

Branched Multi-Task Networks: Deciding What Layers To Share

Simon Vandenhende¹
simon.vandenhende@kuleuven.be
Stamatios Georgoulis²
stamatios.georgoulis@vision.ee.ethz.ch
Bert De Brabandere¹
bert.debrabandere@esat.kuleuven.be
Luc Van Gool^{1,2}
vangool@vision.ee.ethz.ch

¹ PSI-ESAT
KU Leuven
Leuven, Belgium
² CVL/TRACE
ETH Zurich
Zurich, Switzerland

Abstract

In the context of multi-task learning, neural networks with branched architectures have often been employed to jointly tackle the tasks at hand. Such ramified networks typically start with a number of shared layers, after which different tasks branch out into their own sequence of layers. Understandably, as the number of possible network configurations is combinatorially large, deciding what layers to share and where to branch out becomes cumbersome. Prior works have either relied on ad hoc methods to determine the level of layer sharing, which is suboptimal, or utilized neural architecture search techniques to establish the network design, which is considerably expensive. In this paper, we go beyond these limitations and propose an approach to automatically construct branched multi-task networks, by leveraging the employed tasks' affinities. Given a specific budget, i.e. number of learnable parameters, the proposed approach generates architectures, in which shallow layers are task-agnostic, whereas deeper ones gradually grow more task-specific. Extensive experimental analysis across numerous, diverse multi-tasking datasets shows that, for a given budget, our method consistently yields networks with the highest performance, while for a certain performance threshold it requires the least amount of learnable parameters.

1 Introduction

Deep neural networks are usually trained to tackle different tasks in isolation. Humans, in contrast, are remarkably good at solving a multitude of tasks concurrently. Biological data processing appears to follow a multi-tasking strategy too; instead of separating tasks and solving them in isolation, different processes seem to share the same early processing layers in the brain – see e.g. V1 in macaques [15]. Drawing inspiration from such observations, deep learning researchers began to develop multi-task networks with branched architectures.

As a whole, multi-task networks [6] seek to improve generalization and processing efficiency through the joint learning of related tasks. Compared to the typical learning of separate deep neural networks for each of the individual tasks, multi-task networks come

with several advantages. First, due to their inherent layer sharing [14, 20, 21, 26, 31], the resulting memory footprint is typically substantially lower. Second, as features in the shared layers do not need to be calculated repeatedly for the different tasks, the overall inference speed is often higher [31, 34]. Finally, multi-task networks may outperform their single-task counterparts [20, 32, 43, 44]. Evidently, there is merit in utilizing multi-task networks.

When it comes to designing them, however, a significant challenge is to decide on the layers that need to be shared among tasks. Assuming a hard parameter sharing setting¹, the number of possible network configurations grows quickly with the number of tasks. As a result, a trial-and-error procedure to define the optimal architecture becomes unwieldy. Resorting to neural architecture search [11] techniques is not a viable option too, as in this case, the layer sharing has to be jointly optimized with the layers types, their connectivity, etc., rendering the problem considerably expensive. Instead, researchers have recently explored more viable alternatives, like routing [39], stochastic filter grouping [5], and feature partitioning [35], which are, however, closer to the soft parameter sharing setting. Previous works on hard parameter sharing opted for the simple strategy of sharing the initial layers in the network, after which all tasks branch out simultaneously. The point at which the branching occurs is usually determined ad hoc [14, 20, 43]. This situation hurts performance, as a suboptimal grouping of tasks can lead to the sharing of information between unrelated tasks, known as *negative transfer* [47].

In this paper, we go beyond the aforementioned limitations and propose a novel approach to decide on the degree of layer sharing between multiple visual recognition tasks in order to eliminate the need for manual exploration. To this end, we base the layer sharing on measurable levels of *task affinity* or *task relatedness*: two tasks are strongly related, if their single task models rely on a similar set of features. [46] quantified this property by measuring the performance when solving a task using a variable sets of layers from a model pretrained on a different task. However, their approach is considerably expensive, as it scales quadratically with the number of tasks. Recently, [10] proposed a more efficient alternative that uses representation similarity analysis (RSA) to obtain a measure of task affinity, by computing correlations between models pretrained on different tasks. Given a dataset and a number of tasks, our approach uses RSA to assess the task affinity at arbitrary locations in a neural network. The task affinity scores are then used to construct a branched multi-task network in a fully automated manner. In particular, our task clustering algorithm groups similar tasks together in common branches, and separates dissimilar tasks by assigning them to different branches, thereby reducing the negative transfer between tasks. Additionally, our method allows to trade network complexity against task similarity. We provide extensive empirical evaluation of our method, showing its superiority in terms of multi-task performance vs computational resources.

