linear) alternating direction method [8] guarantee globally optimal solutions, and several heuristics such as orthogonal matching pursuit [9] and least angle regression [10] efficiently address this convex program. Another class of choice for $\mathcal{T}(\mathbf{w})$ that has received much attention lately are the non-convex regularizers, which have been shown to result into better solutions than their convex counterpart [11]. Some of these choices are $\|\mathbf{w}\|_p$ with $0 < p < 1$ [12], minimax concave penalty (MCP) [13], smoothly clipped absolute deviation (SCAD) [14], and trimmed lasso [11], [15]. The latter has the additional property that it *exactly* results into a sparsity level k as indicated in the optimization problem (1).

As is evident from above, the work done in this area is extensive, with several different proposed frameworks that address various aspects of the problem such as scalability, computational costs, bias and exactness of sparsity. See [5] for a survey on these methods. Various scenarios such as (overlapping) grouped variables [16], [17], shape constraints [18] and restricted non-zero values [19] impose additional constraints on the design of the sparse vector \mathbf{w} in (1). Though there are methods to address such specific structural constraints, there is, to the best of our understanding, limited work on a generalized framework that effectively *models* and incorporates such constraints in (1).

The sparsity and structural constraints on \mathbf{w} can alternatively be viewed as constraints on the selection of the feature vectors $\{a_j\}_{j=1}^d$ from the matrix A . Thus, a *direct* control over the selection of these feature vectors will provide flexibility in modeling a variety of structural constraints discussed above, and also in enforcing the sparsity level k of the vector \mathbf{w} (which, generally speaking, is also a structural constraint). To this end, we develop a framework that (a) provides a direct control over the desired level of sparsity in the vector \mathbf{w} , (b) is flexible to incorporate a wide-range of application specific structural constraints on \mathbf{w} , and (c) results into an algorithm that is designed to avoid poor local minima of the underlying non-convex optimization problem.

A. Srivastava and A. Bayati contributed equally to this work.

A. Srivastava is with the Department of Mechanical Engineering, Indian Institute of Technology Delhi, India. asrvstv@iitd.ac.in

A. Bayati and S.M. Salapaka are with the Coordinated Science Laboratory, University of Illinois Urbana Champaign, USA. abayati2@salapaka@illinois.edu

The above contribution (a) result from our viewpoint of the optimization problem (1), where we dissociate the k -sparse vector $\mathbf{w} \in \mathbb{R}^d$ into two parts — a binary column stochastic matrix $V \in \{0, 1\}^{d \times k}$ and a vector $\mathbf{x} \in \mathbb{R}^k$. The matrix V is designed to encode the location of the k non-zero values in \mathbf{w} . Equivalently, it directly controls subset of k columns (features) in the design matrix A , that map the non-zero entries in \mathbf{w} to the measurement \mathbf{y} . The vector $\mathbf{x} \in \mathbb{R}^k$ comprises of these non-zero values in $\mathbf{w} \in \mathbb{R}^d$. As elaborated in the Section II, the column stochasticity and the size of the binary matrix V , and the size of the vector \mathbf{x} guarantee that the level k of sparsity in \mathbf{w} is *exactly* achieved.

The contribution (b) also results from the decision matrix V . Since the matrix V governs the choice of the k features in the design matrix A , it explicitly enables modeling several structural constraints that restrict the *permissible choice of subsets of k features* in A . For example, (as demonstrated later) constraints such as selecting only 2 out of the 4 given features, not allowing all features in a given subset $\{a_{l_1}, a_{l_2}, a_{l_3}\}$ to be selected, or modeling existing constraints such as selecting pre-defined groups of features (popularly addressed using group lasso [20]) can be conveniently modelled as structural constraints on V . As far as we are aware, our proposed framework is the most flexible in incorporating such variety of constraints on the permitted choice of the features; primarily owing to the matrix parameter V in our model that determines the choice of the features in A .

The contribution (c) results from the use of Maximum Entropy Principle (MEP) in determining the matrix V and the vector \mathbf{x} . Note that the matrix V is a discrete decision variable that lies in a *combinatorially* large set \mathcal{V} of all possible binary column stochastic matrices. Thus, the sparse linear regression (SLR) problem, with V and \mathbf{x} as the decision variable, can be viewed as a combinatorial optimization problem. In the past, MEP-based frameworks have successfully addressed a variety of such problems; for instance the facility location problem [21], data aggregation [22] and network design [23]. The abstract idea behind all these frameworks is to consider a distribution over the set of all possible values of the discrete variable. Then, determine the distribution that maximizes the Shannon entropy [21] at a pre-specified value of the expected cost function. This results into an iterative process, wherein the pre-specified value is successively lowered to as small value as possible and the solution from the previous iteration forms an initialization for the next. These iterations mimic a homotopy from the convex entropy to the non-convex cost function, which prevents the algorithm from getting stuck in a poor local minima [23]. As described later in Section III, instead of considering the distribution over the combinatorially large set \mathcal{V} , we introduce auxiliary distributions over the individual entries v_{ij} in V (at the cost of an additional constraint); thus, making the resulting optimization problem computationally tractable.

