

where a and b are defined as the location and scale parameters. The Cauchy prior is widely used for feature learning (Carvalho, Polson, and Scott 2009; Daniel, Jose Miguel, and Pierre 2013). Moreover, in order to obtain sparse hyperparameter $\phi_{j(l_i)}$, we place the Jeffreys prior (Qi et al. 2004) over the $\phi_{j(l_i)}$:

$$p(\phi_{j(l_i)}) \propto \frac{1}{\phi_{j(l_i)}}, \quad i = 1, \dots, m. \quad (6)$$

One advantage of using the Jeffreys prior is that the prior has no hyper-parameters.

Eqs. (3)-(6) define the PTS model. In the next section, we discuss how to learn the model parameters.

Parameter Inference

In this section, we develop an EM algorithm (Dempster, Laird, and Rubin 1977) to learn the model parameters except $\sigma = (\sigma_1, \dots, \sigma_m)^T$. In order to achieve model flexibility, σ is treated as a tuning parameter and for simplicity, different σ_i 's are assumed to be identical, i.e., $\sigma_1^2 = \dots = \sigma_m^2 = \lambda$ where λ is chosen via the cross validation method in our experiments. In the EM algorithm, $\phi^{(L)} = \{\phi_{j(l_i)}\}_{j \in \mathbb{N}_d, i \in \mathbb{N}_m}$ are treated as hidden variables and the model parameters are denoted by $\Theta = \{\mathbf{h}, \mathbf{b}, \phi^{(E)}\}$ where \mathbf{h} denotes the set of all component coefficients, $\mathbf{b} = (b_1, \dots, b_m)^T$, and $\phi^{(E)} = \{\phi_{j(e)}\}_{j \in \mathbb{N}_d, e \in \mathcal{I}(\mathcal{T})}$. In the following, we give the details in the EM algorithm.

E-step: we construct the Q -function as

$$\begin{aligned} Q(\Theta|\Theta^{(t)}) &= \mathbb{E}[\ln p(\Theta|\mathbf{y}, \phi)] \\ &= \int p(\phi|\mathbf{y}, \Theta^{(t)}) \ln p(\Theta|\mathbf{y}, \phi) d\phi, \end{aligned}$$

where $\phi = \{\phi_{j(c)}\}_{j \in \mathbb{N}_d, c \in \mathcal{T}}$ and $\Theta^{(t)}$ denotes the estimate of Θ in the t th iteration. By defining $\mathbf{h}^{(L)} = \{h_{j(l_i)}\}_{j \in \mathbb{N}_d, i \in \mathbb{N}_m}$ and $\mathbf{h}^{(E)} = \{h_{j(e)}\}_{j \in \mathbb{N}_d, e \in \mathcal{I}(\mathcal{T})}$, it is easy to get

$$\begin{aligned} \ln p(\Theta|\mathbf{y}, \phi) & \\ &\propto \ln p(\mathbf{y}|\mathbf{h}, \sigma) + \ln p(\mathbf{h}^{(E)}|\phi^{(E)}) + \ln p(\mathbf{h}^{(L)}|\phi^{(L)}) \\ &\propto -\frac{1}{2\lambda} \sum_{i=1}^m \sum_{k=1}^{n_i} (y_i^{(k)} - \sum_{j=1}^d x_{ji}^{(k)} \prod_{c \in \mathbf{H}_i} h_{j(c)} - b_i)^2 \\ &\quad - \sum_{j=1}^d \sum_{e \in \mathcal{I}(\mathcal{T})} \ln \left(\frac{h_{j(e)}^2}{\phi_{j(e)}^2} + 1 \right) - \sum_{j=1}^d \sum_{i=1}^m \frac{h_{j(l_i)}^2}{2\phi_{j(l_i)}^2} \\ &\quad - \sum_{j=1}^d \sum_{e \in \mathcal{I}(\mathcal{T})} \ln \phi_{j(e)}. \end{aligned}$$

In addition, we have

$$p(\phi_{j(l_i)}|\mathbf{y}, \Theta^{(t)}) \propto p(\phi_{j(l_i)})p(h_{j(l_i)}^{(t)}|\phi_{j(l_i)}).$$