2 Related work

Multi-task learning. Multi-task learning (MTL) [6, 41] is associated with the concept of jointly learning multiple tasks under a single model. This comes with several advantages, as described above. Early work on MTL often relied on sparsity constraints [4, 19, 27, 30, 45]

¹In this setting, the input is first encoded through a stack of shared layers, after which tasks branch out into their own sequence of task-specific layers [14, 20, 21, 31, 43]. Alternatively, a set of task-specific networks can be used in conjunction with a feature sharing mechanism [26, 33, 42]. The latter approach is termed soft parameter sharing in the literature.

to select a small subset of features that could be shared among all tasks. However, this can lead to negative transfer when not all tasks are related to each other. A general solution to this problem is to cluster tasks based on prior knowledge about their similarity or relatedness [1, 3, 12, 22, 48].

In the deep learning era, MTL models can typically be classified as utilizing soft or hard parameter sharing. In soft parameter sharing, each task is assigned its own set of parameters and a feature sharing mechanism handles the cross-task talk. Cross-stitch networks [33] softly share their features among tasks, by using a linear combination of the activations found in multiple single task networks. Sluice networks [42] extend cross-stitch networks and allow to learn the selective sharing of layers, subspaces and skip connections. In a different vein, multi-task attention networks [26] use an attention mechanism to share a general feature pool amongst task-specific networks. In general, MTL networks using soft parameter sharing are limited in terms of scalability, as the size of the network tends to grow linearly with the number of tasks.

In hard parameter sharing, the parameter set is divided into shared and task-specific parameters. MTL models using hard parameter sharing are often based on a generic framework with a shared off-the-shelf encoder, followed by task-specific decoder networks [7, 20, 34, 43]. Multilinear relationship networks [29] extend this framework by placing tensor normal priors on the parameter set of the fully connected layers. [14] proposed the construction of a hierarchical network, which predicts increasingly difficult tasks at deeper layers. A limitation of the aforementioned approaches is that the branching points are determined ad hoc, which can easily lead to negative transfer if the predefined task groupings are suboptimal. In contrast, in our branched multi-task networks, the degree of layer sharing is automatically determined in a principled way, based on task affinities.

Our work bears some similarity to fully-adaptive feature sharing [31] (FAFS), which starts from a thin network where tasks initially share all layers, but the final one, and dynamically grows the model in a greedy layer-by-layer fashion. Task groupings, in this case, are decided on the probability of concurrently simple or difficult examples across tasks. Differently, (1) our method clusters tasks based on feature affinity scores, rather than example difficulty, which is shown to achieve better results for a variety of datasets; (2) the tree structure is determined offline using the precalculated affinities for the whole network, and not online in a greedy layer-by-layer fashion, which promotes task groupings that are optimal in a global, rather than local, sense.

Neural architecture search. Neural architecture search (NAS) [11] aims to automate the construction of the network architecture. Different algorithms can be characterized based on their search space, search strategy or performance estimation strategy. Most existing works on NAS, however, are limited to task-specific models [24, 25, 37, 38, 49]. This is to be expected as when using NAS for MTL, layer sharing has to be jointly optimized with the layers types, their connectivity, etc., rendering the problem considerably expensive. To alleviate the heavy computation burden, a recent work [23] implemented an evolutionary architecture search for multi-task networks, while other researchers explored more viable alternatives, like routing [39], stochastic filter grouping [5], and feature partitioning [35]. In contrast to traditional NAS, the proposed methods do not build the architecture from scratch, but rather start from a predefined backbone network for which a layer sharing scheme is automatically determined.

Transfer learning. Transfer learning [36] uses the knowledge obtained when solving one task, and applies it to a different but related task. Our work is loosely related to transfer learning, as we use it to measure levels of task affinity. [46] provided a taxonomy for task

Algorithm 1 Branched Multi-Task Networks - Task clustering

-
- 1: **Input:** Tasks \mathcal{T} , K images \mathcal{I} , a sharable encoder E with D locations where we can branch, a set of task specific decoders D_t and a computational budget \mathcal{C} .
 - 2: **for** t in \mathcal{T} **do**
 - 3: Train the encoder E and task-specific decoder D_t for task t .
 - 4: $RDM^t \leftarrow RDM(E, D, \mathcal{I})$ ▷ RDM for task t .
 - 5: **end for**
 - 6: $A_{d,i,j} \leftarrow r_s \left(\text{triu} \left(RDM_{d,::}^{t_i} \right), \text{triu} \left(RDM_{d,::}^{t_j} \right) \right)$ **for** t_i, t_j in \mathcal{T} and d in locations ▷ Task affinity
 - 7: $D = 1 - A$ ▷ Task dissimilarity
 - 8: **Return:** Task-grouping with minimal task dissimilarity that fits within \mathcal{C}
-

transfer learning to quantify such relationships. However, their approach scales unfavorably w.r.t. the number of tasks, and we opted for a more efficient alternative proposed by [10]. The latter uses RSA to obtain a measure of task affinity, by computing correlations between models pretrained on different tasks. In our method, we use the performance metric from their work to compare the usefulness of different feature sets for solving a particular task.