We observe that the proposed MEP-based framework performs as good as the recent trimmed lasso method on the unconstrained optimization problems [11], and outperforms the convex regularization based methods such as lasso, ridge

regression, LARS, and adaptive lasso. We demonstrate the frameworks flexibility in handling various practical constraints (as discussed above). We also illustrate and analyze the characteristic features of the MEP-based framework such as annealing and the phase transitions, and their utility towards increasing computational efficiency and determining the choice of sparsity level k in the SLR (1).

II. PROBLEM FORMULATION

As briefly stated in the Section I, we begin by re-writing the sparse vector \mathbf{w} as a product of a matrix V and a vector \mathbf{x} , i.e. $\mathbf{w} = V\mathbf{x}$, where the matrix V lies in the set

$$\mathcal{V} := \{V = (v_{ij}) \in \{0, 1\}^{d \times k} : \sum_i v_{ij} = 1 \forall j\} \quad (3)$$

and $\mathbf{x} \in \mathbb{R}^k$. Note that the number of columns in the matrix V , the size of the vector \mathbf{x} , binary entries in V , and column-stochasticity of V (i.e., $\sum_i v_{ij} = 1 \forall j$) ensure that maximum number of non-zero elements in $\mathbf{w} = V\mathbf{x}$ are exactly k . For instance, let $V \in \mathcal{V}$ be such that $v_{rs} = 1$, then the r -th position in \mathbf{w} is non-zero and is taken up by the s -th entry of \mathbf{x} . Further, the column stochasticity of V ensures that the s -th entry of $\mathbf{x} \in \mathbb{R}^k$ does not appear at any other location in \mathbf{w} — thereby, guaranteeing k non-zero values in \mathbf{w} . The fact that the vector \mathbf{x} lies in \mathbb{R}^k , and that the matrix V lies in the set \mathcal{V} in (3) together are equivalent to the sparsity constraint $\|\mathbf{w}\|_0 \leq k$. Thus, we re-write the sparse linear regression (SLR) problem in (1) as

$$\min_{\mathbf{x} \in \mathbb{R}^k, V \in \mathcal{V}} \|\mathbf{y} - AV\mathbf{x}\|_2^2. \quad (4)$$

III. MEP-BASED FRAMEWORK FOR PROBLEM SOLUTION

To make the optimization problem in (4) amenable to an MEP-based framework, we reformulate it as

$$\min_{\substack{\mathbf{x} \in \mathbb{R}^k \\ \{\eta(V|\mathbf{x})\}}} \sum_{V \in \mathcal{V}} \eta(V|\mathbf{x}) \|\mathbf{y} - AV\mathbf{x}\|_2^2 \quad (5a)$$

$$\text{subject to } \eta(V|\mathbf{x}) \in \{0, 1\}, \text{ and } \sum_{V \in \mathcal{V}} \eta(V|\mathbf{x}) = 1, \quad (5b)$$

where $\eta(V|\mathbf{x})$ is an auxiliary binary decision variable that determines the matrix V . We then replace $\eta(V|\mathbf{x}) \in \{0, 1\}$ by the soft decision variable $p(V|\mathbf{x}) \in [0, 1]$, resulting into a relaxed regression cost

$$D := \sum_{V \in \mathcal{V}} p(V|\mathbf{x}) \|\mathbf{y} - AV\mathbf{x}\|_2^2. \quad (6)$$

Note that $p(\cdot|\mathbf{x})$ can also be interpreted as the discrete distribution over the space of all the matrices $V \in \mathcal{V}$ given \mathbf{x} . We use MEP to design this distribution $\{p(V|\mathbf{x})\}$ as well as to determine the vector \mathbf{x} . In particular, their design is based on the principle of maximizing the Shannon entropy H subject to the constraint that the expected cost function D in (6) attains a pre-determined value c_0 . Maximizing H results

into most unbiased estimate of $\{p(V|\mathbf{x})\}$. The corresponding optimization problem is

$$\max_{\substack{\{p(V|\mathbf{x})\} \\ \mathbf{x} \in \mathbb{R}^k}} H := - \sum_V p(V|\mathbf{x}) \log(p(V|\mathbf{x})) \quad (7a)$$

$$\text{subject to } D := \sum_V p(V|\mathbf{x}) \|\mathbf{y} - AV\mathbf{x}\|_2^2 = c_0, \quad (7b)$$

$$\sum_V p(V|\mathbf{x}) = 1, \quad V \in \mathcal{V}. \quad (7c)$$

Since $|\mathcal{V}| = d^k$, the resulting decision variable space $\{p(V|\mathbf{x})\}$ is exponentially large; thus, making the optimization problem (7) intractable in its current form. We trim down the decision variable space to polynomial order by dissociating the decision variable $p(V|\mathbf{x})$ as

