

Adaptive Activation Network and Functional Regularization for Efficient and Flexible Deep Multi-Task Learning

Yingru Liu,¹ Xuewen Yang,¹ Dongliang Xie,² Xin Wang,¹
Li Shen,⁴ Haozhi Huang,⁴ Niranjana Balasubramanian³

¹Department of Electrical and Computer Engineering, Stony Brook University, New York, United States

²Institute of Network Technology, Beijing University of Posts and Telecommunications, Beijing, China

³Department of Computer Science, Stony Brook University, New York, United States

⁴Tencent AI Lab, Shenzhen, China

Abstract

Multi-task learning (MTL) is a common paradigm that seeks to improve the generalization performance of task learning by training related tasks simultaneously. However, it is still a challenging problem to search the flexible and accurate architecture that can be shared among multiple tasks. In this paper, we propose a novel deep learning model called Task Adaptive Activation Network (TAAN) that can automatically learn the optimal network architecture for MTL. The main principle of TAAN is to derive flexible activation functions for different tasks from the data with other parameters of the network fully shared. We further propose two functional regularization methods that improve the MTL performance of TAAN. The improved performance of both TAAN and the regularization methods is demonstrated by comprehensive experiments.

Introduction and Related Works

Multi-task learning (MTL), the process of learning to solve multiple tasks at the same time, allows sharing information across related tasks, thereby improving model performance across all the tasks (Caruana 1993; Zamir et al. 2018). Prior MTL studies can be divided into two categories, one is based on optimization (Li et al. 2017; Murugesan and Carbonell 2017), and the other is based on parameter sharing approaches for deep neural networks.

The central problem in MTL is to model both task-specific and task-shared knowledge across different tasks, where the knowledge is encoded in the parameters of the model. Despite the modeling capability of deep learning models, it is still a challenge to automatically determine the optimal splitting between task-specific and task-shared knowledge, which is critical in the design of neural network architecture for MTL and a fundamental problem to solve.

Most existing work of MTLs with deep learning can be categorized into two types: hard-sharing and soft-sharing. As shown in Figure 1 (a), hard-sharing methods (Caruana 1993) define the low-level hidden layers as the knowledge shared across all tasks and the high-level hidden layers as the task-specific knowledge. In the extreme case, all hidden layers are shared by all the tasks. This method is ineffi-

cient, as it requires user to determine the knowledge sharing architecture. It is tedious and almost impossible to manually search for the optimal architecture for MTL when the depth of the neural network is large. On the other hand, soft-sharing methods define independent but identical network structures for all tasks, but learn the task relationship by adding regularization on network parameters (Long et al. 2017; Yang and Hospedales 2017) or inserting connections across networks (Misra et al. 2016; Xiao et al. 2018; Meyerson and Miikkulainen 2018; Ma et al. 2019; Ruder et al. 2019). Although theoretically many soft-sharing methods are able to learn the task relationship, they are not scalable in real-world applications, as the model complexity explodes when the number of tasks increases. Modern neural networks, such as VGG-16 (Simonyan and Zisserman 2014) and ResNet-50 (He et al. 2016), contain over millions of trainable parameters. It is impractical to define an individual network for each task, especially when the number of tasks is large.

In light of the challenge in dividing the neural network architecture between task-shared and task-specific parts among different tasks, we propose an efficient deep learning architecture called Task Adaptive Activation Network (TAAN) to enable flexible and low-cost multi-task learning. In TAAN, all the tasks share the weight and bias parameters of the neural network, thus its complexity is similar to that of a single task model. During the training, TAAN automatically discovers the optimal sharing structure through the learning of task-adaptive activation functions, which avoids the tedious procedures of determining the splitting point of the network. As there are no fixed constraints on the network structure and the whole model can be trained in the end-to-end manner, TAAN is flexible enough to be used either as an individual network or a sub-network in a large system for a large variety of deep learning applications. Based on the functional analysis of the activation functions learnt, We further propose two functional regularization methods to enhance the performance of TAAN. We also prove a relationship between the activation functions and the hidden features. To the best of our knowledge, such a relationship between the method and the hidden features is not studied in the literature work.

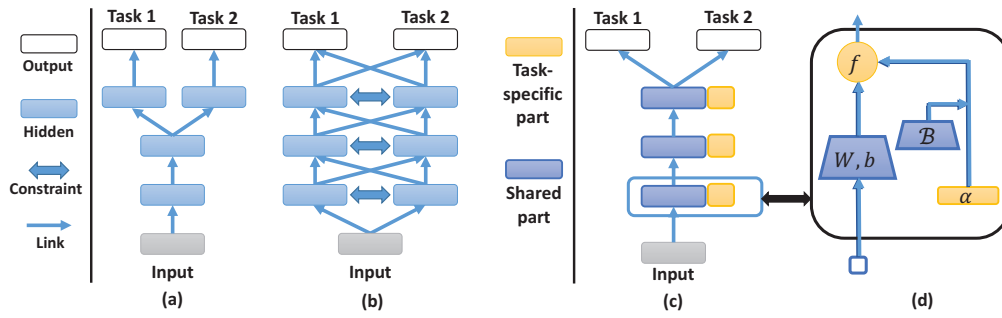


Figure 1: Categories of deep learning MTL: (a) Hard-sharing; (b) Soft-sharing; (c) Task Adaptive Activation Network (Ours); (d) Inner Structure of Adaptive Activation Layer.

