Efficiently Identifying Task Groupings for Multi-Task Learning

Anonymous Author(s)
Affiliation
Address
email

Abstract

Multi-task learning can leverage information learned by one task to benefit the training of other tasks. Despite this capacity, naïvely training all tasks together in one model often degrades performance, and exhaustively searching through combinations of task groupings can be prohibitively expensive. As a result, efficiently identifying the tasks that would benefit from co-training remains a challenging design question without a clear solution. In this paper, we suggest an approach to select which tasks should train together in multi-task learning models. Our method determines task groupings in a single training run by co-training all tasks together and quantifying the effect to which one task's gradient would affect another task's loss. On the large-scale Taskonomy computer vision dataset, we find this method can decrease test loss by 10.0% compared to simply training all tasks together while operating 11.6 times faster than a state-of-the-art task grouping method.

1 Introduction

Many of the forefront challenges in applied machine learning demand that a single model performs well on multiple tasks, or optimizes multiple objectives while simultaneously adhering to unmovable inference-time constraints. For instance, autonomous vehicles necessitate low inference time latency to make multiple predictions on a real-time video feed to precipitate a driving action [24]. Robotic arms are asked to concurrently learn how to pick, place, cover, align, and rearrange various objects to improve learning efficiency [22], and online movie recommendation systems model multiple engagement metrics to facilitate low-latency personalized recommendations [11]. Each of the above applications depends on multi-task learning, and advances which improve multi-task learning performance have the potential to make an out-sized impact on these and many other domains.

Multi-task learning can improve modeling performance by introducing an inductive bias to prefer hypothesis classes which explain multiple objectives and by focusing attention on relevant features [40]. However, it may also lead to severely degraded performance when tasks compete for model capacity or are unable to build a shared representation that can generalize to all objectives. Accordingly, finding groups of tasks that derive benefit from the positives of co-training while mitigating the negatives often improves the modeling performance of multi-task learning systems.

While recent work has developed new multi-task learning optimization schemes [25, 9, 42, 50, 10, 47], the problem of deciding which tasks should be trained together in the first place is an understudied and complex issue that is often left to human experts [53]. However, a human's understanding of similarity is motivated by their intuition and experience rather than a prescient knowledge of the underlying structures learned by a neural network. To further complicate matters, the benefit or detriment induced from co-training relies on many non-trivial decisions including, but not limited to, dataset characteristics, model architecture, hyperparameters, capacity, and convergence [48, 46, 43, 44]. As

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

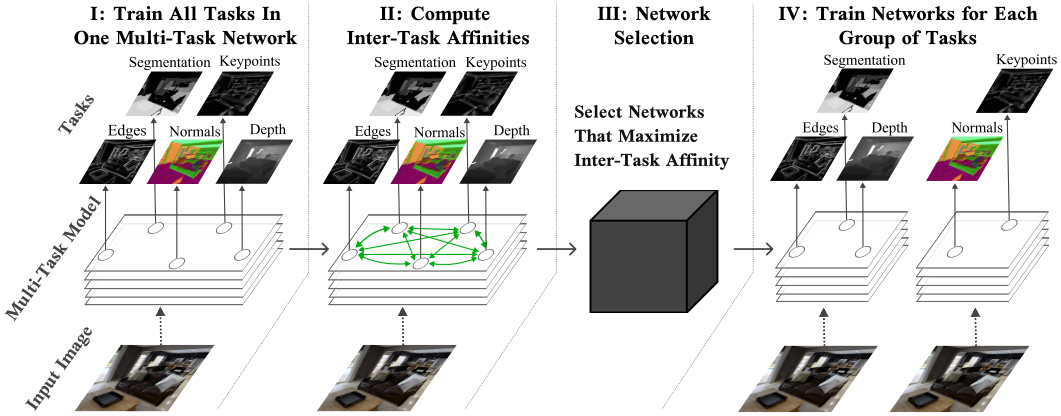

Figure 1: Overview of our suggested approach to efficiently determine task groupings. (I): Train all tasks together in a multi-task learning model. (II): Compute inter-task affinity scores during training. (III): Select multi-task networks that maximize the inter-task affinity score onto each serving-time task. (IV): Train the resulting networks and deploy to inference.

a result, a systematic technique to determine which tasks should train together in a multi-task neural network would be valuable to practitioners and researchers alike [4, 5].

One approach to select task groupings is to exhaustively search over the $2^{|\mathcal{T}|} - 1$ multi-task networks for a set of tasks $\mathcal{T}$. However, the cost associated with this search can be prohibitive, especially when there is a large number of tasks or for models under active development. Moreover, as model scale and complexity continues to increase, even approximate task grouping algorithms which evaluate only a subset of combinations may become prohibitively costly and time-consuming to evaluate.