$$p(V|\mathbf{x}) = \prod_{i,j=1}^{d,k} p_{ij}(v_{ij}|\mathbf{x}), \quad (8)$$

where $p_{ij}(\cdot|\mathbf{x})$ is distribution over all possible values $v_{ij} \in \{0, 1\}$, the ij -th entry in V , takes. The new decision variable space $\{\{p_{ij}(v_{ij}|\mathbf{x})\}, x\}$ is now of the polynomial order $\mathcal{O}(dk)$, which takes us closer to posing the optimization problem (7) in a computationally tractable way. A possible downside of the dissociation is increase in non-linearity, which we explore more in the ongoing research. Substituting (8) in the objective (7a) we obtain

$$\mathcal{H} := \mathbf{1}_d^\top [Q \circ \log Q + \bar{Q} \circ \log \bar{Q}] \mathbf{1}_k, \quad (9)$$

where \circ denotes element wise operation, \log is also element wise, $Q \in [0, 1]^{d \times k}$ and $\bar{Q} \in [0, 1]^{d \times k}$ are defined as

$$Q := (q_{ij}), \quad q_{ij} := p_{ij}(v_{ij} = 1|\mathbf{x}), \quad \bar{Q} = \mathbf{1}_{d \times k} - Q. \quad (10)$$

Please see Appendix of [24] for details on the above (as well as the following) algebraic simplifications. The constraint (7b), in terms of Q and \bar{Q} , transforms into $\mathcal{D} :=$

$$\|\mathbf{y} - AQ\mathbf{x}\|_2^2 + [a_1^\top a_1 \dots a_d^\top a_d][Q \circ \bar{Q}](\mathbf{x} \circ \mathbf{x}) = c_0, \quad (11)$$

and the constraint (7c), which ensures that only one V is selected from the set \mathcal{V} and that V is a column stochastic matrix, is taken care by the fact that we define $\bar{Q} := (\bar{q}_{ij})$ (i.e., $p_{ij}(v_{ij} = 0|x)$) as $1 - Q$ (where $q_{ij} = p_{ij}(v_{ij} = 1|x)$), and that $Q^\top \mathbf{1}_d = \mathbf{1}_k$, i.e., Q is also column stochastic matrix. More precisely, the reformulation of the optimization problem (7) in terms of the tractable decision variables is

$$\max_{\mathbf{x} \in \mathbb{R}^k, Q \in [0,1]^{d \times k}} \mathcal{H} \text{ subject to } \mathcal{D} = c_0, \quad Q^\top \mathbf{1}_d = \mathbf{1}_k. \quad (12)$$

We consider the following augmented Lagrangian corresponding to the optimization problem (12) is $F_T =$

$$\mathcal{H} - \frac{1}{T}(\mathcal{D} - c_0) - \mu^\top (Q^\top \mathbf{1}_d - \mathbf{1}_k) - \frac{1}{2}\rho \|Q^\top \mathbf{1}_d - \mathbf{1}_k\|_2^2, \quad (13)$$

where T and μ denote the Lagrange multipliers corresponding to the constraints in (12), $\|Q^\top \mathbf{1}_d - \mathbf{1}_k\|_2^2$ denotes the penalty term, and ρ denotes the penalty parameter. Due to its close analogy to the MEP-based framework illustrated in

Algorithm 1 Maximum Entropy Sparsity-enforcing Regularization for Linear Regression

- 1: **Input:** $T_{\min}, T_{\max}, \beta < 1$;
 - 2: **Output:** Q and \mathbf{x} .
 - 3: **Initialize:** $T_0 = T_{\max}, Q_0 = [q_1, \dots, q_k], q_j = \frac{1}{d} \mathbf{1}_d, t = 1, \rho_0 = \frac{1}{T_0}, \mu_0 = \mathbf{0}, x_0$ in 15a.
 - 4: **while** $T_t \geq T_{\min}$ **do**
 - 5: **Obtain** Q_t, x_t : Minimize F_T in (13) using a descent method and initial value Q_{t-1}, x_{t-1} .
 - 6: $T_t \leftarrow \beta T_t, t \leftarrow t + 1$, update μ_t in (14), $\rho_t = \frac{1}{T_t}$.
 - 7: **end while**
-

[25], we refer to T as the temperature, and F_T as the free-energy term. It is known from sensitivity analysis [23] that a large value of the Lagrange parameter T corresponds to a large value of c_0 . Similarly, a small value of T corresponds to a small value of c_0 . In our framework, we repeatedly solve (12) at decreasing values of c_0 by maximizing the Lagrangian F_T at iteratively decreasing values of T .

