

Learning task structure via sparsity grouped multitask learning

Meghana Kshirsagar¹, Eunho Yang², and Aurélie C. Lozano³

¹ **Memorial Sloan Kettering Cancer Center,
1275 York Ave., New York, NY
kshirsam@mskcc.org

² School of Computing,
Korea Advanced Inst. of Science and Tech., Daejeon, South Korea.
eunhoy@kaist.ac.kr

³ IBM T. J. Watson research
Yorktown Heights, New York, NY
aclozano@us.ibm.com

Abstract. Sparse mapping has been a key methodology in many high-dimensional scientific problems. When multiple tasks share the set of relevant features, learning them jointly in a group drastically improves the quality of relevant feature selection. However, in practice this technique is used limitedly since such grouping information is usually hidden. In this paper, our goal is to recover the group structure on the sparsity patterns and leverage that information in the sparse learning. Toward this, we formulate a joint optimization problem in the task parameter and the group membership, by constructing an appropriate regularizer to encourage sparse learning as well as correct recovery of task groups. We further demonstrate that our proposed method recovers groups and the sparsity patterns in the task parameters accurately by extensive experiments.

1 Introduction

Humans acquire knowledge and skills by categorizing the various problems/tasks encountered, recognizing how the tasks are related to each other and taking advantage of this organization when learning a new task. Statistical machine learning methods also benefit from exploiting such similarities in learning related problems. Multitask learning (MTL) (Caruana, 1997) is a paradigm of machine learning, encompassing learning algorithms that can share information among related tasks and help to perform those tasks together more efficiently than in isolation. These algorithms exploit task relatedness by various mechanisms. Some works enforce that parameters of various tasks are close to each other in some geometric sense (Evgeniou & Pontil, 2004; Maurer, 2006). Several works leverage the existence of a shared low dimensional subspace (Argyriou et al., 2008; Liu et al., 2009; Jalali et al., 2010; Chen et al., 2012) or manifold (Agarwal et al., 2010) that contains the task parameters. Some bayesian MTL methods assume the same prior on parameters of related tasks (Yu et al., 2005; Daumé III, 2009), while neural networks based methods share some hidden units (Baxter, 2000).

** This work was done while MK and EY were at IBM T.J. Watson research

A key drawback of most MTL methods is that they assume that all tasks are equally related. Intuitively, learning unrelated tasks jointly may result in poor predictive models; i.e tasks should be coupled based on their relatedness. While the coupling of task parameters can sometimes be controlled via hyper-parameters, this is infeasible when learning several hundreds of tasks. Often, knowing the *task relationships* themselves is of interest to the application at hand. While these relationships might sometimes be derived from domain specific intuition (Kim & Xing, 2010; Widmer et al., 2010; Rao et al., 2013), they are either not always known a priori or are pre-defined based on the knowledge of $P(X)$ rather than $P(Y|X)$. We aim to automatically learn these task relationships, while simultaneously learning individual task parameters. This idea of jointly learning task groups and parameters has been explored in prior works. For instance Argyriou et al. (2008) learn a set of kernels, one per group of tasks and Jacob et al. (2009) cluster tasks based on similarity of task parameters. Others (Zhang & Yeung, 2010; Gong et al., 2012) try to identify “outlier” tasks. Kumar & Daume (2012); Kang et al. (2011) assume that task parameters within a group lie in a shared low dimensional subspace. Zhang & Schneider (2010) use a matrix-normal regularization to capture task covariance and feature covariance between tasks and enforce sparsity on these covariance parameters and Fei & Huan (2013) use a similar objective with a structured regularizer. Their approach is however, not suited for high dimensional settings and they do not enforce any sparsity constraints on the task parameters matrix W . A Bayesian approach is proposed in Passos et al. (2012), where parameters are assumed to come from a nonparametric mixture of nonparametric factor analyzers.

Here, we explore the notion of *shared sparsity* as the structure connecting a group of related tasks. More concretely, we assume that tasks in a group all have similar relevant features or analogously, the same zeros in their parameter vectors. Sparsity inducing norms such as the ℓ_1 norm capture the principle of parsimony, which is important to many real-world applications, and have enabled efficient learning in settings with high dimensional feature spaces and few examples, via algorithms like the Lasso (Tibshirani, 1996). When confronted by several tasks where sparsity is required, one modeling choice is for each task to have its’ own sparse parameter vector. At the other extreme is the possibility of enforcing shared sparsity on all tasks via a structured sparsity inducing norm such as ℓ_1/ℓ_2 on the task parameter matrix: $\|W\|_{1,2}$ (Bach et al., 2011).⁴ We choose to enforce sparsity at a group level by penalizing $\|W_g\|_{1,2}$, where $\|W_g\|$ is the parameter matrix for all tasks in group g , while learning group memberships of tasks.

To see why this structure is interesting and relevant, consider the problem of transcription factor (TF) binding prediction. TFs are proteins that bind to the DNA to regulate expression of nearby genes. The binding specificity of a TF to an arbitrary location on the DNA depends on the pattern/sequence of nucleic acids (A/C/G/T) at that location. These sequence preferences of TFs have some similarities among related TFs. Consider the task of predicting TF binding, given segments of DNA sequence (these are the examples), on which we have derived features such as n -grams (called k -mers)⁵. The feature space is very high dimensional and a small set of features typically cap-

⁴ Note: this cross-task structured sparsity is different from the Group Lasso (Yuan & Lin, 2006), which groups covariates within a task ($\min_{w \in \mathbb{R}^d} \sum_g \|w_g\|$, where w_g is a group of parameters)

⁵ e.g.: GTAATTNC is an 8-mer (‘N’ represents a wild card)

ture the binding pattern for a single TF. Given several tasks, each representing one TF, one can see that the structure of the ideal parameter matrix is likely to be group sparse, where TFs in a group have similar binding patterns (i.e similar important features but with different weights). The applicability of task-group based sparsity is not limited to isolated applications, but desirable in problems involving billions of features, as is the case with web-scale information retrieval and in settings with few samples such as genome wide association studies involving millions of genetic markers over a few hundred patients, where only a few markers are relevant.