Loss weighting. One of the known challenges of jointly learning multiple tasks is properly weighting the loss functions associated with the individual tasks. Early work [20] used the homoscedastic uncertainty of each task to weigh the losses. Gradient normalization [7] balances the learning of tasks by dynamically adapting the gradient magnitudes in the network. Liu *et al.* [26] weigh the losses to match the pace at which different tasks are learned. Dynamic task prioritization [14] prioritizes the learning of difficult tasks. [43] cast MTL as a multi-objective optimization problem, with the overall objective of finding a Pareto optimal solution. Note that, addressing the loss weighting issue in MTL is out of the scope of this work. In fact, all our experiments are based on a simple uniform loss weighing scheme.

3 Method

In this paper, we aim to jointly solve N different visual recognition tasks $\mathcal{T} = \{t_1, \dots, t_N\}$ given a computational budget \mathcal{C} , i.e. number of parameters or FLOPS. Consider a backbone architecture: an encoder, consisting of a sequence of shared layers or blocks f_l , followed by a decoder with a few task-specific layers. We assume an appropriate structure for layer sharing to take the shape of a tree. In particular, the first layers are shared by all tasks, while later layers gradually split off as they show more task-specific behavior. The proposed method aims to find an effective task grouping for the sharable layers f_l of the encoder, i.e. grouping related tasks together in the same branches of the tree. When two tasks are strongly related, we expect their single-task models to rely on a similar feature set [46]. Based on this viewpoint, the proposed method derives a task affinity score at various locations in the sharable encoder. The number of locations D can be freely determined as the number of candidate branching locations. As such, the resulting task affinity scores are used for the automated construction of a branched multi-task network that fits the computational budget \mathcal{C} . Fig. 3 illustrates our pipeline, while Algorithm 1 summarizes the whole procedure.

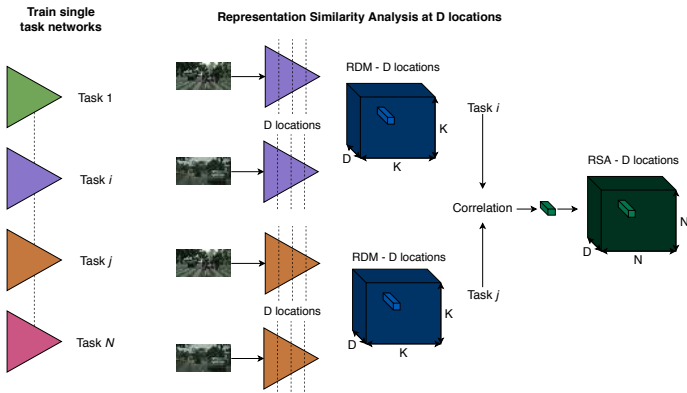


Figure 1: Pipeline overview. (left) We train a single-task model for every task $t \in \mathcal{T}$. (middle) We use RSA to measure the task affinity at D predefined locations in the sharable encoder. In particular, we calculate the representation dissimilarity matrices (RDM) for the features at D locations using K images, which gives a $D \times K \times K$ tensor per task. (right) The affinity tensor \mathbf{A} is found by calculating the correlation between the RDM matrices, which results in a three-dimensional tensor of size $D \times N \times N$, with N the number of tasks.




Figure 2: Our pipeline’s output is a branched multi-task network, similar to how NAS techniques output sample architectures. An example branched multi-task network is visualized here.

Figure 3: The proposed method: (a) calculate task affinities at various locations in the sharable encoder; (b) build a branched multi-task network based on the computed affinities.

3.1 Calculate task affinity scores

As mentioned, we rely on RSA to measure task affinity scores. This technique has been widely adopted in the field of neuroscience to draw comparisons between behavioral models and brain activity. Inspired by how [10] applied RSA to select tasks for transfer learning, we use the technique to assess the task affinity at predefined locations in the sharable encoder. Consequently, using the measured levels of task affinity, tasks are assigned in the same or different branches of a branched multi-task network, subject to the computational budget \mathcal{C} .