Then we compute the following expectation as

$$\begin{aligned} \mathbb{E}\left[\frac{1}{2\phi_{j(l_i)}^2}|\mathbf{y}, \Theta^{(t)}\right] &= \frac{\int_0^\infty \frac{1}{2\phi_{j(l_i)}^2} p(\phi_{j(l_i)})p(h_{j(l_i)}^{(t)}|\phi_{j(l_i)})d\phi_{j(l_i)}}{\int_0^\infty p(\phi_{j(l_i)})p(h_{j(l_i)}^{(t)}|\phi_{j(l_i)})d\phi_{j(l_i)}} \\ &= \frac{1}{[2h_{j(l_i)}^{(t)}]^2}. \end{aligned}$$

By defining $\beta_{ji} = \frac{1}{[2h_{j(l_i)}^{(t)}]^2}$, the Q -function can be simplified as

$$\begin{aligned} Q(\Theta|\Theta^{(t)}) &= -\frac{1}{2\lambda} \sum_{i=1}^m \sum_{k=1}^{n_i} (y_i^{(k)} - \sum_{j=1}^d x_{ji}^{(k)} \prod_{c \in \mathbf{H}_i} h_{j(c)} - b_i)^2 \\ &\quad - \sum_{j=1}^d \sum_{e \in \mathcal{I}(\mathcal{T})} \ln \left(\frac{h_{j(e)}^2}{\phi_{j(e)}^2} + 1 \right) - \sum_{j=1}^d \sum_{i=1}^m \beta_{ji} h_{j(l_i)}^2 \\ &\quad - \sum_{j=1}^d \sum_{e \in \mathcal{I}(\mathcal{T})} \ln \phi_{j(e)}. \end{aligned} \quad (7)$$

M-step: we maximize Eq. (7) to update the estimates of \mathbf{h} , \mathbf{b} , and $\phi^{(E)}$. For the estimation of $\mathbf{h}^{(E)}$ and $\phi^{(E)}$, we solve the following optimization problem:

$$\begin{aligned} \min_{\mathbf{h}^{(E)}, \phi^{(E)}} J &= \frac{1}{2} \sum_{i=1}^m \sum_{k=1}^{n_i} (y_i^{(k)} - \sum_{j=1}^d x_{ji}^{(k)} \prod_{c \in \mathbf{H}_i} h_{j(c)} - b_i)^2 \\ &\quad + \lambda \sum_{j=1}^d \sum_{e \in \mathcal{I}(\mathcal{T})} \ln \left(\frac{h_{j(e)}^2}{\phi_{j(e)}^2} + 1 \right) + \lambda \sum_{j=1}^d \sum_{e \in \mathcal{I}(\mathcal{T})} \ln \phi_{j(e)}. \end{aligned} \quad (8)$$

By setting the derivatives of J with respect to $\phi^{(E)}$ to zero and then get

$$\phi_{j(e)} = |h_{j(e)}|. \quad (9)$$

By plugging the solution in Eq. (9), we can simplify problem (8) as

$$\begin{aligned} \min_{\mathbf{h}^{(E)}} \bar{J} &= \frac{1}{2} \sum_{l_i \in \mathcal{D}(e)} \sum_{k=1}^{n_i} (\bar{y}_i^{(k)} - \sum_{j=1}^d h_{j(e)} \bar{x}_{ji}^{(k)})^2 \\ &\quad + \lambda \sum_{j=1}^d \ln |h_{j(e)}|, \end{aligned} \quad (10)$$

where $e \in \mathcal{I}(\mathcal{T})$, $\bar{x}_{ji}^{(k)} = x_{ji}^{(k)} \prod_{c \in \mathbf{H}_i, c \neq e} h_{j(c)}^{(t)}$, and $\bar{y}_i^{(k)} = y_i^{(k)} - b_i^{(t)}$. Problem (10) is non-convex since the second term of the objective function is non-convex. To solve this problem, we use the majorization-minimization (MM) algorithm (Lange, Hunter, and Yang 2000) to solve this problem. For numerical stability, we slightly modify Eq. (10) by replacing the second term with $\sum_{j=1}^d \ln(|h_{j(e)}| + \alpha)$ where α is a tuning parameter. We denote the solution obtained in the t' th iteration in the MM algorithm as $h_{j(e)}^{(t')}$. Then, in the $(t' + 1)$ th iteration, due to the concavity property of the logarithm function $\ln(\cdot)$, we have

$$\sum_{j=1}^d \ln(|h_{j(e)}| + \alpha) \leq \sum_{j=1}^d \left[\ln(|h_{j(e)}^{(t')}| + \alpha) + \frac{|h_{j(e)}| - |h_{j(e)}^{(t')}|}{|h_{j(e)}^{(t')}| + \alpha} \right].$$