Task Adaptive Activation Network

Model Architecture

Weights and biases are major system parameters in many deep learning models. As the simplest MTL architecture, all tasks can share their weights and biases on the hidden layers. Built upon this simple design, to differentiate tasks, we propose to learn different activation functions. Keeping the complexity to be similar to a single-task model, our network architecture is more scalable than the soft-sharing methods, where the number of network components is proportional to the number of tasks.

We call our network architecture Task Adaptive Activation Network (TAAN). As shown in Figure 1 (c), the core structure of TAAN is the concatenation of multiple Adaptive Activation Layers (AAL). For a task t , given the input h_{l-1}^t from either the previous AAL or data input, the output of the l -th AAL is defined by

$$h_l^t = \mathcal{F}_l^t(W_l h_{l-1} + b_l),$$

where the weight and bias parameters $\{W_l, b_l\}$ are shared across tasks. Unlike most deep learning models that only take a fixed activation function at each layer, AAL parameterizes the activation function by a set of basis functions and assigns the specific coordinates to each task. The task-specific activation function at the layer l is defined as

$$\mathcal{F}_l^t(\cdot) = \sum_{i=1}^M \alpha_l^t(i) \mathcal{B}_i(\cdot),$$

where $\alpha_l^t = [\alpha_l^t(1), \dots, \alpha_l^t(M)] \in \mathbb{R}^M$ denotes the coordinates of bases $\{\mathcal{B}_i\}_{i=1}^M$.

As each task has its own coordinate vector $\alpha_l^t \in \mathbb{R}^M$, there is a coordinate matrix $\alpha_l \in \mathbb{R}^{T \times M}$ attached to each AAL hidden layer, where each row of α_l is the transpose of α_l^t , and T is the number of tasks. While the other model parameters are shared for all tasks, the coordinate matrices of the hidden layers control the level of network sharing among multiple tasks. For instance, if tasks 1 and 2 have more shared knowledge at the 1_{st} hidden layer, $\mathcal{F}_1^1(\cdot)$ and $\mathcal{F}_1^2(\cdot)$ have higher similarity, thus α_1^1 and α_1^2 are more similar. On the other hand, if tasks 1 and 2 share less knowledge at the 2_{nd} hidden layer, their activation functions are more

diverse. During the training phase, the coordinate matrices of all hidden layers are optimized to extract both the shared and task-specific knowledge from data.

Choices of Basis Function

In the framework of TAAN, the parameterization of activation functions plays an important role in the network property. In selecting the basis functions, two essential problems should be considered:

1. The activation functions generated by the selected bases should have equivalent or better performance than existing activation functions.
2. There needs a metric to quantify the similarity of activation functions for the function space spanned by the selected bases. Otherwise, the MTL property of TAAN will be unexplainable.

We choose to parameterize $\mathcal{F}_l^t(\cdot)$ with the adaptive piecewise linear (APL) activation unit (Agostinelli et al. 2014):

$$\mathcal{F}_l^t(x) = \max(0, x) + \sum_{i=1}^M \alpha_l^t(i) \max(0, -x + b_i). \quad (1)$$

whose performance has been demonstrated in (Agostinelli et al. 2014; Hou et al. 2017). In our MTL design, we expect that the activation function to be able to evaluate the similarity of knowledge drawn on a neural network layer for two tasks. However, it is nontrivial to measure the similarity of the functions expressed by APL. As $\{\max(0, x), \dots, \max(0, -x + b_M)\}$ are non-orthonormal and unbounded, the inner-product and Euclidean distance of the coordinate vectors $\{\alpha_l^t\}$ do not provide a valid measurement for the learned activation functions. In the next section, we will not only define the solid measurement metrics for the similarity of APLs, but also propose two efficient functional regularization methods to further enhance the performance of TAAN.

Theoretical Analysis and Functional Regularization

Although TAAN can be fully trained by a data-driven optimization, the non-convex optimization process may make

the training performance unstable. As most multi-task learning applications are supposed to be related, knowledge sharing is generally desirable. In this section, we first provide some basic theory to address the challenge of the non-orthogonality and unboundness of the basis functions in TAAN. We then further propose two functional regularization methods to help TAAN to learn more correlated activation functions for different tasks.

Metrics to Measure the Relationship of Activation Functions

In order to understand how TAAN captures the relationship of multiple tasks, it is essential to define effective metrics to measure the difference between two activation functions. As the basis functions of APL are unbounded and non-orthonormal, the coordinate vectors do not reveal much property about the functions. Besides, the commonly-used L^2 norm $\|f\|_2 = \left(\int |f(x)|^2 dx\right)^{\frac{1}{2}}$ and inner product $\langle f, g \rangle = \int f(x)g(x)dx$ are infinite almost everywhere, it is impossible to use them as metrics. Before redefining the finite inner product and norm, we first propose the following lemma, which can be straightly proved by the integral of the Gaussian function.