The procedure to calculate the task affinity scores is the following. As a first step, we train a single-task model for each task $t_i \in \mathcal{T}$. The single-task models use an identical encoder E - made of all sharable layers f_l - followed by a task-specific decoder D_{t_i} . The decoder contains only task-specific operations and is assumed to be significantly smaller in size compared to the encoder. As an example, consider jointly solving a classification and a dense prediction task. Some fully connected layers followed by a softmax operation are typically needed for the classification task, while an additional decoding step with some upscaling operations is required for the dense prediction task. Of course, the appropriate loss functions are applied in each case. Such operations are part of the task-specific decoder D_{t_i} . The different single-task networks are trained under the same conditions.

At the second step, we choose D locations in the sharable encoder E where we calculate a two-dimensional task affinity matrix of size $N \times N$. When concatenated, this results in a

three-dimensional tensor \mathbf{A} of size $D \times N \times N$ that holds the task affinities at the selected locations. To calculate these task affinities, we have to compare the representation dissimilarity matrices (RDM) of the single-task networks - trained in the previous step - at the specified D locations. To do this, a held-out subset of K images is required. The latter images serve to compare the dissimilarity of their feature representations in the single-task networks for every pair of images. Specifically, for every task t_i , we characterize these learned feature representations at the selected locations by filling a tensor of size $D \times K \times K$. This tensor contains the dissimilarity scores $1 - \rho$ between feature representations, with ρ the Pearson correlation coefficient. Specifically, $\mathbf{RDM}_{d,i,j}$ is found by calculating the dissimilarity score between the features at location d for image i and j . The 3-D tensors are linearized to 1-D tensors to calculate the Pearson correlation coefficient.

For a specific location d in the network, the computed RDMs are symmetrical, with a diagonal of zeros. For every such location, we measure the similarity between the upper or lower triangular part of the RDMs belonging to the different single-task networks. We use the Spearman’s correlation coefficient r_s to measure similarity. When repeated for every pair of tasks, at a specific location d , the result is a symmetrical matrix of size $N \times N$, with a diagonal of ones. Concatenating over the D locations in the sharable encoder, we end up with the desired task affinity tensor of size $D \times N \times N$. Note that, in contrast to prior work [31], the described method focuses on the features used to solve the single tasks, rather than the examples and how easy or hard they are across tasks, which is shown to result in better task groupings in Section 4. Furthermore, the computational overhead to determine the task affinity scores based on feature correlations is negligible. We conclude that the computational cost of the method boils down to pre-training N single task networks. A detailed computational cost analysis can be found in the supplementary materials.

Other measures of task similarity [2, 9] probed the features from a network pre-trained on ImageNet. This avoids the need to pre-train a set of single-task networks first. However, in this case, the task dictionaries only consisted of various, related classification problems. Differently, we consider more diverse, and loosely related task (see Section 4). In our case, it is arguably more important to learn the task-specific information needed to solve a task. This motivates the use of pre-trained single-task networks.

3.2 Construct a branched multi-task network

Given a computational budget \mathcal{C} , we need to derive how the layers (or blocks) f_l in the sharable encoder E should be shared among the tasks in \mathcal{T} . Each layer $f_l \in E$ is represented as a node in the tree, i.e. the root node contains the first layer f_0 , and nodes at depth l contain layer(s) f_l . The granularity of the layers f_l corresponds to the intervals at which we measure the task affinity in the sharable encoder, i.e. the D locations. When the encoder is split into b_l branches at depth l , this is equivalent to a node at depth l having b_l children. The leaves of the tree contain the task-specific decoders D_t . Fig. 2 shows an example of such a tree using the aforementioned notation. Each node is responsible for solving a unique subset of tasks.

The branched multi-task network is built with the intention to separate dissimilar tasks by assigning them to separate branches. To this end, we define the dissimilarity score between two tasks t_i and t_j at location d as $1 - \mathbf{A}_{d,i,j}$, with \mathbf{A} the task affinity tensor². The branched multi-task network is found by minimizing the sum of the task dissimilarity scores at every location in the sharable encoder. In contrast to prior work [31], the task affinity (and

²This is not to be confused with the dissimilarity score used to calculate the RDM elements $\mathbf{RDM}_{d,i,j}$.

dissimilarity) scores are calculated a priori. This allows us to determine the task clustering offline. Since the number of tasks is finite, we can enumerate all possible trees that fall within the given computational budget \mathcal{C} . Finally, we select the tree that minimizes the task dissimilarity score. The task dissimilarity score of a tree is defined as $C_{cluster} = \sum_l C_{cluster}^l$, where $C_{cluster}^l$ is found by averaging the maximum distance between the dissimilarity scores of the elements in every cluster. The use of the maximum distance encourages the separation of dissimilar tasks. By taking into account the clustering cost at all depths, the procedure can find a task grouping that is considered optimal in a global sense. This is in contrast to the greedy approach in [31], which only minimizes the task dissimilarity locally, i.e. at isolated locations in the network.