The main contributions of this work are:

- We present a new approach towards learning task group structure in a multitask learning setting that simultaneously learns both the task parameters W and a clustering over the tasks.
- We define a regularizer that divides the set of tasks into groups such that all tasks within a group share the same sparsity structure. Though the ideal regularizer is discrete, we propose a relaxed version and we carefully make many choices that lead to a feasible alternating minimization based optimization strategy. We find that several alternate formulations result in substantially worse solutions.
- We evaluate our method through experiments on synthetic datasets and two interesting real-world problem settings. The first is a regression problem: QSAR, quantitative structure activity relationship prediction (see (Ma et al., 2015) for an overview) and the second is a classification problem important in the area of regulatory genomics: transcription factor binding prediction (described above). On synthetic data with known group structure, our method recovers the correct structure. On real data, we perform better than prior MTL group learning baselines.

1.1 Relation to prior work

Our work is most closely related to Kang et al. (2011), who assume that each group of tasks shares a latent subspace. They find groups so that $\|W_g\|_*$ for each group g is small, thereby enforcing sparsity in a *transformed* feature space. Another approach, GO-MTL (Kumar & Daume, 2012) is based on the same idea, with the exception that the latent subspace is shared among all tasks, and a low-rank decomposition of the parameter matrix $W = LS$ is learned. Subsequently, the coefficients matrix S is clustered to obtain a grouping of tasks. Note that, learning group memberships is not the goal of their approach, but rather a post-processing step upon learning their model parameters.

To understand the distinction from prior work, consider the weight matrix W^* in Figure 4(a), which represents the true group sparsity structure that we wish to learn. While each task group has a low-rank structure (since s of the d features are non-zero, the rank of any W_g is bounded by s), it has an additional property that $(d - s)$ features are zero or irrelevant for all tasks in this group. Our method is able to exploit this additional information to learn the correct sparsity pattern in the groups, while that of Kang et al. (2011) is unable to do so, as illustrated on this synthetic dataset in Figure 5 (details of this dataset are in Sec 5.1). Though Kang et al. (2011) recovers some of the block diagonal structure of W , there are many non-zero features which lead to an incorrect group structure. We present a further discussion on how our method is sample efficient as compared to Kang et al. (2011) for this structure of W in Sec 3.1.

We next present the setup and notation, and lead to our approach by starting with a straight-forward combinatorial objective and then make changes to it in multiple steps (Sec 2-4). At each step we explain the behaviour of the function to motivate the particular choices we made; present a high-level analysis of sample complexity for our method and competing methods. Finally we show experiments (Sec 5) on four datasets.

2 Setup and Motivation

We consider the standard setting of multi-task learning in particular where tasks in the same group share the sparsity patterns on parameters. Let $\{T_1, \dots, T_m\}$ be the set of m tasks with training data \mathcal{D}_t ($t = 1 \dots m$). Let the parameter vectors corresponding to each of the m tasks be $\mathbf{w}^{(1)}, \mathbf{w}^{(2)}, \dots, \mathbf{w}^{(m)} \in \mathbb{R}^d$, d is the number of covariates/features. Let $\mathcal{L}(\cdot)$ be the loss function which, given \mathcal{D}_t and $\mathbf{w}^{(t)}$ measures deviation of the predictions from the response. Our goal is to learn the task parameters where i) each $\mathbf{w}^{(t)}$ is assumed to be sparse so that the response of the task can be succinctly explained by a small set of features, and moreover ii) there is a partition $\mathcal{G}^* := \{G_1, G_2, \dots, G_N\}$ over tasks such that all tasks in the same group G_i have the same sparsity patterns. Here N is the total number of groups learned. If we learn every task independently, we solve m independent optimization problems:

$$\underset{\mathbf{w}^{(t)} \in \mathbb{R}^d}{\text{minimize}} \mathcal{L}(\mathbf{w}^{(t)}; \mathcal{D}_t) + \lambda \|\mathbf{w}^{(t)}\|_1$$

where $\|\mathbf{w}^{(t)}\|_1$ encourages sparse estimation with regularization parameter λ . However, if \mathcal{G}^* is given, jointly estimating all parameters together using a group regularizer (such as ℓ_1/ℓ_2 norm), is known to be more effective. This approach requires fewer samples to recover the sparsity patterns by sharing information across tasks in a group:

$$\underset{\mathbf{w}^{(1)}, \dots, \mathbf{w}^{(m)}}{\text{minimize}} \sum_{t=1}^m \mathcal{L}(\mathbf{w}^{(t)}; \mathcal{D}_t) + \sum_{g \in \mathcal{G}^*} \lambda_g \|\mathbf{W}_g\|_{1,2} \quad (1)$$

where $\mathbf{W}_g \in \mathbb{R}^{d \times |g|}$, where $|g|$ is the number of tasks in the group g and $\|\cdot\|_{1,2}$ is the sum of ℓ_2 norms computed over row vectors. Say $t_1, t_2 \dots t_k$ belong to group g , then $\|\mathbf{W}_g\|_{1,2} := \sum_{j=1}^d \sqrt{(w_j^{(t_1)})^2 + (w_j^{(t_2)})^2 + \dots + (w_j^{(t_k)})^2}$. Here $w_j^{(t)}$ is the j -th entry of vector $\mathbf{w}^{(t)}$. Note that here we use ℓ_2 norm for grouping, but any ℓ_α norm $\alpha \geq 2$ is known to be effective.