Lemma 1. *If X is a random variable with Gaussian distribution $p(x) = \mathcal{N}(\mu, \sigma^2)$, the 2_{nd} moments of the outputs of the basis functions $\mathcal{B}_0(x) = \max(0, x)$ and $\mathcal{B}_b(x) = \max(0, -x + b)$ are denoted as*

$$\begin{aligned} \mathbb{E}_{\mu, \sigma}(\mathcal{B}_0^2) &= (\mu^2 + \sigma^2) \left(1 - \Phi\left(-\frac{\mu}{\sigma}\right)\right) + \frac{\mu\sigma}{\sqrt{2\pi}} \exp\left(-\frac{\mu^2}{2\sigma^2}\right), \\ \mathbb{E}_{\mu, \sigma}(\mathcal{B}_{b_i}\mathcal{B}_{b_j}) &= (\mu^2 + \sigma^2 + b_i b_j - (b_i + b_j)\mu) \Phi\left(\frac{\tilde{b} - \mu}{\sigma}\right) \\ &\quad + (b_i + b_j - \mu - \tilde{b}) \frac{\sigma}{\sqrt{2\pi}} \exp\left(-\frac{(\tilde{b} - \mu)^2}{2\sigma^2}\right), \\ \mathbb{E}_{\mu, \sigma}(\mathcal{B}_0\mathcal{B}_b) &= \begin{cases} (b\mu - \mu^2 - \sigma^2) \left(\Phi\left(\frac{b - \mu}{\sigma}\right) - \Phi\left(-\frac{\mu}{\sigma}\right)\right) \\ \quad + \frac{\sigma\mu}{\sqrt{2\pi}} \exp\left(-\frac{(b - \mu)^2}{2\sigma^2}\right) \\ \quad + \frac{\sigma(b - 3\mu)}{\sqrt{2\pi}} \exp\left(-\frac{\mu^2}{2\sigma^2}\right), & \text{if } b > 0 \\ 0, & \text{if } b \leq 0 \end{cases} \end{aligned}$$

where $\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x \exp(-\frac{t^2}{2}) dt$ is the CDF of the standard Gaussian distribution and $\tilde{b} = \min(b_i, b_j)$. \square

With the introduction of the Gaussian density as the weighted factor, the 2_{nd} moments in Lemma 1 are finite. Given Lemma 1, it is straightforward to define an finite inner product and distance of two activation functions as the integration weighted by a Gaussian Mixture Model (GMM).

Definition 1 (Inner product induced by GMM). *Given a Gaussian Mixture Model (GMM) as $\mathcal{G}(\cdot) = \sum \pi_i \mathcal{N}(\mu_i, \sigma_i^2)$, $\pi_i \geq 0$, $\sum \pi_i = 1$, the inner product of two APL functions expressed by Eq. (1) is defined as*

$$\langle \mathcal{F}_l^{t_1}, \mathcal{F}_l^{t_2} \rangle = \sum_i \pi_i \langle \mathcal{F}_l^{t_1}, \mathcal{F}_l^{t_2} \rangle_{\mu_i, \sigma_i}, \quad (2)$$

where

$$\begin{aligned} \langle \mathcal{F}_l^{t_1}, \mathcal{F}_l^{t_2} \rangle_{\mu, \sigma} &= \int_{-\infty}^{+\infty} \mathcal{F}_l^{t_1}(x) \mathcal{F}_l^{t_2}(x) \mathcal{N}(\mu, \sigma^2) dx \\ &= \mathbb{E}_{\mu, \sigma}(\mathcal{B}_0^2) + \sum_{i=1}^M \left(\alpha_l^{t_1}(i) + \alpha_l^{t_2}(i)\right) \mathbb{E}_{\mu, \sigma}(\mathcal{B}_0 \mathcal{B}_{b_i}) \\ &\quad + \sum_{i=1}^M \sum_{j=1}^M \alpha_l^{t_1}(i) \alpha_l^{t_2}(j) \mathbb{E}_{\mu, \sigma}(\mathcal{B}_{b_i} \mathcal{B}_{b_j}). \end{aligned}$$

\square

Definition 2 (Distance between \mathcal{F}_l^t). *Given a Gaussian Mixture Model (GMM) as $\mathcal{G}(\cdot) = \sum \pi_i \mathcal{N}(\mu_i, \sigma_i^2)$, $\pi_i \geq 0$, $\sum \pi_i = 1$, the distance between two APL functions expressed by Eq. (1) is defined as*

$$\begin{aligned} d(\mathcal{F}_l^{t_1}, \mathcal{F}_l^{t_2}) &= \int (\mathcal{F}_l^{t_1}(x) - \mathcal{F}_l^{t_2}(x))^2 \mathcal{G}(x) dx \\ &= \sum_i \pi_i \sum_{i=1}^M \sum_{j=1}^M \left(\alpha_l^{t_1}(i) - \alpha_l^{t_2}(i)\right) \left(\alpha_l^{t_1}(j) - \alpha_l^{t_2}(j)\right) \mathbb{E}_{\mu, \sigma}(\mathcal{B}_{b_i} \mathcal{B}_{b_j}). \end{aligned} \quad (3)$$

\square

Definitions 1 and 2 exploit GMM to represent the optimal weight function for the inner products, as it is general and can universally approximate the possible weight functions that define a finite inner product for Eq. (1). \mathcal{G} can be either defined by prior knowledge or trained together with TAAN. With the valid definition of the inner product, it is straightforward to define the L^2 norm of two APL functions as follows.

$$\|\mathcal{F}_l^t\|_2 = \left(\langle \mathcal{F}_l^t, \mathcal{F}_l^t \rangle\right)^{\frac{1}{2}}, \quad (4)$$

Eqs. (2), (3) and (4) are effective tools to analyze the learned activation functions. Eq. (3) is a measurement of the difference between two task-specific activation functions. Given the inner product, a cosine similarity can also be defined to evaluate the correlation of the functions. At the end of this paper, we show that Eq. (2) and Eq. (4) are helpful in analyzing the relation between the activation function and the hidden feature. Although this relation does not directly influence the design of our methods, it exposes how the learned task-adaptive activation functions modulate the extracted features to capture the task relationship.