More precisely, let T_k be the temperature value at the k -th iteration of the algorithm, the penalty parameter ρ_k be equal to T_k , and the multiplier μ_k be given by the iteration

$$\mu_k = \mu_{k-1} + \rho_k (Q^\top \mathbf{1}_d - \mathbf{1}_k), \quad (14)$$

where $\mu_0 = 0$. We vary T_k from a large value ($\rightarrow \infty$) to a small value (≈ 0). At large values of T_k , the Lagrangian F_{T_k} is dominated by the convex entropy function \mathcal{H} and the penalty parameter. As T_k becomes small, the other terms including the non-convex \mathcal{D} gets more weightage. As in other MEP-based frameworks, it is this homotopy from a convex function to the non-convex cost function \mathcal{D} that helps the algorithm avoid getting stuck in a poor local minima. It is also known that given $T_{k+1} \geq T_k$, the above iterations converge to a local minima of the optimization problem (12) (see [26] for details). Please see Algorithm 1 for details on implementation.

Convergence of Q to a binary matrix : As our objective is to solve the optimization problem (5), we want that as $T \rightarrow 0$, the matrix Q converges to a binary matrix, i.e., $Q \rightarrow \{0, 1\}^{d \times k}$. This in turn enables $p(V|\mathbf{x}) \rightarrow \{0, 1\}$, i.e., the soft decision variables converge to the binary solution required in (5). The structure of the optimal Q is amenable to this desired aspect of the Algorithm 1. More precisely, by setting $\frac{\partial F_T}{\partial \mathbf{x}} = 0$ and $\frac{\partial F_T}{\partial Q} = 0$, we obtain

$$\mathbf{x} = \left[Q^\top A^\top A Q + \text{diag}[\lambda^\top (Q \circ \bar{Q})] \right]^{-1} Q^\top A^\top \mathbf{y}, \quad (15a)$$

$$Q = \frac{\exp(\frac{2}{T} H_m) \circ}{\exp(\frac{2}{T} H_m) + \exp(-\frac{2}{T} H_p)}, \quad (15b)$$

where \circ is elementwise operation, \exp is elementwise

$$H_m = \min \{ \Xi, \mathbf{0} \}, \quad H_p = \max \{ \Xi, \mathbf{0} \}, \quad (16a)$$

$$\Xi = A^\top (\mathbf{y} - AQ\mathbf{x})\mathbf{x}^\top - \frac{1}{2} \lambda_a (\mathbf{x} \circ \mathbf{x})^\top \circ (1 - 2Q) \quad (16b)$$

$$- \frac{1}{2} \mathbf{1}_d \mu^\top - \frac{1}{2T} \mathbf{1}_d \mathbf{1}_d^\top Q + \frac{1}{2T} \mathbf{1}_d \mathbf{1}_k^\top, \quad (16c)$$




Fig. 1: Phase Transition. Plot of k_d versus T , where k_d are the distinct number of non-zero values in $\mathbf{x} \in \mathbb{R}^k$ (or, equivalently the distinct columns in $Q \in [0, 1]^{d \times k}$).

and $\lambda_a = [a_1^\top a_1 \ \dots \ a_d^\top a_d]$. Note that Q in (15b) resembles a Gibb's distribution, whose entries are identical at large values of T ($\rightarrow \infty$), and converge to either 0 or 1 as $T \rightarrow 0$ (with Ξ bounded); thus, achieving the above objective.

IV. FLEXIBILITY IN MODELING CONSTRAINTS

As briefly discussed in the Section I, our proposed framework explicitly allows the control over the selection of the features in the design matrix A . More precisely, the j -th feature a_j is selected if and only if the sum of j -th row in the binary matrix V is non-zero. More precisely, if the j -th feature is selected, then $\sum_{t=1}^k v_{jt} \geq 1$. This attribute of the framework allows us to conveniently model several structural constraints in the design of the sparse vector \mathbf{w} , which otherwise are difficult to model in the existing literature when explicit control over the selection of the feature is not possible. Below we elucidate some of these scenarios.

1- Correlated feature vectors: The columns of a given design matrix A may be linearly dependent on each other, where the extent of their linear dependence is measured by the Pearson correlation coefficient [27]. For highly correlated set of features, it is desirable to have only one of the features to be selected, i.e., to have a non-zero entry in the sparse vector \mathbf{w} only for one of such features, and to have zero values in \mathbf{w} corresponding to all other features in this set. One straightforward solution is to drop all the features in this set, except the one that highly correlates with the output \mathbf{y} . However, such a methodology is sub-optimal. On the other hand, our proposed framework explicitly models this constraint in terms of the matrix V . In particular, let the r features in the set $\{a_{l_1}, a_{l_2}, \dots, a_{l_r}\}$ be highly correlated. Then the constraint $\sum_{t=1}^m v_{l_1 t} + v_{l_2 t} + \dots + v_{l_r t} \leq 1$, enforces that at most one of the above r -features is picked. To incorporate this in the optimization problem (12), we begin by re-writing it as $\sum_{V \in \mathcal{V}} \eta(V|x) (\sum_{t=1}^k v_{l_1 t} + v_{l_2 t} + \dots + v_{l_r t}) \leq 1$. As done before, we replace the above $\eta(V|x) \in \{0, 1\}$ with soft weights $p(V|x) \in [0, 1]$, and dissociate them as in (8). Subsequently, the algebraic manipulations (similar to that in