Thus, in the $(t' + 1)$ th iteration of the MM algorithm, we only need to solve a weighted ℓ_1 minimization problem (Candes, Wakin, and Boyd 2008; Wipf and Nagarajan 2010):

$$\frac{1}{2} \sum_{l_i \in \mathcal{D}(e)} \sum_{k=1}^{n_i} (\bar{y}_i^{(k)} - \sum_{j=1}^d h_{j(e)} \bar{x}_{ji}^{(k)})^2 + \lambda \sum_{j=1}^d \frac{|h_{j(e)}|}{|h_{j(e)}^{(t')}| + \alpha}. \quad (11)$$

Problem (11) can be solved efficiently by some mature Lasso-style solvers. Moreover, according to (Lange, Hunter, and Yang 2000), the MM algorithm is guaranteed to converge to a local optimum.

For the estimation of $\mathbf{h}^{(L)}$, we need to solve m weighted ridge regression problems with each one formulated as

$$\min_{\mathbf{h}^{(L)}} \tilde{J} = \frac{1}{2} \sum_{k=1}^{n_i} (\bar{y}_i^{(k)} - \sum_{j=1}^d h_{j(l_i)} \tilde{x}_{ji}^{(k)})^2 + \lambda \sum_{j=1}^d \beta_{ji} h_{j(l_i)}^2, \quad (12)$$

where $i \in \mathbb{N}_m$ and $\tilde{x}_{ji}^{(k)} = x_{ji}^{(k)} \prod_{c \in \mathcal{A}(l_i)} h_{j(c)}^{(t)}$. Problem (12) has an analytical solution as

$$h_{j(l_i)} = (\tilde{\mathbf{X}}_i^T \tilde{\mathbf{X}}_i + \mathbf{D}_i)^{-1} \tilde{\mathbf{X}}_i^T \bar{\mathbf{y}}_i,$$

where $\tilde{\mathbf{X}}_i = (\tilde{\mathbf{x}}_i^{(1)}, \dots, \tilde{\mathbf{x}}_i^{(n_i)})^T$, $\bar{\mathbf{y}}_i = (\bar{y}_i^{(1)}, \dots, \bar{y}_i^{(n_i)})^T$, and \mathbf{D}_i is a $d \times d$ diagonal matrix with $\lambda \beta_{ji}$ as the j th diagonal element.

For the estimation of \mathbf{b} , we set the derivative of $Q(\Theta | \Theta^{(t)})$ with respect to \mathbf{b} to zero and get the solution of $b_i^{(t+1)}$ as

$$b_i^{(t+1)} = \frac{1}{n_i} \sum_{k=1}^{n_i} (y_i^{(k)} - \sum_{j=1}^d x_{ji}^{(k)} \prod_{c \in \mathbf{H}_i} h_{j(c)}).$$

Note that the Cauchy prior used in the internal nodes is to make the solution sparse as shown in Lasso-style problem (11). For the leaf nodes, the normal prior is used to get non-sparse solutions as shown in problems (12), since there is no need for sparsity.

Comparison between the Multi-Component SBD and PBD Techniques

In this section, we discuss the relation between the multi-component SBD and PBD techniques utilized in MTL. For the SBD model, we take the multi-task cascade model in (Zweig and Weinshall 2013) as a representative model. For the PBD model with multiple component coefficients, to the best of our knowledge, our PTS method is the first one.

We first compare the model complexity. For the multi-component SBD model, component coefficients have the same size as the model coefficients due to the sum operator. As in the multi-task cascade model, the model involves Ldm parameters, where $L \geq 2$ is the number of the predefined cascade. For the multi-component PBD model as introduced in this paper, there are only $d|\mathcal{T}|$ parameters, where $|\mathcal{T}| = m + |\mathcal{I}(\mathcal{T})| \ll Lm$. So the complexity of multi-component PBD models is lower than that of the multi-component SBD counterparts in terms of the model parameters.