Functional Regularization

For each layer of TAAN, the coordinate matrix $\alpha_l \in \mathbb{R}^{T \times M}$ can be learned directly from the training data. As the tasks in MTL are generally considered to be related, it is reasonable to encourage sharing more than splitting. This insight can be incorporated into TAAN by introducing regularization term on α_l during training. We propose two functional regularization methods to further enhance the performances of TAAN.

We first define a baseline method that is introduced from literature work. In the experiment section, we will show that this baseline has little benefit for the MTL performance of TAAN.

Approach 1 (Baseline: Trace-Norm). *The first regularization hypothesis is that the matrix α_l is low-rank, as the tasks in MTL often have high correlation. The low-rank assumption of the model parameters is often considered in MTL literature work (Argyriou, Evgeniou, and Pontil 2008; Pong et al. 2010). Thus, we introduce a regularization term to α_l :*

$$\mathcal{L}_{tn}(\alpha_l) = \text{trace}(\sqrt{\alpha_l \alpha_l^T}),$$

where $\sqrt{\cdot}$ denotes the square root of matrix. The trace-norm $\text{trace}(\cdot)$ is proven to be the convex envelope of the matrix rank (Andersson, Carlsson, and Olsson 2017). \square

The proposed regularization term \mathcal{L}_{tn} corresponds to the hypothesis on linear dependency. As the relationship between the coordinate α_l and the corresponding function is ambiguous, \mathcal{L}_{tn} may not introduce beneficial structure into the activation functions.

In this paper, we propose two functional regularization methods that penalize the dissimilarity of the task-specific activation functions.

Approach 2 (Functional regularization by cosine similarity). *Given the definition of inner product as Eq. (2), the similarity of two task-specific activation functions can be defined by the cosine similarity, which is computed as*

$$\mathcal{L}_{cos}(\alpha_l) = -\frac{1}{T^2} \sum_{ij} C_{ij}(\alpha_l),$$

$$C_{ij}(\alpha_l) = \frac{\langle \mathcal{F}_l^i, \mathcal{F}_l^j \rangle}{\sqrt{\langle \mathcal{F}_l^i, \mathcal{F}_l^i \rangle \langle \mathcal{F}_l^j, \mathcal{F}_l^j \rangle}},$$

where T is the number of tasks. \square

\mathcal{L}_{cos} measures the correlation between two functions, but neglects the difference between the outputs of them. To consider the variance of function output, we propose a straight regularization method that directly regularizes the distance Eq. (3).

Approach 3 (Functional regularization by distance). *Given the coordinate matrix α_l for the l_{th} layer of network, we compress the distance function Eq. (3) between task-specific activation functions with the following regularization*

$$\mathcal{L}_{dis}(\alpha_l) = \frac{1}{T^2} \sum_{ij} \mathcal{D}_{ij}(\alpha_l), \quad \mathcal{D}_{ij}(\alpha_l) = d^2(\mathcal{F}_l^i, \mathcal{F}_l^j).$$

\square

Given the definition of regularization terms, the training loss of a TAAN with L task-specific activation layers becomes

$$\mathcal{L}_{total} = \sum_{t=1}^T \mathcal{L}_t(\mathcal{M}_t, \{x_i^t, y_i^t\}) + c \sum_{l=1}^L \mathcal{L}_{MTL}(\alpha_l),$$

where $\mathcal{L}_{MTL} \in \{\mathcal{L}_{tn}, \mathcal{L}_{dis}, \mathcal{L}_{cos}\}$ and c is the regularization coefficient. During the training process, \mathcal{L}_{dis} and \mathcal{L}_{cos} requires GMM \mathcal{G} in Definition 1. From our observations of several dimensions for input and the pre-activation of each hidden layer in the experiment, the shapes of their histograms follow the density curves of Gaussian distribution. In this paper, we fix \mathcal{G} as a standard Gaussian distribution for simplicity.

Experiments

In this section, we compare the performance of our model with several methods on two challenging datasets, Youtube-8M (Abu-El-Hajja et al. 2016) and Omniglot (Lake, Salakhutdinov, and Tenenbaum 2015).

Multi-Domain Multi-Label Classification

We first conduct experiments on Youtube-8M, a large dataset that consists of over 6.1 billion of Youtube videos. Each video has multiple labels from a vocabulary of 3800 topical entities, which can be further grouped into 24 top-level categories.