We introduce a membership parameter $u_{g,t}$: $u_{g,t} = 1$ if task T_t is in a group g and 0 otherwise. Since we are only allowing a hard membership without overlapping (though this assumption can be relaxed in future work), we should have exactly one active membership parameter for each task: $u_{g,t} = 1$ for some $g \in \mathcal{G}$ and $u_{g',t} = 0$ for all other $g' \in \mathcal{G} \setminus \{g\}$. For notational simplicity, we represent the group membership parameters for a group g in the form of a matrix \mathbf{U}_g . This is a diagonal matrix where $\mathbf{U}_g := \text{diag}(u_{g,1}, u_{g,2}, \dots, u_{g,m}) \in \{0, 1\}^{m \times m}$. In other words, $[\mathbf{U}_g]_{ii} = u_{g,i} = 1$ if task T_i is in group g and 0 otherwise. Now, incorporating \mathbf{U} in (1), we can derive the




Fig. 1. Toy examples with two fixed parameter vectors: (Left) $\mathbf{w}^{(1)} = (1, 0, 0)^\top$, $\mathbf{w}^{(2)} = (a, 0, 0)^\top$, and (Right) $\mathbf{w}^{(1)} = (1/2, 1, 0)^\top$, $\mathbf{w}^{(2)} = (0, 1, a)^\top$, where we only vary one coordinate a fixing all others to visualize the norm values in 2-d space. Functions show the group norms, $\sum_g \|\mathbf{W}\mathbf{U}_g\|_{1,2}$ in (2) where two tasks belong to a single group (solid) or to separate groups (dotted). In both cases, this group regularizer favors the case with a single group.

optimization problem for learning the task parameters $\{\mathbf{w}^{(t)}\}_{t=1,\dots,m}$ and \mathbf{U} simultaneously as follows:

$$\begin{aligned} & \underset{\mathbf{W}, \mathbf{U}}{\text{minimize}} \quad \sum_{t=1}^m \mathcal{L}(\mathbf{w}^{(t)}; \mathcal{D}_t) + \sum_{g \in \mathcal{G}} \lambda_g \|\mathbf{W}\mathbf{U}_g\|_{1,2} \\ & \text{s. t.} \quad \sum_{g \in \mathcal{G}} \mathbf{U}_g = \mathbf{I}^{m \times m}, \quad [\mathbf{U}_g]_{ii} \in \{0, 1\}. \end{aligned} \quad (2)$$

where $\mathbf{W} \in \mathbb{R}^{d \times m} := [\mathbf{w}^{(1)}, \mathbf{w}^{(2)}, \dots, \mathbf{w}^{(m)}]$. Here $\mathbf{I}^{m \times m}$ is the $m \times m$ identity matrix. After solving this problem, \mathbf{U} encodes which group the task T_t belongs to. It turns out that this simple extension in (2) fails to correctly infer the group structure as it is biased towards finding a smaller number of groups. Figure 1 shows a toy example illustrating this. The following proposition generalizes this issue.

Proposition 1. *Consider the problem of minimizing (2) with respect to \mathbf{U} for a fixed $\widehat{\mathbf{W}}$. The assignment such that $\widehat{\mathbf{U}}_g = \mathbf{I}^{m \times m}$ for some $g \in \mathcal{G}$ and $\widehat{\mathbf{U}}_{g'} = \mathbf{0}^{m \times m}$ for all other $g' \in \mathcal{G} \setminus \{g\}$, is a minimizer of (2).*

Proof: Please refer to the appendix.

3 Learning Groups on Sparsity Patterns

In the previous section, we observed that the standard group norm is beneficial when the group structure \mathcal{G}^* is known but not suitable for inferring it. This is mainly because it is basically aggregating groups via the ℓ_1 norm; let $\mathbf{v} \in \mathbb{R}^N$ be a vector of $(\|\mathbf{W}\mathbf{U}_1\|_{1,2}, \|\mathbf{W}\mathbf{U}_2\|_{1,2}, \dots, \|\mathbf{W}\mathbf{U}_N\|_{1,2})^\top$, then the regularizer of (2) can be understood as $\|\mathbf{v}\|_1$. By the basic property of ℓ_1 norm, \mathbf{v} tends to be a sparse vector, making \mathbf{U} have a small number of active groups (we say some group g is active if there exists a task T_t such that $u_{g,t} = 1$.)

Based on this finding, we propose to use the ℓ_α norm ($\alpha \geq 2$) for summing up the regularizers from different groups, so that the final regularizer as a whole forces most




Fig. 2. For the two toy examples from Figure 1, we show the behaviour of $(\sum_g (\|\mathbf{W}\mathbf{U}_g\|_{1,2})^2)^{0.5}$ (the group regularizer in (3) with $\alpha = 0.5$). See the caption of Figure 1 for the choice of $\mathbf{W} = [\mathbf{w}^{(1)}, \mathbf{w}^{(2)}]$. In the example on the right, the regularizer now favors putting the tasks in separate groups.

of $\|\mathbf{W}\mathbf{U}_g\|_{1,2}$ to be non-zeros:

$$\begin{aligned} & \underset{\mathbf{W}, \mathbf{U}}{\text{minimize}} \sum_{t=1}^m \mathcal{L}(\mathbf{w}^{(t)}; \mathcal{D}_t) + \sum_{g \in \mathcal{G}} \lambda_g \left(\|\mathbf{W}\mathbf{U}_g\|_{1,2} \right)^\alpha \\ & \text{s. t. } \sum_{g \in \mathcal{G}} \mathbf{U}_g = \mathbf{I}^{m \times m}, \quad [\mathbf{U}_g]_{ik} \in \{0, 1\}. \end{aligned} \quad (3)$$

Note that strictly speaking, $\|\mathbf{v}\|_\alpha$ is defined as $(\sum_{i=1}^N |v_i|^\alpha)^{1/\alpha}$, but we ignore the relative effect of $1/\alpha$ in the exponent. One might want to get this exponent back especially when α is large. ℓ_α norms give rise to exactly the opposite effects in distributions of $u_{g,t}$, as shown in Figure 2 and Proposition 2.

Proposition 2. Consider a minimizer $\hat{\mathbf{U}}$ of (3), for any fixed $\widehat{\mathbf{W}}$. Suppose that there exist two tasks in a single group such that $\hat{w}_i^{(s)} \hat{w}_j^{(t)} \neq \hat{w}_j^{(s)} \hat{w}_i^{(t)}$. Then there is no empty group g such that $\hat{\mathbf{U}}_g = \mathbf{0}^{m \times m}$.