In this paper, we aim to develop an efficient framework to select task groupings without sacrificing performance. We propose to measure inter-task affinity by training all tasks together in a single multi-task network and quantifying the effect to which one task's gradient update would affect another task's loss. This per-step quantity is averaged across training, and tasks are then grouped together to maximize the affinity onto each task. A visual depiction of the method is shown in Figure 1. Our suggested approach makes no assumptions regarding model architecture, and is applicable to any paradigm in which shared parameters are updated with respect to multiple losses.

In summary, our primary contribution is to suggest a measure of inter-task affinity that can be used to systematically and efficiently determine task groupings for multi-task learning. Our theoretical analysis shows that grouping tasks by maximizing inter-task affinity will outperform any other task grouping in the convex setting under mild conditions. Further on two challenging multi-task image benchmarks, our empirical analysis finds this approach outperforms training all tasks independently, training all tasks together (with and without training augmentations), and is competitive with a state-of-the-art task grouping method while decreasing runtime by more than an order of magnitude.

2 Related Work

Task Groupings. Prevailing wisdom suggests tasks which are similar or share a similar underlying structure may benefit from co-training in a multi-task system [8, 7, 3]. Early work in this domain pertaining to the convex setting assume all tasks share a common latent feature representation, and find that model performance can be significantly improved by clustering tasks based on the basis vectors they share in this latent space [23, 27]. However, early convex methods to determine task groupings often make prohibitive assumptions that do not scale to deep neural networks.

Deciding which tasks should train together in multi-task neural networks has traditionally been addressed with costly cross-validation techniques or high variance human intuition. An altogether different approach may leverage recent advances in transfer learning focused on understanding task relationships [51, 2, 13, 55, 1]; however, [43] show transfer learning algorithms which determine task similarity do not carry over to the multi-task learning domain and instead propose a multi-task specific framework which trains between $\binom{|\mathcal{T}|}{2} + |\mathcal{T}|$ and $2^{|\mathcal{T}|} - 1$ models to approximate exhaustive search performance. Our approach differs from [43] in that it computes task groupings from only a single training run.

Architectures and Training Dynamics. A plethora of multi-task methods addressing what parameters to share among tasks in a model have been developed, such as Neural Architecture Search [17, 44, 46, 32, 19, 36], Soft-Parameter Sharing [37, 12, 49], and asymmetric information transfer [28, 41, 29] to improve multi-task performance. Although this direction is promising, we direct our focus towards *when to share* tasks in a multi-task network rather than architecture modifications to maximize the benefits of co-training. Nevertheless, both approaches are complementary, and architecture augmentations seem to perform best when trained with related tasks [40].

Significant effort has also been invested to improve the optimization dynamics of MTL systems. In particular, dynamic loss reweighing has achieved performance superior to using fixed loss weights found with extensive hyperparameter search [25, 16, 34, 9, 42, 30]. Another set of methods seek to mitigate inter-task conflict by manipulating the direction of task gradients rather than simply their magnitude [45, 54, 50, 10, 47, 33]. In our experiments, we compare against Uncertainty Weights [25], GradNorm [9], and PCGrad [50] to contextualize the relative change in performance from splitting tasks into groups. We find that task grouping methods outperform all three training augmentations; nonetheless, they are naturally complementary. In Section 5, our results indicate enhancing the networks found by our method with PCGrad can lead to additional improvements in performance.

Looking into the Future. "Lookahead" methods in deep learning can often be characterized by saving the current state of the model, applying one or more gradient updates to a subset of the parameters, reloading the saved state, and then leveraging the information learned from the future state to modify the current set of parameters. This approach has been used extensively in the meta-learning [14, 39, 6, 15, 26], optimization [38, 18, 52, 20, 21], and recently auxiliary task learning domains [31]. Unlike the above mentioned methods which look into the future to modify optimization processes, our work adapts this central concept to the multi-task learning domain to characterize task interactions and assign tasks to groups of networks.

3 Task Grouping Problem Definition

We draw a distinction between inference-time latency constraints and inference-time budget. The former characterizes the speed at which predictions can be computed, with similarly sized models running in parallel having latency roughly equivalent to a single model running by itself. The latter relates to the number of parameters used by all models during inference, with a n-times parameter model having similar budget to an n-group of normal sized models. We configure our analysis to span both dimensions, but also provide analysis into only the latter in the Appendix.