Experiment Setting To create an MTL experiment, we consider each top-level category as a specific domain. For each domain, we have to define a multi-label classifier to recognize various attributes of the data. As some domains have too small amount of data, we use only the top 16 for our experiments. The task IDs and their corresponding domains are given in Figure 2. We used the training set provided in the original dataset, but split its validation set equally into two parts, one used for our own validation and the other for testing. This setting is the same as (Ma et al. 2019). We use the mean average precision at 10 (mAP%10) as the metric in the performance evaluation.

Model Configuration All the models are composed of three feed-forward layers, whose output sizes are 1024. For TAAN, the number of basis functions is 64. During the training, we choose the Adam optimizer (Kingma and Ba 2014; Zou et al. 2019), with $\beta_1 = 0.9$ and $\beta_2 = 0.98$. The initial learning rate is 0.0001 and the batch size is set to 256 for each task. We compare our method with the following three deep multi-task learning models:

- **MRN (Long et al. 2017):** Multilinear Relationship Network (MRN) is a soft-sharing method. For each layer of the neural network, MRN concatenates the task-specific weight matrix into a 3-D tensor and posits Tensor Gaussian distribution for it. Multi-task learning is achieved by introducing Bayesian regularization term and learning the covariance matrices of the Gaussian distribution.
- **DMTRL (Yang and Hospedales 2017):** DMTRL is another soft-sharing method. It assumes that weight tensors have low-rank structure and parameterizes the weight tensors by several tensor decomposition methods. Inherited from the tensor decomposition methods, DMTRL requires user-predefined ranks of the weights, which correspond to the independent groups of tasks.
- **Soft-Order (Meyerson and Miikkulainen 2018):** The Soft Ordering method defines a set of identical hidden layers and achieves multi-task learning by shuffling the order of the hidden layers to build the network for each task. There are two strong constraints of the Soft Ordering method. First, the sizes of the hidden layers in the neural network should be identical. Second, it can not simultaneously shuffle the feed-forward layers and the convolutional layers.

Single-task learning (STL) and hard-sharing model (Hard-Share) are also included as baselines.

Table 1: The number of network parameters and classification performance on the Youtube-8M dataset

	# of params	Task mAP%10															Mean mAP%10	
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15		16
Benchmarks																		
STL	56.1 M	.683	.807	.835	.689	.781	.431	.559	.831	.596	.779	.551	.916	.505	.383	.699	.759	0.675
Hard-Share	6.93 M	.770	.838	.870	.806	.824	.598	.676	.846	.754	.808	.716	.940	.629	.680	.806	.839	0.775
Soft-Order	6.93 M	.785	.845	.877	.819	.825	.611	.683	.848	.764	.809	.727	.943	.648	.694	.819	.838	0.783
MRN _t	56.1 M	.664	.806	.829	.608	.778	.413	.502	.828	.576	.775	.542	.916	.383	.355	.687	.695	0.647
MRN _{full}	61.6 M	.670	.804	.828	.607	.777	.416	.506	.828	.578	.775	.543	.916	.383	.355	.688	.696	0.648
DMTRL-Tucker	39.1 M	.660	.817	.812	.651	.794	.473	.490	.793	.569	.742	.594	.944	.475	.475	.697	.710	0.669
DMTRL-LAF	20.1 M	.851	.869	.890	.875	.841	.691	.757	.873	.839	.854	.809	.948	.746	.792	.875	.896	0.838
DMTRL-TT	46.2 M	.860	.874	.893	.888	.846	.710	.778	.879	.849	.860	.822	.950	.770	.807	.881	.903	0.848
TAAN (Ours)																		
TAAN	6.93 M	.882	.896	.910	.902	.857	.739	.804	.879	.859	.828	.808	.934	.717	.774	.833	.848	0.842
TAAN + \mathcal{L}_{tn}	6.93 M	.741	.824	.861	.779	.813	.572	.636	.831	.711	.785	.681	.937	.581	.637	.783	.809	0.749
TAAN + \mathcal{L}_{cos}	6.93 M	.896	.906	.915	.915	.859	.769	.830	.885	.879	.843	.828	.937	.756	.805	.854	.876	0.860
TAAN + \mathcal{L}_{dis}	6.93 M	.889	.899	.912	.910	.859	.752	.820	.886	.873	.836	.821	.933	.731	.789	.845	.863	0.851

Results We conducted a set of experiments on Youtube-8M dataset and reported the mAP%10 results for each task in Table 1.

- **MRN:** According to Table 1, MRN’s performance can not benefit from its Bayesian learning on Tensor Gaussian distribution. This may be because its learning of the covariance matrices is inefficient, which causes the inaccuracy in the modeling of the task relationship. The SVD operation required for covariance computing also makes the training process slow. Furthermore, serious negative transferring is observed for MRN, as its scores are even worse than a single task model.
- **Soft-Order:** In our experiment, the Soft-Order method can only be used in the last two layers of the neural network. While the layer re-ordering may not efficiently capture task relation, its improvement with respect to the hard-sharing model is not significant.
- **DMTRL:** The performance of DMTRL depends on the tensor decomposition method. Serious negative transferring is observed for Tucker Decomposition. DMTRL-LAF and DMTRL-TT achieve the best performance among the benchmark models. The performance gain over TAAN is only observed for 6 tasks.
- **TAAN:** TAAN significantly outperforms existing methods, when the number of parameters is as small as a simple hard-sharing model. It achieves the highest mAP%10 for the first 9 tasks and task 11. For the remaining 6 tasks, TAAN achieves comparable results with state-of-the-art model DMTRL-TT, with much fewer parameters (6.93 million compared with 46.2 million) and faster training and inference speeds (Table 2). As DMTRL is trained with the assumption of the full knowledge of the rank of the weight tensor, its performance may be further compromised when the assumed rank is inappropriate.