Proof: Please refer to the appendix.

Figure 3 visualizes the unit surfaces of different regularizers derived from (3) (i.e. $\sum_{g \in \mathcal{G}} (\|\mathbf{W}\mathbf{U}_g\|_{1,2})^\alpha = 1$ for different choices of α .) for the case where we have two groups, each of which has a single task. It shows that a large norm value on one group (in this example, on G_2 when $a = 0.9$) does not force the other group (i.e. G_1) to have a small norm as α becomes larger. This is evidenced in the bottom two rows of the third column (compare it with how ℓ_1 behaves in the top row). In other words, we see the benefits of using $\alpha \geq 2$ to encourage more active groups.

While the constraint $[\mathbf{U}_g]_{ik} \in \{0, 1\}$ in (3) ensures hard group memberships, solving it requires integer programming which is intractable in general. Therefore, we relax the constraint on \mathbf{U} to $0 \leq [\mathbf{U}_g]_{ik} \leq 1$. However, this relaxation along with the ℓ_α norm over groups prevents both $\|\mathbf{W}\mathbf{U}_g\|_{1,2}$ and also individual $[\mathbf{U}_g]_{ik}$ from being zero. For example, suppose that we have two tasks (in \mathbb{R}^2) in a single group, and $\alpha = 2$. Then, the regularizer for any g can be written as $\left(\sqrt{(w_1^{(1)})^2 u_{g,1}^2 + (w_1^{(2)})^2 u_{g,2}^2} + \sqrt{(w_2^{(1)})^2 u_{g,1}^2 + (w_2^{(2)})^2 u_{g,2}^2} \right)^2$. To simplify the situation, assume further that all entries of \mathbf{W} are uniformly a constant w . Then, this regularizer for a single g would be




Fig. 3. Unit balls of the regularizer in (3) for different values of α . Suppose we have $\mathbf{w}^{(1)} = (x, y)^\top$ in group G_1 and $\mathbf{w}^{(2)} = (z, a)^\top$ in G_2 . In order to visualize in 3-d space, we vary 3 variables x, y and z , and fix $w_2^{(2)}$ to some constant a . The first row is using ℓ_1 norm for summing the groups: $(|x| + |y|) + (|z| + a)$, the second row is using ℓ_2 norm: $\sqrt{(|x| + |y|)^2 + (|z| + a)^2}$, and the last row is using ℓ_5 norm: $((|x| + |y|)^5 + (|z| + a)^5)^{0.2}$. As a increases (from the first column to the third one), \mathbf{w}_1 quickly shrinks to zero in case of ℓ_1 summation. On the other hand, in case of ℓ_2 summation, x and y in $\mathbf{w}^{(1)}$ are allowed to be non-zero, while z shrinks to zero. This effect gets clearer as α increases.

simply reduced to $4w^2(u_{g,1}^2 + u_{g,2}^2)$, and therefore the regularizer over all groups would be $4w^2(\sum_{t=1}^m \sum_g u_{g,t}^2)$. Now it is clearly seen that the regularizer has an effect of grouping over the group membership vector $(u_{g_1,t}, u_{g_2,t}, \dots, u_{g_N,t})$ and encouraging the set of membership parameters for each task to be uniform.

To alleviate this challenge, we re-parameterize $u_{g,t}$ with a new membership parameter $u'_{g,t} := \sqrt{u_{g,t}}$. The constraint does not change with this re-parameterization: $0 \leq u'_{g,t} \leq 1$. Then, in the previous example, the regularization over all groups would be (with some constant factor) the sum of ℓ_1 norms, $\|(u_{g_1,t}, u_{g_2,t}, \dots, u_{g_N,t})\|_1$ over all tasks, which forces them to be sparse. Note that even with this change, the activations of groups are not sparse since the sum over groups is still done by the ℓ_2 norm.

Toward this, we finally introduce the following problem to jointly estimate \mathbf{U} and \mathbf{W} (specifically with focus on the case when α is set to 2):

$$\begin{aligned} & \underset{\mathbf{W}, \mathbf{U}}{\text{minimize}} \sum_{t=1}^m \mathcal{L}(\mathbf{w}^{(t)}; \mathcal{D}_t) + \sum_{g \in \mathcal{G}} \lambda_g \left(\|\mathbf{W} \sqrt{\mathbf{U}}_g\|_{1,2} \right)^2 \\ & \text{s. t. } \sum_{g \in \mathcal{G}} \mathbf{U}_g = \mathbf{I}^{m \times m}, \quad 0 \leq [\mathbf{U}_g]_{ik} \leq 1. \end{aligned} \quad (4)$$

where $\sqrt{\mathbf{M}}$ for a matrix \mathbf{M} is obtained from element-wise square root operations of \mathbf{M} . Note that (3) and (4) are not equivalent, but the minimizer $\hat{\mathbf{U}}$ given any fixed \mathbf{W} usually is actually binary or has the same objective with some other binary $\hat{\mathbf{U}}'$ (see




Fig. 4. Comparisons on different regularizers. (a) \mathbf{W} for 30 tasks with 21 features, which is assumed to be *known and fixed*. Three groups are clearly separated: (T_1-T_{10}) , $(T_{11}-T_{20})$, $(T_{21}-T_{30})$ whose nonzero elements are block diagonally located (black is negative, white is positive). (b) Learned \mathbf{U} from (2) (with a relaxing of discrete constraints, and the square root reparameterization). All tasks are assigned to a single group. (c) Estimation on \mathbf{U} using (2) after relaxing the integer constraints on \mathbf{U} . All these groups are active (have tasks in them), but most of $u_{g,t}$ are not binary. (d) Estimation on \mathbf{U} using (4). All group memberships are recovered correctly. Note that the order of groups does not matter. For (b)-(d), white is 1 and black is 0.