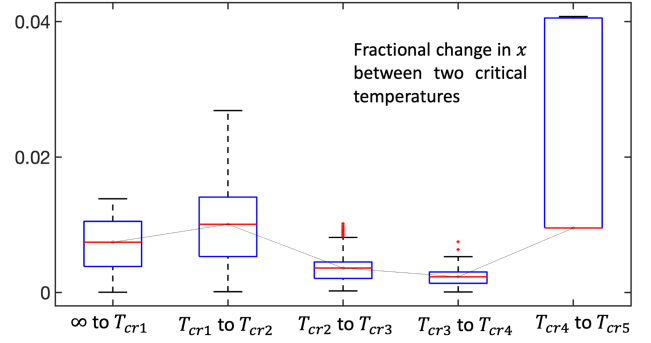


Fig. 2: Plot of fractional change in x in between two consecutive critical temperatures.

(9) and (11)) result into a constraint in $Q = (q_{ij})$ as

$$\sum_{t=1}^k q_{l_1 t} + q_{l_2 t} + \dots + q_{l_r t} \leq 1. \quad (17)$$

The above constraint can be incorporated easily in (12).

2- A priori knowledge: Expert insights play a substantial role in shaping the sparse solution \mathbf{w} , often suggesting the inclusion of at least one feature from each group when working with multiple groups of features. For instance, in the context of medical diagnosis, it is recommended to utilize data from distinct diagnostic groups, such as radiological imaging, clinical laboratory tests, and patient medical history, to ensure a comprehensive evaluation of a patient's condition. Our framework allows us to incorporate such a priori information into the existing problem. To elaborate, suppose there are r features in the set $\{a_{l_1}, a_{l_2}, \dots, a_{l_r}\}$ originating from the same group of features. In this context, the constraint $\sum_{t=1}^k v_{l_1 t} + v_{l_2 t} + \dots + v_{l_r t} \geq 1$ serves to ensure that at least one of the r features is selected. Similar to the previous scenario, it can be expressed in terms of the Q as follows $\sum_{t=1}^k q_{l_1 t} + q_{l_2 t} + \dots + q_{l_r t} \geq 1$.

3- Grouping constraints: Here a group of features needs to be selected as a single unit. Algorithms such as group lasso addresses such instances by introducing a regularization term $\mathcal{T}(\mathbf{x})$ to the cost function in (1). In the proposed framework, such group constraints can be easily modeled. For instance, let $\{a_{l_1}, a_{l_2}, \dots, a_{l_r}\}$ be a group of features that occurs as a single unit, i.e., if one of them is picked then all of them should be picked. This constraint is modeled as $v_{l_1 t} = v_{l_2 t} = \dots = v_{l_r t}$. As above, this constraint in terms of the decision variable Q is given by $q_{l_1 t} = q_{l_2 t} = \dots = q_{l_r t}$.

Note that the above constraints can be easily incorporated into (12), and the resulting optimization problem can be addressed using existing methods such as interior points algorithm and various other penalty methods [26].

Phase Transition: Algorithm 1 is characterized by a unique trait wherein at large values of temperatures T all non-zero values in \mathbf{x} are *identical*; equivalently, the columns in Q are identical (15a). As T decreases, there are specific instances at which the number of distinct non-zero values in \mathbf{x} increase (equivalently, distinct columns in Q increase). We refer to these instances as *phase transitions*, and the

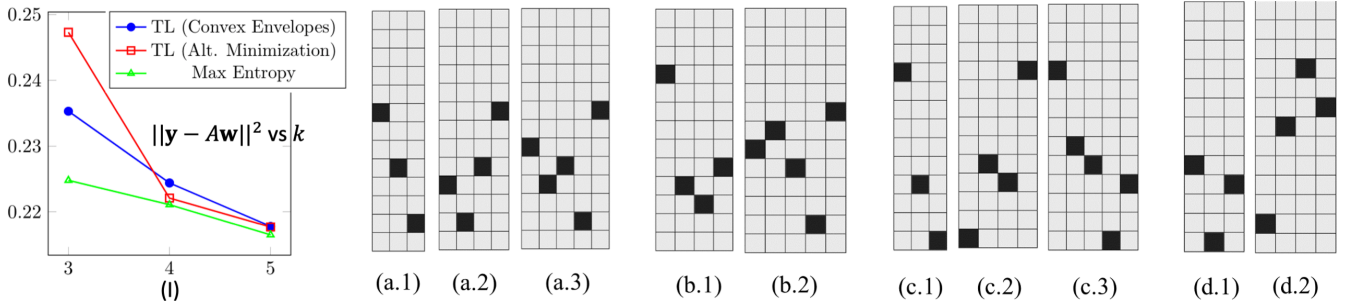


Fig. 3: (I) The figure illustrates the cost values associated with solutions obtained using the Alternating Minimization and Convex Envelopes heuristics of Trimmed Lasso, alongside our proposed Maximum-Entropy methodology. (a)-(d) The figure exhibits the obtained V matrices for different k values (3, 4, and 5), encompassing various scenarios: unconstrained (a.1, a.2, a.3), correlated features removal (b.1, b.2), a priori knowledge imposition (c.1, c.2, c.3), and grouping constraints (d.1, d.2). In this representation, the color black denotes that the element has a value of one, while gray is used to indicate a value of zero.