Furthermore, our proposed functional regularization methods improve the performance of TAAN from 0.842 to 0.860 and 0.851. \mathcal{L}_{cos} and \mathcal{L}_{dis} improve the performance of TAAN on all the tasks. Instead, the baseline trace-norm approach is harmful for the model performance. Although \mathcal{L}_{cos} has better performance than \mathcal{L}_{dis}

in Youtube-8M dataset, the next experiment on multi-domain alphabet classification shows that the performance variance of \mathcal{L}_{cos} is larger than \mathcal{L}_{dis} .

Table 2: Training and inference speed evaluation on Youtube-8M (Note: the dimension of hidden layer is 512)

Model	training time↓	inference time↓
MRN _t	156.1 s	8.83 s
MRN _{full}	157.6 s	8.83 s
DMTRL-Tucker	66.02 s	1.53 s
DMTRL-LAF	24.96 s	0.30 s
DMTRL-TT	52.75 s	1.44 s
TAAN	20.05 s	0.88 s
TAAN + \mathcal{L}_{tn}	22.86 s	0.88 s
TAAN + \mathcal{L}_{cos}	40.34 s	0.88 s
TAAN + \mathcal{L}_{dis}	38.35 s	0.88 s

Scalability and Speed TAAN has much fewer parameters than MRN and DMTRL. We further evaluate the training and inference speeds of these models. The training time (per 100 iterations) and inference time (per 100 iterations) are shown in Table 2. Compared with DMTRL-TT, both the training and inference speeds of TAAN are faster. DMTRL-LAF has the fastest inference speed, but its performance is compromised. The use of regularization increases the training and inference time of TAAN. The computation of the basis functions also increases the inference time of TAAN. Acceleration may be achieved by adding L1 regularization and pruning into its coordinate matrices.

Visualization We visualize the distances of the activation functions to analyze the architecture of the best trained TAAN. For each hidden layer of TAAN, we use a $T \times T$ distance matrix to measure the difference between the task-specific activation functions on that layer. According to the results displayed in Figure 2, TAAN is able to capture the complicated knowledge sharing for the tasks on the Youtube-8M dataset. For instance, domain “Food & Drink”

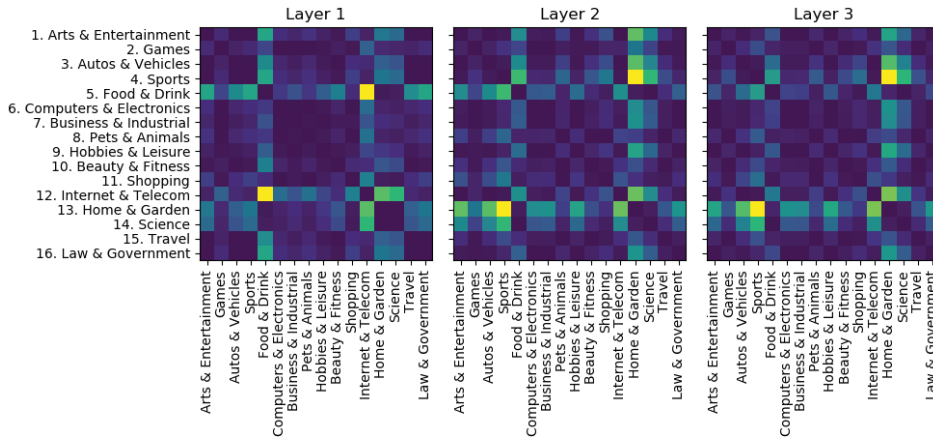


Figure 2: Distance matrices of the activation functions in TAAN. Light colors denote less similarity.

shares all the hidden layers with domain “*Home & Garden*”. TAAN also discover that the domains “*Food & Drink*” and “*Internet & Telecom*” are the most unrelated, as the distances between their activation functions are always high.

Multi-Domain Alphabet Classification

Model performance To further verify the effect of functional regularization, we conduct an auxiliary experiment on Omniglot (Lake, Salakhutdinov, and Tenenbaum 2015) dataset, which contains 1623 different handwritten characters from 50 different alphabets. We define 50 tasks, each classifies characters from one alphabet.

Table 3: Model performance in Omniglot dataset

Model	Parameters	Accuracy
DMTRL-Tucker	--	68.89 %
DMTRL-LAF	--	66.63 %
DMTRL-TT	--	69.39 %
Soft-Order	1,478 K	75.89 %
TAAN (Ours)	1,380 K	84.11 %

We first search for a TAAN architecture, that has comparable model complexity and accuracy with the state-of-the-art models. The built TAAN contains 4 convolutional and 1 dense layers. Each layer is set as AAL, where the number of bases is 32. The classification performance is shown in Table 3. TAAN outperforms the second-best model (Soft-Order) and increases the accuracy from 75.89 % to 84.11 %.