Theorem 1 of Kang et al. (2011) for details). As a toy example, we show in Figure 4, the estimated \mathbf{U} (for a known \mathbf{W}) via different problems presented so far ((2), (4)).




Fig. 5. Comparison of the regularizer from Kang et al. (2011) when learned on our synthetic dataset (set-1). Fig 4 (a) shows the \mathbf{W}^* . (a) the learned \mathbf{Q} , analogous to \mathbf{U} in our notation (white is 1 and black is 0) and (b) $\widehat{\mathbf{W}}$

Fusing the Group Assignments The approach derived so far works well when the number of groups $N \ll m$, but can create many singleton groups when N is very large. We add a final modification to our objective to encourage tasks to have similar group membership wherever warranted. This makes the method more robust to the mis-specification of the number of groups, ‘ N ’ as it prevents the grouping from becoming too fragmented when $N \gg N^*$. For each task $t = 1, \dots, m$, we define $N \times N$ matrix $\mathbf{V}_t := \text{diag}(u_{1,t}, \dots, u_{N,t})$. Note that the \mathbf{V}_t are entirely determined by the \mathbf{U}_g matrices, so no actual additional variables are introduced. Equipped with this additional notation, we obtain the following objective where $\|\cdot\|_F$ denotes the Frobenius norm (the

element-wise ℓ_2 norm), and μ is the additional regularization parameter that controls the number of active groups.

$$\begin{aligned}
& \underset{\mathbf{W}, \mathbf{U}}{\text{minimize}} \sum_{t=1}^m \mathcal{L}(\mathbf{w}^{(t)}; \mathcal{D}_t) + \sum_{g \in \mathcal{G}} \lambda_g \left(\|\mathbf{W} \sqrt{\mathbf{U}_g}\|_{1,2} \right)^2 \\
& \quad + \mu \sum_{t < t'} \|\mathbf{V}_t - \mathbf{V}_{t'}\|_{\text{F}}^2 \\
& \text{s. t. } \sum_{g \in \mathcal{G}} \mathbf{U}_g = \mathbf{I}^{m \times m}, \quad [\mathbf{U}_g]_{ik} \in [0, 1]
\end{aligned} \tag{5}$$

3.1 Theoretical comparison of approaches:

It is natural to ask whether enforcing the shared sparsity structure, when groups are unknown leads to any efficiency in the number of samples required for learning. In this section, we will use intuitions from the high-dimensional statistics literature in order to compare the sample requirements of different alternatives such as independent lasso or the approach of Kang et al. (2011). Since the formal analysis of each method requires making different assumptions on the data X and the noise, we will instead stay intentionally informal in this section, and contrast the number of samples each approach would require, assuming that the desired structural conditions on the x 's are met. We evaluate all the methods under an idealized setting where the structural assumptions on the parameter matrix W motivating our objective (2) hold exactly. That is, the parameters form N groups, with the weights in each group taking non-zero values only on a common subset of features of size at most s . We begin with our approach.

Complexity of Sparsity Grouped MTL Let us consider the simplest inefficient version of our method, a generalization of subset selection for Lasso which searches over all feature subsets of size s . It picks one subset S_g for each group g and then estimates the weights on S_g independently for each task in group g . By a simple union bound argument, we expect this method to find the right support sets, as well as good parameter values in $O(Ns \log d + ms)$ samples. This is the complexity of selecting the right subset out of $\binom{d}{s}$ possibilities for each group, followed by the estimation of s weights for each task. We note that there is no direct interaction between m and d in this bound.

Complexity of independent lasso per task An alternative approach is to estimate an s -sparse parameter vector for each task independently. Using standard bounds for ℓ_1 regularization (or subset selection), this requires $O(s \log d)$ samples per task, meaning $O(ms \log d)$ samples overall. We note the multiplicative interaction between m and $\log d$ here.

Complexity of learning all tasks jointly A different extreme would be to put all the tasks in one group, and enforce shared sparsity structure across them using $\|\cdot\|_{1,2}$ regularization on the entire weight matrix. The complexity of this approach depends on the sparsity of the union of all tasks which is Ns , much larger than the sparsity of individual groups. Since each task requires to estimate its own parameters on this

Algorithm 1 SG-MTL (eqn (4))

Input: $\{\mathcal{D}_i\}_{i=1}^m$
Initialize \mathbf{W} using single task learning
Initialize \mathbf{U} to random matrix
repeat
 Update \mathbf{U} by solving a projected gradient descent
 for all tasks $t = 1, 2, \dots, m$ **do**
 Update $\mathbf{w}^{(t)}$ using a coordinate descent for all features $j = 1, 2, \dots, d$
 end for
until stopping criterion is satisfied

shared sparse basis, we end up requiring $O(msN \log d)$ samples, with a large penalty for ignoring the group structure entirely.

Complexity of Kang et al. (2011) As yet another baseline, we observe that an s -sparse weight matrix is also naturally low-rank with rank at most s . Consequently, the weight matrix for each group has rank at most s , plausibly making this setting a good fit for the approach of Kang et al. (2011). However, appealing to the standard results for low-rank matrix estimation (see e.g. Negahban & Wainwright (2011)), learning a $d \times n_g$ weight matrix of rank at most s requires $O(s(n_g + d))$ samples, where n_g is the number of tasks in the group g . Adding up across tasks, we find that this approach requires a total of $O(s(m + md))$, considerably higher than all other baselines even if the groups are already provided. It is easy to see why this is unavoidable too. Given a group, one requires $O(ms)$ samples to estimate the entries of the s linearly independent rows. A method utilizing sparsity information knows that the rest of the columns are filled with zeros, but one that only knows that the matrix is low-rank assumes that the remaining $(d - s)$ rows all lie in the linear span of these s rows, and the coefficients of that linear combination need to be estimated giving rise to the additional sample complexity. In a nutshell, this conveys that estimating a sparse matrix using low-rank regularizers is sample inefficient, an observation hardly surprising from the available results in high-dimensional statistics but important in comparison with the baseline of Kang et al. (2011).