Next, we examine the sparsity property. For the multi-component SBD model, each model coefficient w_{ji} equals 0 if and only if all of its component coefficients are 0 or the sum of the component coefficients is 0. In contrast, in the multi-component PBD model, when any one of its component coefficient equals 0, w_{ji} will become 0. From this point of view, the multi-component PBD model is more likely to learn sparse solutions.

Related Work

Besides the existing SDB and PBD models, the tree-guided group Lasso (TGGL) in (Kim and Xing 2010) addresses a similar problem for multi-output regression where the structure among outputs is encoded in a tree as a priori information. The TGGL method is formulated as a group lasso with the groups induced by the tree structure but the coefficient decomposition technique is not utilized in that model. In TGGL, the sparsity pattern is induced by the groups defined in the tree structure, which differs from the tree sparsity in our method. In addition, the formulation of TGGL is only for the multi-output case and not easy to extend to the general multi-task setting. Some hierarchical MTL models (Daumé III 2009; Görnitz et al. 2011) are proposed to exploit the hierarchical relations between tasks. However, those methods are not for sparse learning and so their objectives are different from this work.

Experiments

In this section, we conduct empirical experiments on both synthetic and real-world problems to study the proposed PTS method. Baselines used for comparison include the multi-task feature learning (MTFL) model (Liu, Ji, and Ye 2009), the dirty model (DM) (Jalali et al. 2010), the multi-level Lasso (MLL) model (Lozano and Swirszcz 2012), the tree-guided group Lasso (TGGL) model (Kim and Xing 2010), and the multi-task cascade (Cascade) model (Zweig and Weinshall 2013).

Synthetic Data

We first evaluate the proposed method on synthetic data. We simulate a multi-task regression problem with $m = 8$ tasks and $d = 30$ features. We assume all the learning tasks share the same data matrix \mathbf{X} , where \mathbf{X} is a $n \times d$ matrix with $n = 100$ samples. Each column of \mathbf{X} is generated from a normal distribution $\mathcal{N}(\mathbf{0}, \mathbf{I}_n)$ where $\mathbf{0}$ denotes a zero vector or matrix with appropriate size and \mathbf{I}_n is an $n \times n$ identity matrix. To simulate the tree sparsity pattern in the solutions, we generate \mathbf{W} according to Figure 2(b) where the dark elements denote non-zero coefficients and the white ones are zero. In Figure 2(b), the columns correspond to the tasks and the rows represent the features. Moreover, the corresponding tree structure is plotted in Figure 2(a). Figures 2(c)-2(h) give the sparse patterns in \mathbf{W} with respect to different features. The dark nodes in the trees represent the non-zero component coefficients in Eq. (2). All non-zero elements in \mathbf{W} are set to 0.8. The label matrix \mathbf{Y} is generated as $\mathbf{Y} = \mathbf{X}\mathbf{W} + \epsilon$, where ϵ is a noise matrix with each column generated from $\mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_n)$.

We use the mean square error (MSE) to measure the performance of the estimation on \mathbf{W} , which is defined as

$$\text{MSE}(\mathbf{W}) = \frac{\sum_{i=1}^m (\mathbf{w}_i - \mathbf{w}_i^*)^T \mathbf{X}^T \mathbf{X} (\mathbf{w}_i - \mathbf{w}_i^*)}{mn},$$

where $\mathbf{W}^* = (\mathbf{w}_1^*, \dots, \mathbf{w}_m^*)$ is the ground truth in Figure 2(b). We also use two additional metrics to measure the performance of sparsity recovery: the number of zero coefficients that are estimated to be nonzero (i.e., $\#(Z \rightarrow NZ)$)




Figure 2: Illustrations in synthetic data. (a) Tree used to simulate \mathbf{W} ; (b) The simulated \mathbf{W} ; (c)-(h) Sparse patterns for different features.

Table 1: The performance of various methods over 30 simulations on the synthetic data in terms of mean \pm standard deviation.