The test accuracies of TAAN with respect to the proposed functional regularizations are presented in Figure 3. Only \mathcal{L}_{dis} can enhance the model performance, as it is the most explicit measurement to evaluate the output difference of two activation functions. Although \mathcal{L}_{cos} measures the function difference from the angular perspective, the absolute variance of the function outputs plays a more important and fundamental role in the MTL performance of TAAN. The trace-norm regularization \mathcal{L}_{tn} is extremely harmful for the model performance, which indicates that the direct hypoth-

esis on the coordinates is not beneficial, as the bases are not orthogonal.




Figure 3: Test accuracy of TAAN with different functional regularization and coefficient on Omniglot dataset

Visualization An example of distance matrices of the 4 convolutional layers and the single dense layer are displayed in Figure 4. The matrices of *conv1* to *conv3* consist of a large number of small values, indicating that most tasks share the lower 3 layers of the network. This training result coincides with the previous study, which shows that lower layers are more transferable and sharable across multiple tasks for multi-domain alphabet classification (Yang and Hospedales 2017). As higher layers generally extract task-specific features, the distances between the task-specific activations become larger in *conv4* and *dense1*.

In addition, different task pairs share various levels of knowledge, and our proposed model can flexibly and automatically determine the sharing level in different cases. Most tasks share the first 3 layers, while some others share the whole 5 layers. As a special case, one task does not share any hidden layers with the other ones. The learned activation functions of this task are highly different from the other activation functions learned. The corresponding distances are represented as the light-color crosses in all the matrices. This

specific task is defined as the letter recognition for the alphabet “Brallie”. Brallie is written by dots while other alphabets (tasks) are written mainly by strokes. During the training, our model specifies the characteristics of Brallie and assigns a specific set of activation functions for the corresponding task.




Figure 4: An example of the distance matrix for each hidden layers. Light colors denote less similarity.

Relation between Activation Functions and Hidden Features

In deep neural network, the output of each hidden layer is considered as a disentangled representation or feature of the data information, which is essential for computing the output of each task. As TAAN computes the hidden features through the task-adaptive activation functions, it is important to study how the learned activation functions modulate the hidden features to learn the domain knowledge. Given the framework of TAAN, we are able to prove a relation between the activation functions and the hidden features. Under mild assumption, the following inequalities hold.

Lemma 2. For the l_{th} hidden layer of TAAN, we assume:

- if $l = 1$, there exists a GMM \mathcal{G}_1 with diagonal covariance matrix in each component and a positive scalar C_1 such that the pre-activation $a_1 = W_1 x + b_1$ satisfies $P(a_1) \leq C_1 \mathcal{G}(a_1)$, where x is the data input.
- if $l > 1$, for each task t , there exists a GMM \mathcal{G}_l^t with diagonal covariance matrix in each component and a positive scalar C_l^t such that the pre-activation $a_l^t = W_l h_{l-1}^t + b_l$ satisfies $P(a_l^t) \leq C_l^t \mathcal{G}_l^t(a_l^t)$.

Then, the following inequalities hold for $l = 1$:

$$\mathbb{E}((h_1^{t_1})^T h_1^{t_2}) \leq C_1 \sum_{n=1}^{N_1} \langle F_1^{t_1}, F_1^{t_2} \rangle_{\mathcal{G}_{1,n}},$$

$$\mathbb{E}(\|h_1^{t_1} - h_1^{t_2}\|^2) \leq C_1 \sum_{n=1}^{N_1} d_{\mathcal{G}_{1,n}}(\mathcal{F}_1^{t_1}, \mathcal{F}_1^{t_2}),$$

where $\langle \cdot, \cdot \rangle_{\mathcal{G}_{1,n}}$ is the inner-product in Eq. (2) and $d_{\mathcal{G}_{1,n}}(\cdot, \cdot)$ is the distance in Eq. (3).

For $l > 1$, we have the following inequalities:

$$\mathbb{E}((h_l^{t_1})^T h_l^{t_2}) \leq \frac{C_l^{t_1} (1 + \|a_l^{t_1}\|_1)}{2} \sum_{n=1}^{N_l} \|F_l^{t_1}\|_{\mathcal{G}_{l,n}^{t_1}}^2$$

$$+ \frac{(1 + \|a_l^{t_1}\|_1)}{2} \lambda_{\max} \mathbb{E}(\|h_l^{t_1} - h_l^{t_2}\|^2)$$

$$+ C_l^{t_1} \sum_{n=1}^{N_l} \langle F_l^{t_1}, F_l^{t_2} \rangle_{\mathcal{G}_{l,n}^{t_1}}, \quad (5)$$

$$\mathbb{E}(\|h_l^{t_1} - h_l^{t_2}\|^2) \leq 2C_l^{t_1} \sum_{n=1}^{N_l} d(F_l^{t_1}, F_l^{t_2})_{\mathcal{G}_{l,n}^{t_1}}$$

$$+ 2C_l^{t_1} (1 + \|a_l^{t_1}\|_1)^2 \lambda_{\max} \mathbb{E}(\|h_{l-1}^{t_1} - h_{l-1}^{t_2}\|^2), \quad (6)$$

where N_l the output dimension, $\|\cdot\|_1$ is the L_1 norm of vector, and $\|\cdot\|_{\mathcal{G}_{l,n}^{t_1}}$ is the norm defined in Eq. (4), and λ_{\max} is the maximum eigenvalue of $W_l^T W_l$. \square

To the best of our knowledge, our paper is the first deep multi-task learning work that attempts to connect the proposed method to the hidden feature on deep learning MTL. According to Eq. (5), the inner product of the task-specific output vectors of a hidden layer are bounded by the inner-product of the activation functions and the difference of input. During the training, TAAN can learn to either shrink or expand the correlation between task-specific hidden features. Moreover, the formulation of Lemma 2 demonstrates that it is effective to use GMM to develop the valid measurements for the activation functions. Similarly, the distance between the hidden feature is also controlled by the distance of the activation function and the difference of input.