Given a set of tasks $\mathcal{T}$, a fixed inference-time budget b, and a fixed inference-time latency constraint c, our aim is to assign tasks to networks such that combined task performance is maximized. Additionally, each network must have parameter count less than c; the networks must span our set of tasks $\mathcal{T}$; and the total number of networks is less than or equal to our inference time budget b. Moreover, tasks can be trained in a model where they are not served to assist in the training of other tasks. Formulating our task grouping framework in this manner aligns our work with industry trends where inference-time budget is limited and inference-time latency unmovable.

More formally, for a set of n tasks $\mathcal{T} = \{\tau_1, \tau_2, .., \tau_n\}$, we would like to construct a group of k multi-task neural networks $M = \{m_1, m_2, ..., m_k\}$ such that $\forall \tau_i \in \mathcal{T}$, $\exists$ **exactly one** multi-task network $m_j \in M$, parameter count of $m_j < c$, such that m_j makes an inference-time prediction for t_i subject to $k \leq b$ where b is our fixed inference-time budget. m_j is the jth multi-task network which takes an input X, and concurrently trains a set of tasks $\{\tau_a, \tau_c, ..., \tau_f\}$, but only serves a subset of those tasks at inference. For a given performance measure $\mathcal{P}$, we can then define the aggregate performance of our task grouping as $\sum_{i=1}^{n} \mathcal{P}(\tau_i|M)$ where $\mathcal{P}(\tau_i|M)$ computes the performance of task τ_i from the set of models M using the model m_j which the task grouping algorithm predicts the performance of τ_i will be highest.

4 Grouping Tasks by Measuring Inter-Task Affinity

We propose a method to group tasks by examining the effect to which one task's gradient would increase or decrease another task's loss. We formally define the method in Section 4.1, describe a systematic procedure to go from inter-task affinity scores to a grouping of tasks in Section 4.2, and provide theoretical analysis in Section 4.3.

4.1 Inter-Task Affinity

Within the context of a hard-parameter sharing paradigm, tasks collaborate to build a shared feature representation which is then specialized by individual task-specific heads to output a prediction. Specifically, through the process of successive gradient updates to the shared parameters, tasks implicitly transfer information to each other. As a consequence, we propose to view the extent to which a task's successive gradient updates on the shared parameters affect the objective of other tasks in the network as a proxy measurement of inter-task affinity.

Consider a multitask loss function parameterized by $\{\theta_s\} \cup \{\theta_i | \, i \in \mathcal{T}\}$ where θ_s represents the shared parameters and θ_i represents the task $i \in \mathcal{T}$ specific parameters. Given a batch of examples $\mathcal{X}$, let

$$L_{\text{total}}(\mathcal{X}, \theta_s, \{\theta_i\}) = \sum_{i \in \mathcal{T}} L_i(\mathcal{X}, \theta_s, \theta_i) \,,$$

denote the total loss where L_i represents the non-negative loss of task i. For simplicity of notation, we set the loss weight of each task to be equal to 1, though our construction generalizes to arbitrary weightings.

For a given training batch $\mathcal{X}^t$ at time-step t, define the quantity $\theta^{t+1}_{s|i}$ to represent the updated shared parameters after a gradient step with respect to the task i. Assuming stochastic gradient descent for simplicity, we have

$$\theta^{t+1}_{s|i} \coloneqq \theta^t_s - \eta \nabla_{\theta^t_s} L_i(\mathcal{X}^t, \theta^t_s, \theta^t_i) \,.$$

We can now calculate a *lookahead* loss for each task by using the using the updated shared parameters while keeping the task-specific parameters as well as the input batch unchanged. That is, in order to assess the effect of the gradient update of task i on a given task j, we can compare the loss of task j before and after applying the gradient update from task i onto the shared parameters. To eliminate the scale discrepancy among different task losses, we consider the ratio of a task's loss before and after the gradient step on the shared parameters as a scale invariant measure of relative progress. We can then define an asymmetric measure for calculating the *affinity* of task i at a given time-step t on task j as

$$\mathcal{Z}^t_{i \to j} = 1 - \frac{L_j(\mathcal{X}^t, \theta^{t+1}_{s|i}, \theta^t_j)}{L_j(\mathcal{X}^t, \theta^t_s, \theta^t_j)} \,. \tag{1}$$

Notice that a positive value of $\mathcal{Z}^t_{i \to j}$ indicates that the update on the shared parameters results in a lower loss on task j than the original parameter values, while a negative value of $\mathcal{Z}^t_{i \to j}$ indicates that the shared parameter update is antagonistic for this task's performance. Our suggested measure of inter-task affinity is computed at a per-step level of granularity, but can be averaged across all steps, every n steps, or a contiguous subset of steps to derive a "training-level" score:

$$\hat{\mathcal{Z}}_{i \to j} = \frac{1}{T} \sum_{t=1}^{T} \mathcal{Z}^t_{i \to j}$$

In Section 5, we find $\hat{\mathcal{Z}}_{i \to j}$ is empirically effective in selecting high performance task groupings and provide an ablation study along this dimension in Section 5.2.