		MTFL	DM	MLL	TGGL	Cascade	PTS
$\sigma = 2$	MSE	1.728 \pm 0.156	1.727 \pm 0.156	1.632 \pm 0.158	1.448 \pm 0.188	1.431 \pm 0.137	1.374\pm0.161
	$\#(Z \rightarrow NZ)$	51.333 \pm 5.365	51.367 \pm 5.549	37.533\pm7.655	49.467 \pm 5.770	51.033 \pm 5.014	43.100 \pm 7.801
	$\#(NZ \rightarrow Z)$	0.000\pm0.000	0.000\pm0.000	0.200 \pm 0.407	0.000\pm0.000	0.000\pm0.000	0.000\pm0.000
$\sigma = 3$	MSE	3.814 \pm 0.429	3.804 \pm 0.442	3.669 \pm 0.422	2.850 \pm 0.307	2.815 \pm 0.412	2.758\pm0.322
	$\#(Z \rightarrow NZ)$	50.633 \pm 4.530	50.833 \pm 4.488	40.500\pm7.807	48.600 \pm 5.203	50.500 \pm 4.687	41.467\pm4.273
	$\#(NZ \rightarrow Z)$	0.000\pm0.000	0.000\pm0.000	2.300 \pm 2.103	0.000\pm0.000	0.000\pm0.000	0.000\pm0.000

and the number of nonzero coefficients that are estimated to be zero (i.e., $\#(NZ \rightarrow Z)$).

We generate n samples for training as well as n samples for testing. The tuning parameters in all models are selected via another validation set with n samples. Table 1 shows the performance of all the methods over 30 simulations. As indicated in Table 1, PTS outperforms all the other methods in terms of MSE. All the methods have nice $\#(NZ \rightarrow Z)$, while the MLL method misidentifies some non-zero coefficients to be zero. Moreover, both the MLL and PTS models have lower $\#(Z \rightarrow NZ)$ than the other methods, implying that the PBD models are more likely to obtain sparse solutions. Figure 3 shows the \mathbf{W} 's learned by different methods on the synthetic data when $\sigma = 2$. The black area denotes non-zero coefficients, while the white one corresponds to zero coefficients. Again the PBD models, i.e., the MLL and PTS models, show better recovery of the sparse pattern compared with other methods.

Microarray Data

In this experiment, we report results on microarray data (Wille et al. 2004).¹ The data is a gene expression dataset with microarray data related to isoprenoid biosynthesis in plant organism. All the learning tasks are regression problems which aim to find the cross-talks from the expression levels of 21 genes in the mevalonate pathway (data features)

¹<http://www.ncbi.nlm.nih.gov/pmc/articles/PMC545783/>




Figure 3: \mathbf{W} learned from various methods in synthetic data when $\sigma = 2$. (a) MTFL; (b) DM; (c) MLL; (d) TGGL; (e) Cascade; (f) PTS.

to the expression levels of 18 genes in the plastidial pathway (labels). There are 118 samples. We perform 10 random splits, each of which uses 60%, 20% and 20% samples for training, testing and validation separately. All the variables are log transformed and standardized to zero mean and unit variance. Since there is no tree structure about the task relation, in the first setting we utilize the two-level tree in Figure 1(b) as the priori information. Moreover, we run a hierarchical clustering algorithm similar to (Kim and Xing 2010) on the labels of all tasks to get a tree as the second setting. The tree learned from the hierarchical clustering algorithm is shown in Figure 4, where the leaves denote the tasks (i.e., the genes).

Table 2 gives the average MSE. Except the TGGL and




Figure 4: The tree learned from the hierarchical clustering algorithm in the microarray Data.

Table 2: Average MSE of various methods over 10 random splits on the microarray data.