For ease of reference, we collect all these results in Table 1 below.

	SG-MTL	Lasso	Single group	Kang et al. (2011)
Samples	$O(Ns \log d + ms)$	$O(ms \log d)$	$O(msN \log d)$	$O(s(m + md))$

Table 1. Sample complexity estimates of recovering group memberships and weights using different approaches

4 Optimization

We solve (4) by alternating minimization: repeatedly solve one variable fixing the other until convergence (Algorithm 1) We discuss details below.

Solving (4) w.r.t \mathbf{U} : This step is challenging since we lose convexity due the reparameterization with a square root. The solver might stop with a premature \mathbf{U} stuck in

a local optimum. However, in practice, we can utilize the random search technique to get the minimum value over multiple re-trials. Our experimental results reveal that the following projected gradient descent method performs well.

Given a fixed \mathbf{W} , solving for \mathbf{U} only involves the regularization term i.e $R(\mathbf{U}) = \sum_{g \in \mathcal{G}} \lambda_g \left(\sum_{j=1}^d \|\mathbf{W}_j \sqrt{\mathbf{U}_g}\|_2 \right)^2$ which is differentiable w.r.t \mathbf{U} . The derivative is shown in the appendix along with the extension for the fusion penalty from (5). Finally after the gradient descent step, we project $(u_{g_1,t}, u_{g_2,t}, \dots, u_{g_N,t})$ onto the simplex (independently repeat the projection for each task) to satisfy the constraints on it. Note that projecting a vector onto a simplex can be done in $O(m \log m)$ (Chen & Ye, 2011).

Solving (4) w.r.t \mathbf{W} This step is more amenable in the sense that (4) is convex in \mathbf{W} given \mathbf{U} . However, it is not trivial to efficiently handle the complicated regularization terms. Contrast to \mathbf{U} which is bounded by $[0, 1]$, \mathbf{W} is usually unbounded which is problematic since the regularizer is (not always, but under some complicated conditions discovered below) non-differentiable at 0.

While it is challenging to directly solve with respect to the entire \mathbf{W} , we found out that the coordinate descent (in terms of each element in \mathbf{W}) has a particularly simple structured regularization.

Consider any $w_j^{(t)}$ fixing all others in \mathbf{W} and \mathbf{U} ; the regularizer $R(\mathbf{U})$ from (4) can be written as

$$\sum_{g \in \mathcal{G}} \lambda_g \left\{ u_{g,t} (w_j^{(t)})^2 + 2 \left(\sum_{j' \neq j} \sqrt{\sum_{t'=1}^m u_{g,t'} (w_{j'}^{(t')})^2} \right) \sqrt{\sum_{t'=1}^m u_{g,t'} (w_j^{(t')})^2} \right\} + C(j, t) \quad (6)$$

where $w_j^{(t)}$ is the only variable in the optimization problem, and $C(j, t)$ is the sum of other terms in (4) that are constants with respect to $w_j^{(t)}$.

For notational simplicity, we define $\kappa_{g,t} := \sum_{t' \neq t} u_{g,t'} (w_j^{(t')})^2$ that is considered as a constant in (6) given \mathbf{U} and $\mathbf{W} \setminus \{w_j^{(t)}\}$. Given $\kappa_{g,t}$ for all $g \in \mathcal{G}$, we also define \mathcal{G}^0 as the set of groups such that $\kappa_{g,t} = 0$ and \mathcal{G}^+ for groups s.t. $\kappa_{g,t} > 0$. Armed with this notation and with the fact that $\sqrt{x^2} = |x|$, we are able to rewrite (6) as

$$\begin{aligned} & \sum_{g \in \mathcal{G}} \lambda_g u_{g,t} (w_j^{(t)})^2 + 2 \sum_{g \in \mathcal{G}^+} \lambda_g \left(\sum_{j' \neq j} \sqrt{\sum_{t'=1}^m u_{g,t'} (w_{j'}^{(t')})^2} \right) \sqrt{\sum_{t'=1}^m u_{g,t'} (w_j^{(t')})^2} \quad (7) \\ & + 2 \sum_{g \in \mathcal{G}^0} \lambda_g \left(\sum_{j' \neq j} \sqrt{\sum_{t'=1}^m u_{g,t'} (w_{j'}^{(t')})^2} \right) \sqrt{u_{g,t}} |w_j^{(t)}| \end{aligned}$$

where we suppress the constant term $C(j, t)$. Since $\sqrt{x^2 + a}$ is differentiable in x for any constant $a > 0$, the first two terms in (7) are differentiable with respect to $w_j^{(t)}$, and the only non-differentiable term involves the absolute value of the variable, $|w_j^{(t)}|$. As a result, (7) can be efficiently solved by proximal gradient descent followed by an

element-wise soft thresholding. Please see appendix for the gradient computation of \mathcal{L} and soft-thresholding details.

5 Experiments

We conduct experiments on two synthetic and two real datasets and compare with the following approaches.

- 1) Single task learning (STL): Independent models for each task using elastic-net regression/classification.
- 2) AllTasks: We combine data from all tasks into a single task and learn an elastic-net model on it
- 3) Clus-MTL: We first learn STL for each task, and then cluster the task parameters using k-means clustering. For each task cluster, we then train a multitask lasso model.
- 4) GO-MTL: group-overlap MTL (Kumar & Daume, 2012)
- 5) Kang et al. (2011): nuclear norm based task grouping
- 6) SG-MTL : our approach from equation 4
- 7) Fusion SG-MTL: our model with a fusion penalty (see Section 3)

Table 2. Synthetic datasets: (Upper table) Average MSE from 5 fold CV. (Lower table) Varying group sizes and the corresponding average MSE. For each method, lowest MSE is highlighted.