Conclusion and Discussion

In this paper, we propose an efficient and flexible deep multi-task learning framework, TAAN, which can automatically explore the task relationship by learning a set of task-adaptive activation functions. We also propose two functional regularization methods to enhance the performance of TAAN. Experiment results demonstrate that TAAN is more accurate and flexible than existing soft sharing methods. Furthermore, TAAN can explicitly show the relationship among tasks through the similarity level of the parameters of the activation functions learnt for different tasks. This can serve a base for the advancing of the machine learning field.

Acknowledgments

The work is supported by the National Science Foundation (NSF) under the grant ECCS 1731238.

References

- Abu-El-Hajja, S.; Kothari, N.; Lee, J.; Natsev, P.; Toderici, G.; Varadarajan, B.; and Vijayanarasimhan, S. 2016. Youtube-8m: A large-scale video classification benchmark. *CoRR* abs/1609.08675.
- Agostinelli, F.; Hoffman, M. D.; Sadowski, P. J.; and Baldi, P. 2014. Learning activation functions to improve deep neural networks. *CoRR* abs/1412.6830.

- Andersson, F.; Carlsson, M.; and Olsson, C. 2017. Convex envelopes for fixed rank approximation. *Optimization Letters* 11(8):1783–1795.
- Argyriou, A.; Evgeniou, T.; and Pontil, M. 2008. Convex multi-task feature learning. *Machine Learning* 73(3):243–272.
- Caruana, R. 1993. Multitask learning: A knowledge-based source of inductive bias. In *The 10th International Conference on Machine Learning*, 41–48.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.
- Hou, L.; Samaras, D.; Kurc, T.; Gao, Y.; and Saltz, J. 2017. ConvNets with Smooth Adaptive Activation Functions for Regression. In *The 20th International Conference on Artificial Intelligence and Statistics*, volume 54, 430–439.
- Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization. *CoRR* abs/1412.6980.
- Lake, B. M.; Salakhutdinov, R.; and Tenenbaum, J. B. 2015. Human-level concept learning through probabilistic program induction. *Science* 350(6266):1332–1338.
- Li, C.; Yan, J.; Wei, F.; Dong, W.; Liu, Q.; and Zha, H. 2017. Self-paced multi-task learning. In *AAAI Conference on Artificial Intelligence*.
- Long, M.; Cao, Z.; Wang, J.; and Yu, P. S. 2017. Learning multiple tasks with multilinear relationship networks. In *Advances in Neural Information Processing Systems 30*, 1594–1603.
- Ma, J.; Zhao, Z.; Chen, J.; Li, A.; Hong, L.; and Chi, E. H. 2019. SNR: Sub-network routing for flexible parameter sharing in multi-task learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 216–223.
- Meyerson, E., and Miikkulainen, R. 2018. Beyond shared hierarchies: Deep multitask learning through soft layer ordering. In *International Conference on Learning Representations*.
- Misra, I.; Shrivastava, A.; Gupta, A.; and Hebert, M. 2016. Cross-stitch networks for multi-task learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 3994–4003.
- Murugesan, K., and Carbonell, A. 2017. Self-paced multitask learning with shared knowledge. In *The 26th International Joint Conference on Artificial Intelligence*, 2522–2528.
- Pong, T.; Tseng, P.; Ji, S.; and Ye, J. 2010. Trace norm regularization: Reformulations, algorithms, and multi-task learning. *SIAM Journal on Optimization* 20(6):3465–3489.
- Ruder, S.; Bingel, J.; Augenstein, I.; and Søgaard, A. 2019. Latent multi-task architecture learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 4822–4829.
- Simonyan, K., and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *CoRR* abs/1409.1556.
- Xiao, L.; Zhang, H.; Chen, W.; Wang, Y.; and Jin, Y. 2018. Learning what to share: Leaky multi-task network for text classification. In *The 27th International Conference on Computational Linguistics*, 2055–2065.
- Yang, Y., and Hospedales, T. M. 2017. Deep multi-task representation learning: A tensor factorisation approach. In *The 5th International Conference on Learning Representations*.
- Zamir, A. R.; Sax, A.; Shen, W.; Guibas, L. J.; Malik, J.; and Savarese, S. 2018. Taskonomy: Disentangling task transfer learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Zou, F.; Shen, L.; Jie, Z.; Zhang, W.; and Liu, W. 2019. A sufficient condition for convergences of adam and rmsprop. In *IEEE Conference on Computer Vision and Pattern Recognition*.