associated temperature values at which they occur as critical temperatures T_{cr} . Figure 1 illustrates this phenomenon on randomly generated data $\mathbf{y} \in \mathbb{R}^8$, $A \in \mathbb{R}^{8 \times 15}$, $\mathbf{w} \in \mathbb{R}^{15}$, and $\|\mathbf{w}\|_0 = 3$. For the purpose of illustration we set $k = 5$ (though the *true* sparsity is 3, it is not known a priori in general). The Algorithm 1 begins with $k_d = 1$ distinct non-zero value in $\mathbf{x} \in \mathbb{R}^5$ at high temperatures (i.e., low values of $\log(1/T)$). As T decreases, the number of distinct values remain unchanged for sometime before a critical temperature is reached, where $\mathbf{x} \in \mathbb{R}^5$ contains $k_d = 2$ distinct values. This process continues, till the Algorithm 1 determines $k_d = 5$ distinct values in $\mathbf{x} \in \mathbb{R}^5$. See [24] for computation of T_{cr} .

Annealing schedule: Phase transition plays a key role in designing the annealing schedule for the temperature T in the Algorithm 1. We observe that between two consecutive critical temperatures the change in the vector of non-zero values \mathbf{x} as determined by the Algorithm 1 is small. Let $\Delta \mathbf{x}(n) = \|\mathbf{x}(n) - \mathbf{x}_{\text{mean}}\| / \|\mathbf{x}_{\text{mean}}\|$ denote the fractional change in $\mathbf{x}(n)$ — the non-zero values determined by the Algorithm 1 at temperature T_n — where \mathbf{x}_{mean} is the average of all the non-zero value vectors determined by the Algorithm 1 in between the two consecutive critical temperatures $T_{cr_{n_1}}$ and $T_{cr_{n_2}}$ such that $T_{cr_{n_1}} < T_n < T_{cr_{n_2}}$. Figure 2 illustrates the boxplot of this fractional change in $\Delta \mathbf{x}(n)$ observed in between two consecutive T_{cr} 's for the example considered earlier in this section. Note that $\Delta \mathbf{x}(n)$ is quite small and the median roughly lies between 0% to 1%; whereas the change observed at the phase transition is considerable, as it adds a distinct non-zero value in \mathbf{x} (see Figure 1). In particular, in Figure 2, we observe a change of 84% at the first critical temperature T_{cr_1} , 32% at T_{cr_2} , 8% at T_{cr_3} and 9% at T_{cr_4} in the non-zero values \mathbf{x} determined by the Algorithm 1. Thus, the solution given by the algorithm undergoes a drastic change only at T_{cr} 's and largely remains unchanged in between any two consecutive T_{cr} 's. The above characteristic of our MEP-based algorithm is significant to determining an appropriate schedule for the annealing parameter T . In particular, it motivates a geometrical annealing of temperature T as in other MEP-based frameworks (see [21] for details).

V. SIMULATIONS AND RESULTS

In this section, we apply our methodology to a dataset containing 205 data points related to automobile features, as referenced in [28]. Notably, 195 of these records have complete information, and we carefully select 13 continuous features for sparse linear regression, with the automobile price as the model's output. These features encompass a wide range of attributes, including car dimensions, weight, engine specifications, and fuel efficiency. To enhance the quality of results, we normalize the columns to have a 2-norm of 1.

Our primary aim is to develop a predictive model for automobile prices, with an emphasis on sparsity. This involves the selection of a small-sized subset of these 13 features, accompanied by their respective coefficients. In essence, our goal is to determine the values of the matrix V and vector \mathbf{x} in (4) while maintaining a predefined sparsity level indicated by k . Here, V represents the selected features, and \mathbf{x} signifies their corresponding coefficients. For the unconstrained scenario, we consider three instances of sparsity $k \in \{3, 4, 5\}$. The resultant V matrices are visually represented in Figure 3 (a.1, a.2, and a.3). As illustrated in Section II, a feature a_j is selected if the j -th row of V sums up to a value ≥ 1 . Thus, as illustrated in the Figure, the selected features correspond to columns $\{a_6, a_9, a_{12}\}$, $\{a_6, a_9, a_{10}, a_{12}\}$, and $\{a_6, a_8, a_9, a_{10}, a_{12}\}$, respectively for the above three sparsity levels. The corresponding cost function values for these degrees of sparsity are 0.2248, 0.2211, and 0.2165.