4 Experiments

In this section, we evaluate the proposed method on a number of diverse multi-tasking datasets, that range from real to semi-real data, from few to many tasks, from dense prediction to classification tasks, and so on. For every experiment, we describe the most important elements of the setup. We report the number of parameters (#P) for every model to facilitate a fair comparison. Additional implementation details can be found in the supplementary materials.

4.1 Cityscapes

Dataset. The Cityscapes dataset [8] considers the scenario of urban scene understanding. The train, validation and test set contain respectively 2975, 500 and 1525 real images, taken by driving a car in Central European cities. It considers a few dense prediction tasks: semantic segmentation (S), instance segmentation (I) and monocular depth estimation (D). As in prior works [20, 43], we use a ResNet-50 encoder with dilated convolutions, followed by a Pyramid Spatial Pooling (PSP) [17] decoder. Every input image is rescaled to 512 x 256 pixels. We reuse the approach from [20] for the instance segmentation task, i.e. we consider the proxy task of regressing each pixel to the center of the instance it belongs to.

Results. We measure the task affinity after every block (1 to 4) in the ResNet-50 model (see Fig. 4). The task affinity decreases in the deeper layers, due to the features becoming more task-specific. We compare the performance of the task groupings generated by our method with those by other approaches. As in [32], the performance of a multi-task model m is defined as the average per-task performance drop/increase w.r.t. a single-task baseline b .

We trained all possible task groupings that can be derived from branching the model in the last three ResNet blocks. Fig. 5 visualizes performance vs number of parameters for the trained architectures. Depending on the available computational budget \mathcal{C} , our method generates a specific task grouping. We visualize these generated groupings as a blue path in Fig. 5, when gradually increasing the computational budget \mathcal{C} . Similarly, we consider the task groupings when branching the model based on the task affinity measure proposed by FAFS [31] (green path). We find that, in comparison, the task groupings devised by our method achieve higher performance within a given computational budget \mathcal{C} . Furthermore, in the majority of cases, for a fixed budget \mathcal{C} the proposed method is capable of selecting the best performing task grouping w.r.t. performance vs parameters metric (blue vs other).

We also compare our branched multi-task networks with cross-stitch networks [33], NDDR-CNNs [13] and MTAN [26] in Table 2. While cross-stitch nets and NDDR-CNNs




Figure 4: Task affinity scores measured after each ResNet-50 block on Cityscapes.




Figure 5: Number of parameters versus multi-task performance on Cityscapes for different task groupings. The ‘Other Groupings’ contain any remaining tree structures that can be found by randomly branching the model in the last three ResNet blocks.

give higher multi-task performance, attributed to their computationally expensive soft parameter sharing setting, our branched networks can strike a better trade-off between the performance and number of parameters. In particular, we can effectively sample architectures which lie between the extremes of a baseline multi-task model and a cross-stitch or NDDR-CNN architecture. Finally, our models provide a more computationally efficient alternative to the MTAN model, which reports similar performance while using more parameters.

4.2 Taskonomy

Dataset. The Taskonomy dataset [46] contains semi-real images of indoor scenes, annotated for 26 (dense prediction, classification, etc.) tasks. Out of the available tasks, we select scene categorization (C), semantic segmentation (S), edge detection (E), monocular depth estimation (D) and keypoint detection (K). The task dictionary was selected to be as diverse as possible, while still keeping the total number of tasks reasonable for all computations. We use the tiny split of the dataset, containing 275k train, 52k validation and 54k test images. We reuse the architecture and training setup from [46]: the encoder is based on ResNet-50; a 15-layer fully-convolutional decoder is used for the pixel-to-pixel prediction tasks.

Results. The task affinity is again measured after every ResNet block. Since the number of tasks increased to five, it is very expensive to train all task groupings exhaustively, as done above. Instead, we limit ourselves to three architectures that are generated when gradually increasing the parameter budget. As before, we compare our task groupings against the method from [31]. The numerical results can be found in Table 1. The task groupings themselves are shown in the supplementary materials.