Dataset	STL	ClusMTL	Kang	SG-MTL	Fusion SG-MTL
set-1 (3 groups)	1.067	1.221	1.177	0.682	0.614
set-2 (5 groups)	1.004	1.825	0.729	0.136	0.130

Synthetic data-2 with 30% feature overlap across groups					
Method	Number of groups, N				
	2	4	$N^*=5$	6	10
ClusMTL	1.900	1.857	1.825	1.819	1.576
Kang et al. (2011)	0.156	0.634	0.729	0.958	1.289
SG-MTL	0.145	0.135	0.136	0.137	0.137
Fusion SG-MTL	0.142	0.139	0.130	0.137	0.137

5.1 Results on synthetic data

The first setting is similar to the synthetic data settings used in Kang et al. (2011) except for how W is generated (see Fig 4(a) for our parameter matrix and compare it with Sec 4.1 of Kang et al. (2011)). We have 30 tasks forming 3 groups with 21 features and 15 examples per task. Each group in W is generated by first fixing the zero components and then setting the non-zero parts to a random vector w with unit variance. Y_t for task t is $X_t W_t + \epsilon$. For the second dataset, we generate parameters in a similar manner as above, but with 30 tasks forming 5 groups, 100 examples per task, 150 features and a

30% overlap in the features across groups. In table 2, we show 5-fold CV results and in Fig 6 we show the groups (U) found by our method.

How many groups? Table 2 (Lower) shows the effect of increasing group size on three methods (smallest MSE is highlighted). For our methods, we observe a dip in MSE when N is close to N^* . In particular, our method with the fusion penalty gets the lowest MSE at $N^* = 5$. Interestingly, Kang et al. (2011) seems to prefer the smallest number of clusters, possibly due to the low-rank structural assumption of their approach, and hence cannot be used to learn the number of clusters in a sparsity based setting.




Fig. 6. Groupings found for the case where $N^* = 5$. We show results of a typical run of our method with $N=2,4,6$. On the x-axis are the 30 tasks and on the y-axis are the group ids.

	μ_{MSE}	σ_{mse}	μ_{R^2}	σ_{R^2}
STL	0.811	0.02	0.223	0.01
AllTasks	0.908	0.01	0.092	0.01
ClusMTL	0.823	0.02	0.215	0.02
GOMTL	0.794	0.01	0.218	0.01
Kang	1.011	0.03	0.051	0.03
Fusion SG-MTL	0.752	0.01	0.265	0.01

Table 3. QSAR prediction: average MSE and R^2 over 10 train:test splits with 100 examples per task in the training split (i.e $n=100$). The standard deviation of MSE is also shown. For all group learning methods, we use $N=5$.




Table 4. Avg. MSE of our method (Fusion SG-MTL) as a function of the number of clusters and training data size. The best average MSE is observed with 7 clusters and $n=300$ training examples (green curve with squares). The corresponding R^2 for this setting is 0.401.

5.2 Quantitative Structure Activity Relationships (QSAR) Prediction: Merck dataset

Given features generated from the chemical structures of candidate drugs, the goal is to predict their molecular activity (a real number) with the target. This dataset from

Method	AUC-PR
STL	0.825
AllTasks	0.709
ClusMTL	0.841
GOMTL	0.749
Kang	0.792
Fusion SG-MTL	0.837

Table 5. TFBS prediction: average AUC-PR, with training data size of 200 examples, test data of ≈ 1800 for number of groups $N=10$.




Table 6. Matrix indicating the groups learned by two methods. The rows show the task names and columns are clusters. A white entry at position (i, j) indicates that task i belongs to cluster j . (a) Groups learned by SG-MTL on the TFBS problem (b) Groups learned by the Clus-MTL baseline show that it tends to put all tasks in the same cluster.

Kaggle consists of 15 molecular activity data sets, each corresponding to a different target, giving us 15 tasks. There are between 1500 to 40000 examples and ≈ 5000 features per task, out of which 3000 features are common to all tasks. We create 10 train:test splits with 100 examples in training set (to represent a setting where $n \ll d$) and remaining in the test set. We report R^2 and MSE aggregated over these experiments in Table 3, with the number of task clusters N set to 5 (for the baseline methods we tried $N=2,5,7$). We found that Clus-MTL tends to put all tasks in the same cluster for any value of N . Our method has the lowest average MSE.

In Figure 4, for our method we show how MSE changes with the number of groups N (along x-axis) and over different sizes of the training/test split. The dip in MSE for $n = 50$ training examples (purple curve marked ‘x’) around $N = 5$ suggests there are 5 groups. The learned groups are shown in the appendix in Fig ?? followed by a discussion of the groupings.

5.3 Transcription Factor Binding Site Prediction (TFBS)

This dataset was constructed from processed ChIP-seq data for 37 transcription factors (TFs) downloaded from ENCODE database (Consortium et al., 2012). Training data is generated in a manner similar to prior literature (Setty & Leslie, 2015). Positive examples consist of ‘peaks’ or regions of the DNA with binding events and negatives are regions away from the peaks and called ‘flanks’. Each of the 37 TFs represents a task. We generate all 8-mer features and select 3000 of these based on their frequency

in the data. There are ≈ 2000 examples per task, which we divide into train:test splits using 200 examples (100 positive, 100 negative) as training data and the rest as test data. We report AUC-PR averaged over 5 random train:test splits in Table 5. For our method, we found the number of clusters giving the best AUC-PR to be $N = 10$. For the other methods, we tried $N=5,10,15$ and report the best AUC-PR.

Though our method does marginally better (not statistically significant) than the STL baseline, which is a ridge regression model, in many biological applications such as this, it is desirable to have an interpretable model that can produce biological insights. Our MTL approach learns groupings over the TFs which are shown in Fig 6(a). Overall, ClusMTL has the best AUC-PR on this dataset however it groups too many tasks into a single cluster (Fig 6(b)) and forces each group to have at least one task. Note how our method leaves some groups empty (column 5 and 7) as our objective provides a trade-off between adding groups and making groups cohesive.