To assess our method's accuracy and establish a benchmark, we conduct a comparative analysis, evaluating our results against Trimmed Lasso, as presented in [15], using both the alternating minimization and convex envelopes heuristics. This analysis employs the same dataset and consistent sparsity levels, ensuring a fair comparison. Trimmed Lasso has consistently demonstrated superior performance compared to various other Lasso variants, making it an ideal point of comparison for our method.

The cost values associated with the solutions are graphically illustrated in Figure 3(I). Notably, across this dataset and for three distinct sparsity levels (k values of 3, 4, and 5), our maximum entropy approach and both Trimmed Lasso heuristics exhibit quite similar performance. It's worth

mentioning that our method holds a slight advantage in terms of accuracy. The features selected by both heuristics of Trimmed Lasso correspond to columns $\{a_6, a_{10}, a_{12}\}$, $\{a_6, a_9, a_{10}, a_{12}\}$, and $\{a_6, a_8, a_9, a_{10}, a_{12}\}$ respectively, for the three mentioned sparsity levels.

A distinctive advantage of our approach, setting it apart from Trimmed Lasso, is its capacity to integrate diverse constraints, as discussed in Section IV and simulated below.

1- Correlated feature vectors: We classify features as correlated if their absolute correlation coefficient exceeds 0.8. This criterion identifies the following sets of correlated features: $\{a_1, a_2, a_4\}$, $\{a_1, a_5\}$, $\{a_2, a_5\}$, $\{a_4, a_6\}$, $\{a_{10}, a_{12}, a_{13}\}$, $\{a_{10}, a_6\}$, and $\{a_{13}, a_4\}$. As a result, it is advisable to choose, at most, one feature from each correlated set. As shown in Figure 3, for both $k = 4$ and $k = 5$ (a.2 and a.3), the unconstrained solution includes correlated features $\{a_6, a_{10}\}$ and $\{a_{10}, a_{12}\}$. Upon imposing the constraint (17), the solution changes to selecting $\{a_4, a_9, a_{10}, a_{11}\}$ for $k = 4$ and $\{a_6, a_7, a_8, a_9, a_{12}\}$ for $k = 5$, as demonstrated in Figure 3 (b.1 and b.2). The corresponding cost values are 0.2538 and 0.2214, respectively, which are (naturally) a bit larger than the unconstrained scenario illustrated above.

2- A priori knowledge: Suppose we possess prior knowledge indicating that from various feature groups, we must include at least one feature. For example, in our dataset, it is essential to select at least one feature related to the vehicle's size, one related to its engine, and one associated with its fuel efficiency. In practical terms, this constraint implies that within each group of columns $\{a_1, a_2, a_3, a_4\}$, $\{a_6, a_7, a_8, a_9, a_{10}, a_{11}\}$, and $\{a_{12}, a_{13}\}$, a minimum of one column must be selected. The features selected under this constraint are depicted in Figure 3 (c.1, c.2, and c.3), which correspond to $\{a_4, a_{10}, a_{13}\}$, $\{a_4, a_9, a_{10}, a_{13}\}$, and $\{a_4, a_8, a_9, a_{10}, a_{13}\}$ for sparsity of 3, 4, and 5, respectively. The respective cost values are 0.2657, 0.2550, and 0.2223.

3- Grouping constraints: Solely for demonstrative purposes, we have assumed that columns $\{a_6, a_7\}$ and $\{a_9, a_{10}\}$ are treated as unified groups. In other words, it's an all-or-nothing selection within each group. Illustrated in Figure 3 (a.1 and a.2), the initial unconstrained solution contradicts the imposed constraint. After incorporating the grouping constraint, the newly selected features for $k = 3$ and $k = 4$ are $\{a_9, a_{10}, a_{13}\}$ and $\{a_4, a_6, a_7, a_{12}\}$, as depicted in Figure 3 (d.1 and d.2) with cost of 0.2655 and 0.2268, respectively.