The effect of the employed task grouping technique can be seen from comparing the performance of our models against the corresponding FAFS models, generated by [31]. The latter are consistently outperformed by our models. Compared to the results on Cityscapes

Method	D (L1)↓	S (IoU)↑	C (Top-5)↑	E (L1)↓	K (L1)↓	#P (M)	Δ_{MTL} (%)↑
Single task	0.60	43.5	66.0	0.99	0.23	224	+ 0.00
MTL baseline	0.75	47.8	56.0	1.37	0.34	130	- 22.50
MTAN	0.71	43.8	59.6	1.86	0.40	158	-37.36
Cross-stitch	0.61	44.0	58.2	1.35	0.50	224	- 32.29
NDDR-CNN	0.66	45.9	64.5	1.05	0.45	258	- 21.02
FAFS - 1	0.74	46.1	62.7	1.30	0.39	174	- 24.5
FAFS - 2	0.80	39.9	62.4	1.68	0.52	188	- 48.32
FAFS - 3	0.74	46.1	64.9	1.05	0.27	196	- 8.48
Ours - 1	0.76	47.6	63.3	1.12	0.29	174	- 11.88
Ours - 2	0.74	48.0	63.6	0.96	0.35	188	- 12.66
Ours - 3	0.74	47.9	64.5	0.94	0.26	196	- 4.93

Table 1: Results on the tiny Taskonomy test set. The results for edge (E) and keypoints (K) detection were multiplied by a factor of 100 for better readability. The FAFS models refer to generating the task groupings with the task affinity technique proposed by [31].

(Fig. 5), we find that the multi-task performance is much more susceptible to the employed task groupings, possibly due to negative transfer. Furthermore, we observe that none of the soft parameter sharing models can handle the larger, more diverse task dictionary: the performance decreases when using these models, while the number of parameters increases. This is in contrast to our branched multi-task networks, which seem to handle the diverse set of tasks rather positively. As opposed to [46], but in accordance with [32], we show that it is possible to solve many heterogeneous tasks simultaneously when the negative transfer is limited, by separating dissimilar tasks from each other in our case. In fact, our approach is the first to show such consistent performance across different multi-tasking scenarios and datasets. Existing approaches seem to be tailored for particular cases, e.g. few/correlated tasks, synthetic-like data, binary classification only tasks, etc., whereas we show stable performance across the board of different experimental setups.

4.3 CelebA

Dataset. The CelebA dataset [28] contains over 200k images of celebrities, labeled with 40 facial attribute categories. The training, validation and test set contain 160k, 20k and 20k images respectively. We treat the prediction of each facial attribute as a single binary classification task, as in [18, 31, 43]. To ensure a fair comparison: we reuse the thin- ω model from [31] in our experiments; the parameter budget \mathcal{C} is set for the model to have the same amount of parameters as prior work.

Results. Table 3 shows the results on the CelebA test set. Our branched multi-task networks outperform earlier works [18, 31] when using a similar amount of parameters. Since the Ours-32 model (i.e. ω is 32) only differs from the FAFS model on the employed task grouping technique, we can conclude that the proposed method devises more effective task groupings for the attribute classification tasks on CelebA. Furthermore, the Ours-32 model performs on par with the VGG-16 model, while using 64 times less parameters. We also compare our results with the ResNet-18 model from [43]. The Ours-64 model performs 1.35% better compared to the ResNet-18 model when trained with a uniform loss weighing scheme. More noticeably, the Ours-64 model performs on par with the state-of-the-art ResNet-18 model that was trained with the MGDA loss weighing scheme from [43], while at the same time using 31% less parameters (11.2 vs 7.7 M).

Method	S (IoU) \uparrow	I (px) \downarrow	D (px) \downarrow	#P (M)	Δ_{MTL} (%) \uparrow
Single task	65.2	11.7	2.57	138	+0.00
MTL baseline	61.5	11.8	2.66	92	-3.33
MTAN	62.8	11.8	2.66	113	-2.53
Cross-stitch	65.1	11.6	2.55	140	+0.42
NDDR-CNN	65.6	11.6	2.54	190	+0.89
Ours - 1	62.1	11.7	2.66	107	-2.68
Ours - 2	62.7	11.7	2.62	114	-1.84
Ours - 3	64.1	11.6	2.62	116	-0.96

Table 2: Results on the Cityscapes validation set.

Method	Acc. (%)	#P (M)
MOON [40]	90.94	119.73
Independent Group [16]	91.06	-
MCNN [16]	91.26	-
MCNN-AUX [16]	91.29	-
VGG-16 [31]	91.44	134.41
FAFS [31]	90.79	2.09
<i>GNAS [16]</i>	<i>91.63</i>	<i>7.73</i>
Res-18 (Uniform) [43]	90.38	11.2
Res-18 (MGDA) [43]	91.75	11.2
Ours-32	91.46	2.20
<i>Ours-64</i>	<i>91.73</i>	<i>7.73</i>

Table 3: Results on the CelebA test set. The **Ours-32**, *Ours-64* architectures are found by optimizing the task clustering for the parameter budget that is used in the **FAFS**, *GNAS* model respectively.