Methods	Two-level tree	Clustering based tree
MTFL	0.6342±0.0913	0.6342±0.0913
DM	0.6141±0.1104	0.6141±0.1104
MLL	0.6118±0.0908	0.6118±0.0908
TGGL	0.6139±0.0867	0.6220±0.0941
Cascade	0.6112±0.0826	0.6112±0.0826
PTS	0.6091±0.0804	0.5986±0.0885

PTS methods, the other models cannot utilize the given tree structure and thus have the same results in both settings. In the first setting, the DM, MLL, TGGL, Cascade and PTS models are comparable with each other, and all of them outperform the MTFL model under the significance t-test with 95% confidence. In the second setting, the performance of the PTS method has some improvement compared with the first setting, and is better than that of the other methods, implying that the tree structure brings useful information about the task relation for the PTS method. However, the prediction error of the TGGL method which also utilizes the tree structure increases after incorporating the priori information. One reason is that the TGGL and PTS models utilize the tree structure in different ways. The TGGL model imposes a group Lasso regularizer for all the regression coefficients in a subtree rooted at each internal node. In the PTS model, the component coefficient for an internal node only determines whether the subtree rooted at this node should be removed from the solution or not. If not, the sparse pattern can be further determined by its descendants. Thus, the PTS model is more flexible than the TGGL model by using the multi-component PBD technique.

Cifar-100

In this problem, we test the performance of various methods on the Cifar-100 dataset.² This dataset consists of 50000 color images belonging to 100 classes, with 500 images per class. The 100 classes are further grouped into 20 super-classes, from which we could obtain a tree structure over the classes as a priori information. We test the data in both multi-task and multi-class settings. For the multi-task setting we consider the 100 1-vs-rest classification tasks. For each binary classification task, we use all 500 images of each

²<http://www.cs.toronto.edu/~kriz/cifar.html>

Table 3: Average accuracy of various methods over 5 random splits on the Cifar-100 data under different settings.

Methods	Multi-Task	Multi-Class
MTFL	64.45±0.56	14.18±0.46
DM	67.31±0.63	14.48±0.42
MLL	67.74±0.94	15.48±0.41
TGGL	67.80±0.56	-
Cascade	68.94±0.62	16.10±0.45
PTS	70.22±0.52	16.22±0.53

class as the positive set, and 5 examples from each class in the rest of classes as the negative set. We randomly split the data with 60%, 20% and 20% samples for training, testing and validation separately. For the multi-class setting, we just formulate it as multi-class classification problem which can be viewed a special case of the multi-task setting with all tasks sharing the same training data.

The average results over 5 random splits are reported in Table 3. Since the TGGL model is only designed for the multi-output problem as introduced previously, we only report its result in the 1-vs-rest setting. As shown in Table 3, our PTS model has the best classification accuracy in both multi-task and multi-class settings.

Conclusion and Future Work

In this paper, we proposed a probabilistic tree sparsity model in which the sparsity is induced from a product decomposition of the model coefficients into multiple component coefficients. By using Gaussian and Cauchy distributions as the priors on the component coefficients, we developed an efficient EM algorithm to learn the model coefficients. Experimental results on synthetic data and two real-world datasets demonstrate the usefulness of the proposed method.

As one future direction, we will focus on alleviating the assumption that the tree structure over the tasks should be given as a priori information. We are interested in the utilization of some Bayesian models, e.g. Dirichlet process, to identify the tree structure automatically from data. We are also interested in the comparison between the proposed model and some other sparse feature learning methods, e.g. the sparse coding method (Goodfellow, Courville, and Bengio 2012).

Acknowledgement

This work is supported by the National High Technology Research and Development Program of China (No. 2014AA015103), the National Science and Technology Support Plan (No. 2014BAG01B02), Natural Science Foundations of China (No. 61305071), and HKBU FRG2/13-14/039.

References

- Argyriou, A.; Evgeniou, T.; and Pontil, M. 2008. Convex multi-task feature learning. *Machine Learning* 73(3).
- Arnold, B. C., and Brockett, P. L. 1992. On distributions whose component ratios are cauchy. *The American Statistician* 46(1):25–26.