REFERENCES

- [1] F. Wen, L. Chu, P. Liu, and R. C. Qiu, "A survey on nonconvex regularization-based sparse and low-rank recovery in signal processing, statistics, and machine learning," *IEEE Access*, vol. 6, pp. 69 883–69 906, 2018.
- [2] W. J. Welch, "Algorithmic complexity: three np-hard problems in computational statistics," *Journal of Statistical Computation and Simulation*, vol. 15, no. 1, pp. 17–25, 1982.
- [3] D. L. Donoho, Y. Tsaig, I. Drori, and J.-L. Starck, "Sparse solution of underdetermined systems of linear equations by stagewise orthogonal matching pursuit," *IEEE transactions on Information Theory*, vol. 58, no. 2, pp. 1094–1121, 2012.
- [4] P. Shekhar and A. Patra, "A forward-backward greedy approach for sparse multiscale learning," *Computer Methods in Applied Mechanics and Engineering*, vol. 400, p. 115420, 2022.
- [5] Z. Zhang, Y. Xu, J. Yang, X. Li, and D. Zhang, "A survey of sparse representation: algorithms and applications," *IEEE access*, vol. 3, pp. 490–530, 2015.
- [6] M. A. Figueiredo, R. D. Nowak, and S. J. Wright, "Gradient projection for sparse reconstruction: Application to compressed sensing and other inverse problems," *IEEE Journal of selected topics in signal processing*, vol. 1, no. 4, pp. 586–597, 2007.
- [7] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM journal on imaging sciences*, vol. 2, no. 1, pp. 183–202, 2009.
- [8] J. Yang and X. Yuan, "Linearized augmented lagrangian and alternating direction methods for nuclear norm minimization," *Mathematics of computation*, vol. 82, no. 281, pp. 301–329, 2013.
- [9] G. Davis, S. Mallat, and M. Avellaneda, "Adaptive greedy approximations," *Constructive approximation*, vol. 13, pp. 57–98, 1997.
- [10] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani, "Least angle regression," *The Annals of Statistics*, vol. 32, no. 2, pp. 407–451, 2004. [Online]. Available: <http://www.jstor.org/stable/3448465>
- [11] T. Amir, R. Basri, and B. Nadler, "The trimmed lasso: Sparse recovery guarantees and practical optimization by the generalized soft-min penalty," *SIAM journal on mathematics of data science*, vol. 3, no. 3, pp. 900–929, 2021.
- [12] J. Wang, "Non-convex lp regularization for sparse reconstruction of electrical impedance tomography," *Inverse Problems in Science and Engineering*, vol. 29, no. 7, pp. 1032–1053, 2021.
- [13] P. K. Pokala, R. V. Hemadri, and C. S. Seelamantula, "Iteratively reweighted minimax-concave penalty minimization for accurate low-rank plus sparse matrix decomposition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 12, pp. 8992–9010, 2022.
- [14] Y. Kim, H. Choi, and H.-S. Oh, "Smoothly clipped absolute deviation on high dimensions," *Journal of the American Statistical Association*, vol. 103, no. 484, pp. 1665–1673, 2008. [Online]. Available: <http://www.jstor.org/stable/27640214>
- [15] D. Bertsimas, M. S. Copenhaver, and R. Mazumder, "The trimmed lasso: Sparsity and robustness," *arXiv preprint arXiv:1708.04527*, 2017.
- [16] H. Wang and C. Leng, "A note on adaptive group lasso," *Computational statistics & data analysis*, vol. 52, no. 12, pp. 5277–5286, 2008.
- [17] L. Yuan, J. Liu, and J. Ye, "Efficient methods for overlapping group lasso," *Advances in neural information processing systems*, vol. 24, 2011.
- [18] C. A. Micchelli, J. M. Morales, and M. Pontil, "Regularizers for structured sparsity," *Advances in Computational Mathematics*, vol. 38, pp. 455–489, 2013.
- [19] L. Baldassarre, J. M. Morales, and M. Pontil, "Incorporating additional constraints in sparse estimation," *IFAC Proceedings Volumes*, vol. 45, no. 16, pp. 959–964, 2012.
- [20] L. Jacob, G. Obozinski, and J.-P. Vert, "Group lasso with overlap and graph lasso," in *Proceedings of the 26th annual international conference on machine learning*, 2009, pp. 433–440.
- [21] K. Rose, "Deterministic annealing for clustering, compression, classification, regression, and related optimization problems," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2210–2239, 1998.
- [22] Y. Xu, S. M. Salapaka, and C. L. Beck, "Aggregation of graph models and markov chains by deterministic annealing," *IEEE Transactions on Automatic Control*, vol. 59, no. 10, pp. 2807–2812, 2014.
- [23] A. Srivastava and S. M. Salapaka, "Parameterized mdps and reinforcement learning problems—a maximum entropy principle-based framework," *IEEE Transactions on Cybernetics*, 2021.
- [24] A. Srivastava, A. Bayati, and S. Salapaka, "Sparse linear regression with constraints: A flexible entropy-based framework," *arXiv preprint arXiv:2311.08342*, 2023.
- [25] K. Rose, "Deterministic annealing, clustering, and optimization," Ph.D. dissertation, California Institute of Technology, 1991.
- [26] D. Bertsekas, *Nonlinear Programming*, ser. Athena scientific optimization and computation series. Athena Scientific, 1999. [Online]. Available: <https://books.google.co.in/books?id=TgMpAQAAAJ>
- [27] I. Cohen, Y. Huang, J. Chen, J. Benesty, J. Benesty, J. Chen, Y. Huang, and I. Cohen, "Pearson correlation coefficient," *Noise reduction in speech processing*, pp. 1–4, 2009.
- [28] J. Schlimmer, "Automobile," UCI Machine Learning Repository, 1987, DOI: <https://doi.org/10.24432/C5B01C>.