6 Conclusion

We presented a method to learn group structure in multitask learning problems, where the task relationships are unknown. The resulting non-convex problem is optimized by applying the alternating minimization strategy. We evaluate our method through experiments on both synthetic and real-world data. On synthetic data with known group structure, our method outperforms the baselines in recovering them. On real data, we obtain a better performance while learning intuitive groupings. Code is available at: <https://github.com/meghana-kshirsagar/treemtl/tree/groups>
Full paper with appendix is available at: <https://arxiv.org/abs/1705.04886>

Acknowledgements We thank Alekh Agarwal for helpful discussions regarding Section 3.1. E.Y. acknowledges the support of MSIP/NRF (National Research Foundation of Korea) via NRF-2016R1A5A1012966 and MSIP/IITP (Institute for Information & Communications Technology Promotion of Korea) via ICT R&D program 2016-0-00563, 2017-0-00537.

References

- Agarwal, Arvind, Gerber, Samuel, and Daume, Hal. Learning multiple tasks using manifold regularization. In *Advances in neural information processing systems*, pp. 46–54, 2010.
- Argyriou, A., Evgeniou, T., and Pontil, M. Convex multi-task feature learning. *Machine Learning*, 2008.
- Bach, Francis, Jenatton, Rodolphe, Mairal, Julien, Obozinski, Guillaume, et al. Convex optimization with sparsity-inducing norms. *Optimization for Machine Learning*, 5:19–53, 2011.
- Baxter, Jonathan. A model of inductive bias learning. *J. Artif. Intell. Res.(JAIR)*, 12:149–198, 2000.
- Caruana, Rich. Multitask learning. *Mach. Learn.*, 28(1):41–75, July 1997. ISSN 0885-6125.
- Chen, Jianhui, Liu, Ji, and Ye, Jieping. Learning incoherent sparse and low-rank patterns from multiple tasks. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 5(4):22, 2012.
- Chen, Yunmei and Ye, Xiaoqing. Projection onto a simplex. *arXiv preprint arXiv:1101.6081*, 2011.

- Consortium, ENCODE Project et al. An integrated encyclopedia of dna elements in the human genome. *Nature*, 489(7414):57–74, 2012.
- Daumé III, Hal. Bayesian multitask learning with latent hierarchies. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, pp. 135–142. AUAI Press, 2009.
- Evgeniou, T. and Pontil, M. Regularized multi-task learning. *ACM SIGKDD*, 2004.
- Fei, Hongliang and Huan, Jun. Structured feature selection and task relationship inference for multi-task learning. *Knowledge and information systems*, 35(2):345–364, 2013.
- Gong, Pinghua, Ye, Jieping, and Zhang, Changshui. Robust multi-task feature learning. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 895–903. ACM, 2012.
- Jacob, L., Vert, J.P., and Bach, F.R. Clustered multi-task learning: A convex formulation. In *Advances in neural information processing systems (NIPS)*, pp. 745–752, 2009.
- Jalali, Ali, Sanghavi, Sujay, Ruan, Chao, and Ravikumar, Pradeep K. A dirty model for multi-task learning. *Advances in Neural Information Processing Systems*, pp. 964–972, 2010.
- Kang, Zhuoliang, Grauman, Kristen, and Sha, Fei. Learning with whom to share in multi-task feature learning. In *International Conference on Machine Learning (ICML)*, 2011.
- Kim, Seyoung and Xing, Eric P. Tree-guided group lasso for multi-task regression with structured sparsity. *The Proceedings of the International Conference on Machine Learning (ICML)*, 2010.
- Kumar, Abhishek and Daume, Hal. Learning task grouping and overlap in multi-task learning. In *ICML*, 2012.
- Liu, Jun, Ji, Shuiwang, and Ye, Jieping. Multi-task feature learning via efficient $l_{2,1}$ -norm minimization. In *Proceedings of the twenty-fifth conference on uncertainty in artificial intelligence (UAI)*, pp. 339–348, 2009.
- Ma, Junshui, Sheridan, Robert P, Liaw, Andy, Dahl, George E, and Svetnik, Vladimir. Deep neural nets as a method for quantitative structure–activity relationships. *Journal of chemical information and modeling*, 55(2):263–274, 2015.
- Maurer, Andreas. Bounds for linear multi-task learning. *The Journal of Machine Learning Research*, 7:117–139, 2006.
- Negahban, Sahand and Wainwright, Martin J. Estimation of (near) low-rank matrices with noise and high-dimensional scaling. *The Annals of Statistics*, pp. 1069–1097, 2011.
- Passos, Alexandre, Rai, Piyush, Wainer, Jacques, and Daume III, Hal. Flexible modeling of latent task structures in multitask learning. *The Proceedings of the International Conference on Machine Learning (ICML)*, 2012.
- Rao, Nikhil, Cox, Christopher, Nowak, Rob, and Rogers, Timothy T. Sparse overlapping sets lasso for multitask learning and its application to fmri analysis. In *Advances in neural information processing systems*, pp. 2202–2210, 2013.
- Setty, Manu and Leslie, Christina S. Seqgl identifies context-dependent binding signals in genome-wide regulatory element maps. *PLoS Comput Biol*, 11(5):e1004271, 2015.
- Tibshirani, Robert. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 267–288, 1996.
- Widmer, C., Leiva, J., Altun, Y., and Ratsch, G. Leveraging sequence classification by taxonomy-based multitask learning. *RECOMB*, 2010.
- Yu, Kai, Tresp, Volker, and Schwaighofer, Anton. Learning gaussian processes from multiple tasks. In *Proceedings of the 22nd international conference on Machine learning*, pp. 1012–1019. ACM, 2005.
- Yuan, Ming and Lin, Yi. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2006.
- Zhang, Yi and Schneider, Jeff G. Learning multiple tasks with a sparse matrix-normal penalty. In *Advances in Neural Information Processing Systems*, pp. 2550–2558, 2010.
- Zhang, Yu and Yeung, Dit-Yan. A convex formulation for learning task relationships in multi-task learning. 2010.