5 Conclusion

In this paper, we introduced a principled approach to automatically construct branched multi-task networks for a given computational budget. To this end, we leverage the employed tasks’ affinities as a quantifiable measure for layer sharing. The proposed approach can be seen as an abstraction of NAS for MTL, where only layer sharing is optimized, without having to jointly optimize the layers types, their connectivity, etc., as done in traditional NAS, which would render the problem considerably expensive. Extensive experimental analysis shows that our method outperforms existing ones w.r.t. the important metric of multi-tasking performance vs number of parameters, while at the same time showing consistent results across a diverse set of multi-tasking scenarios and datasets.

Acknowledgment This work is sponsored by the Flemish Government under the Artificiele Intelligentie (AI) Vlaanderen programme. The authors also acknowledge support by Toyota via the TRACE project and MACCHINA (KU Leuven, C14/18/065).

References

- [1] Jacob Abernethy, Francis Bach, Theodoros Evgeniou, and Jean-Philippe Vert. A new approach to collaborative filtering: Operator estimation with spectral regularization. *JMLR*, 10(Mar):803–826, 2009.
- [2] Alessandro Achille, Michael Lam, Rahul Tewari, Avinash Ravichandran, Subhransu Maji, Charless C Fowlkes, Stefano Soatto, and Pietro Perona. Task2vec: Task embedding for meta-learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6430–6439, 2019.
- [3] Arvind Agarwal, Samuel Gerber, and Hal Daume. Learning multiple tasks using manifold regularization. In *NIPS*, 2010.
- [4] Andreas Argyriou, Theodoros Evgeniou, and Massimiliano Pontil. Multi-task feature learning. In *NIPS*, 2007.
- [5] Felix JS Bragman, Ryutaro Tanno, Sebastien Ourselin, Daniel C Alexander, and M Jorge Cardoso. Stochastic filter groups for multi-task cnns: Learning specialist and generalist convolution kernels. In *ICCV*, 2019.
- [6] Rich Caruana. Multitask learning. *Machine learning*, 28(1):41–75, 1997.
- [7] Zhao Chen, Vijay Badrinarayanan, Chen-Yu Lee, and Andrew Rabinovich. Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In *ICML*, 2018.
- [8] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, 2016.
- [9] Yin Cui, Yang Song, Chen Sun, Andrew Howard, and Serge Belongie. Large scale fine-grained categorization and domain-specific transfer learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4109–4118, 2018.
- [10] Kshitij Dwivedi and Gemma Roig. Representation similarity analysis for efficient task taxonomy & transfer learning. In *CVPR*, pages 12387–12396, 2019.
- [11] Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. Neural architecture search: A survey. *JMLR*, 20(55):1–21, 2019.
- [12] Theodoros Evgeniou and Massimiliano Pontil. Regularized multi-task learning. In *KDD*, 2004.
- [13] Yuan Gao, Jiayi Ma, Mingbo Zhao, Wei Liu, and Alan L Yuille. Nddr-cnn: Layerwise feature fusing in multi-task cnns by neural discriminative dimensionality reduction. In *CVPR*, pages 3205–3214, 2019.
- [14] Michelle Guo, Albert Haque, De-An Huang, Serena Yeung, and Li Fei-Fei. Dynamic task prioritization for multitask learning. In *ECCV*, 2018.