- Bakker, B., and Heskes, T. 2003. Task clustering and gating for Bayesian multitask learning. *The Journal of Machine Learning Research* 4:83–99.
- Candes, E. J.; Wakin, M. B.; and Boyd, S. P. 2008. Enhancing sparsity by reweighted ℓ_1 minimization. *Journal of Fourier Analysis and Applications* 14(5-6):877–905.
- Carvalho, C. M.; Polson, N. G.; and Scott, J. G. 2009. Handling sparsity via the horseshoe. In *International Conference on Artificial Intelligence and Statistics*, 73–80.
- Chen, J.; Liu, J.; and Ye, J. 2012. Learning incoherent sparse and low-rank patterns from multiple tasks. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 5(4):22.
- Chen, J.; Zhou, J.; and Ye, J. 2011. Integrating low-rank and group-sparse structures for robust multi-task learning. In *KDD*.
- Daniel, H.-L.; Jose Miguel, H.-L.; and Pierre, D. 2013. Generalized spike-and-slab priors for bayesian group feature selection using expectation propagation. In *International Conference on Artificial Intelligence and Statistics*.
- Daumé III, H. 2009. Bayesian multitask learning with latent hierarchies. In *UAI*.
- Dempster, A. P.; Laird, N. M.; and Rubin, D. B. 1977. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*.
- Gong, P.; Ye, J.; and Zhang, C. 2012. Robust multi-task feature learning. In *KDD*.
- Goodfellow, I.; Courville, A.; and Bengio, Y. 2012. Large-scale feature learning with spike-and-slab sparse coding. In *ICML*.
- Görnitz, N.; Widmer, C. K.; Zeller, G.; Kahles, A.; Rätsch, G.; and Sonnenburg, S. 2011. Hierarchical multitask structured output learning for large-scale sequence segmentation. In *NIPS*.
- Jacob, L.; Bach, F.; and Vert, J.-P. 2008. Clustered multi-task learning: A convex formulation. In *NIPS*.
- Jalali, A.; Ravikumar, P.; Sanghavi, S.; and Ruan, C. 2010. A dirty model for multi-task learning. In *NIPS*.
- Kim, S., and Xing, E. P. 2010. Tree-guided group lasso for multi-task regression with structured sparsity. In *ICML*.
- Kumar, A., and Daume III, H. 2012. Learning task grouping and overlap in multi-task learning. In *ICML*.
- Lange, K.; Hunter, D. R.; and Yang, I. 2000. Optimization transfer using surrogate objective functions. *Journal of computational and graphical statistics* 9(1):1–20.
- Liu, J.; Ji, S.; and Ye, J. 2009. Multi-task feature learning via efficient $\ell_{2,1}$ -norm minimization. In *UAI*.
- Lozano, A. C., and Swirszcz, G. 2012. Multi-level lasso for sparse multi-task regression. In *ICML*.
- Obozinski, G.; Taskar, B.; and Jordan, M. I. 2006. Multi-task feature selection. *Statistics Department, UC Berkeley, Tech. Rep.*
- Qi, Y. A.; Minka, T. P.; Picard, R. W.; and Ghahramani, Z. 2004. Predictive automatic relevance determination by expectation propagation. In *ICML*.
- Quattoni, A.; Carreras, X.; Collins, M.; and Darrell, T. 2009. An efficient projection for $\ell_{1,\infty}$ regularization. In *ICML*.
- Wille, A.; Zimmermann, P.; Vranová, E.; Fürholz, A.; Laule, O.; Bleuler, S.; Hennig, L.; Prelic, A.; von Rohr, P.; Thiele, L.; et al. 2004. Sparse graphical gaussian modeling of the isoprenoid gene network in arabidopsis thaliana. *Genome Biol* 5(11):R92.
- Wipf, D., and Nagarajan, S. 2010. Iterative reweighted ℓ_1 and ℓ_2 methods for finding sparse solutions. *IEEE Journal of Selected Topics in Signal Processing* 4(2):317–329.
- Xiong, T.; Bi, J.; Rao, B.; and Cherkassky, V. 2006. Probabilistic joint feature selection for multi-task learning. In *SIAM International Conference on Data Mining*.
- Xue, Y.; Liao, X.; Carin, L.; and Krishnapuram, B. 2007. Multi-task learning for classification with Dirichlet process priors. *The Journal of Machine Learning Research* 8:35–63.
- Zhang, Y.; Yeung, D.-Y.; and Xu, Q. 2010. Probabilistic multi-task feature selection. In *NIPS*.
- Zhong, W., and Kwok, J. 2012. Convex multitask learning with flexible task clusters. In *ICML*.
- Zweig, A., and Weinshall, D. 2013. Hierarchical regularization cascade for joint learning. In *ICML*.