- [15] Moshe Gur and D Max Snodderly. Direction selectivity in v1 of alert monkeys: evidence for parallel pathways for motion processing. *The Journal of physiology*, 585(2): 383–400, 2007.
- [16] Emily M Hand and Rama Chellappa. Attributes for improved attributes: A multi-task network utilizing implicit and explicit relationships for facial attribute classification. In *AAAI*, 2017.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. *IEEE transactions on pattern analysis and machine intelligence*, 37(9):1904–1916, 2015.
- [18] Siyu Huang, Xi Li, Zhiqi Cheng, Alexander Hauptmann, et al. Gnas: A greedy neural architecture search method for multi-attribute learning. 2018.
- [19] Ali Jalali, Sujay Sanghavi, Chao Ruan, and Pradeep K Ravikumar. A dirty model for multi-task learning. In *NIPS*, 2010.
- [20] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *CVPR*, 2018.
- [21] Iasonas Kokkinos. Ubertnet: Training a universal convolutional neural network for low-, mid-, and high-level vision using diverse datasets and limited memory. In *CVPR*, 2017.
- [22] Abhishek Kumar and Hal Daume III. Learning task grouping and overlap in multi-task learning. 2012.
- [23] Jason Liang, Elliot Meyerson, and Risto Miikkulainen. Evolutionary architecture search for deep multitask networks. In *GECCO*, 2018.
- [24] Chenxi Liu, Barret Zoph, Maxim Neumann, Jonathon Shlens, Wei Hua, Li-Jia Li, Li Fei-Fei, Alan Yuille, Jonathan Huang, and Kevin Murphy. Progressive neural architecture search. In *ECCV*, 2018.
- [25] Hanxiao Liu, Karen Simonyan, Oriol Vinyals, Chrisantha Fernando, and Koray Kavukcuoglu. Hierarchical representations for efficient architecture search. In *ICLR*, 2018.
- [26] Shikun Liu, Edward Johns, and Andrew J Davison. End-to-end multi-task learning with attention. In *CVPR*, 2019.
- [27] Sulin Liu, Sinno Jialin Pan, and Qirong Ho. Distributed multi-task relationship learning. In *KDD*, 2017.
- [28] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *ICCV*, 2015.
- [29] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and S Yu Philip. Learning multiple tasks with multilinear relationship networks. In *NIPS*, pages 1594–1603, 2017.
- [30] Karim Lounici, Massimiliano Pontil, Alexandre B Tsybakov, and Sara Van De Geer. Taking advantage of sparsity in multi-task learning. In *COLT*, 2009.

- [31] Yongxi Lu, Abhishek Kumar, Shuangfei Zhai, Yu Cheng, Tara Javidi, and Rogerio Feris. Fully-adaptive feature sharing in multi-task networks with applications in person attribute classification. In *CVPR*, 2017.
- [32] Kevis-Kokitsi Maninis, Ilija Radosavovic, and Iasonas Kokkinos. Attentive single-tasking of multiple tasks. In *CVPR*, pages 1851–1860, 2019.
- [33] Ishan Misra, Abhinav Shrivastava, Abhinav Gupta, and Martial Hebert. Cross-stitch networks for multi-task learning. In *CVPR*, 2016.
- [34] Davy Neven, Bert De Brabandere, Stamatios Georgoulis, Marc Proesmans, and Luc Van Gool. Fast scene understanding for autonomous driving. In *IV Workshops*, 2017.
- [35] Alejandro Newell, Lu Jiang, Chong Wang, Li-Jia Li, and Jia Deng. Feature partitioning for efficient multi-task architectures. *arXiv preprint arXiv:1908.04339*, 2019.
- [36] Sinno Jialin Pan, Qiang Yang, et al. A survey on transfer learning. *TKDE*, 22(10): 1345–1359, 2010.
- [37] Hieu Pham, Melody Y Guan, Barret Zoph, Quoc V Le, and Jeff Dean. Efficient neural architecture search via parameter sharing. In *ICML*, 2018.
- [38] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. Regularized evolution for image classifier architecture search. In *AAAI*, 2019.
- [39] Clemens Rosenbaum, Tim Klinger, and Matthew Riemer. Routing networks: Adaptive selection of non-linear functions for multi-task learning. In *ICLR*, 2018.
- [40] Ethan M Rudd, Manuel Günther, and Terrance E Boulton. Moon: A mixed objective optimization network for the recognition of facial attributes. In *ECCV*, 2016.
- [41] Sebastian Ruder. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*, 2017.
- [42] Sebastian Ruder, Joachim Bingel, Isabelle Augenstein, and Anders Søgaard. s. In *AAAI*, 2019.
- [43] Ozan Sener and Vladlen Koltun. Multi-task learning as multi-objective optimization. In *NIPS*, 2018.
- [44] Dan Xu, Wanli Ouyang, Xiaogang Wang, and Nicu Sebe. Pad-net: Multi-tasks guided prediction-and-distillation network for simultaneous depth estimation and scene parsing. In *CVPR*, pages 675–684, 2018.
- [45] Ming Yuan and Yi Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2006.
- [46] Amir R Zamir, Alexander Sax, William Shen, Leonidas J Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling task transfer learning. In *CVPR*, 2018.
- [47] Xiangyun Zhao, Haoxiang Li, Xiaohui Shen, Xiaodan Liang, and Ying Wu. A modulation module for multi-task learning with applications in image retrieval. In *ECCV*, 2018.

- [48] Jiayu Zhou, Jianhui Chen, and Jieping Ye. Malsar: Multi-task learning via structural regularization. *Arizona State University*, 21, 2011.
- [49] Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. In *ICLR*, 2017.