4.2 Network Selection Algorithm

At a high level, our proposed approach trains all tasks together in one model, measures the pairwise task affinities throughout training, identifies task groupings that maximize total inter-task affinity, and then trains the resulting groupings for evaluation on the test set. We denote this framework as Task Affinity Grouping (TAG) and now describe the algorithm to convert inter-task affinity scores into a set of multi-task networks. More formally, given $\binom{|\mathcal{T}|}{2}$ values representing the pairwise inter-task affinity scores collected during a single training run, our network selection algorithm should produce k multi-task networks, $k \leq b$ where b is the inference-time budget, with the added constraint that every task must be served from exactly one network at inference.

For a group composed of a pair of tasks $\{a, b\}$, the affinity score onto task a would simply be $\hat{\mathcal{Z}}_{b \to a}$ and the affinity score onto task b would be $\hat{\mathcal{Z}}_{a \to b}$. For task groupings consisting of three or more tasks, we approximate the inter-task affinity onto a given task by averaging the pairwise affinities onto this

given task. Consider the group consisting of tasks $\{a, b, c\}$. We can compute the total inter-task affinity onto task a by averaging the pair-wise affinities from tasks b and c onto a: $(\hat{\mathcal{Z}}_{b\to a} + \hat{\mathcal{Z}}_{c\to a})/2$.

After approximating higher order affinity scores for each network consisting of three or more tasks, we select a set of k multi-task networks such that the total affinity score onto each serving-time task is maximized. Informally, each task which is being served at inference should train with the tasks which most decrease its loss throughout training. This problem is NP-hard (reduction from Set-Cover), but can be solved efficiently with a branch-and-bound-like algorithm as detailed in [43] or with a binary integer programming solver as done by [51].

4.3 Theoretical Analysis

We now offer theoretical analysis of our measure of inter-task affinity. Specifically, our goal is to provide an answer to the following question: given that task b induces higher inter-task affinity than task c on task a, does training $\{a, b\}$ together result in a lower loss on task a than training $\{a, c\}$? Intuitively, we expect the answer to this question to always be positive. However, it is easy to construct counter examples for simple quadratic loss functions where training $\{a, b\}$ actually induce a higher loss than training $\{a, c\}$ (see Appendix). Nonetheless, we show that in a convex setting and under some mild assumptions, the grouping suggested by our measure of inter-task affinity is guaranteed to induce a lower loss value on task a.

For simplicity of notation, we ignore the task specific parameters and use $L_a(\theta)$ to denote the loss of task a evaluated at shared parameters θ. Additionally, we denote the gradient of tasks a, b, and c at θ as g_a, g_b, and g_c, respectively.

Lemma 1. *Let L_a be a α-strongly convex and β-strongly smooth loss function. Given that task b induces higher transference than task c on task a, the following inequality holds:*

$$g_a \cdot g_c - \frac{\beta\,\eta}{2}\,\|g_c\|^2 + \frac{\alpha\,\eta}{2}\,\|g_b\|^2 \leq g_a \cdot g_b \tag{2}$$

The proof is given in the Appendix.

Proposition 1. *Let L_a be a α-strongly convex and β-strongly smooth loss function. Let $\eta \leq \frac{1}{\beta}$ be the learning rate. Suppose that in a given step, task b has higher inter-task affinity than task c on task a. Moreover, suppose that the gradients have equal norm, i.e. $\|g_a\| = \|g_b\| = \|g_c\|$. Then, taking a gradient step on the parameters using the combined gradient of a and b reduces L_a more so than taking a gradient step on the parameters using the combined gradient of a and c, given that $\cos(g_a, g_c) \leq \frac{\eta}{4}\frac{\alpha\,\beta}{\beta-\alpha} - 1$ where $\cos(u, v) \coloneqq \frac{u \cdot v}{\|u\|\,\|v\|}$ is the cosine similarity between u and v.*

Proof. Applying the strong smoothness upper-bound on the updated loss using the combined gradient $g_a + g_b$, we have

$$\begin{aligned} L_a(\theta - \eta\,(g_a + g_b)) &\leq L_a(\theta) - \eta\, g_a \cdot (g_a + g_b) + \frac{\eta^2\,\beta}{2}\,\|g_a + g_b\|^2 \\ &= L_a(\theta) + (\eta^2\,\beta - \eta)\, g_a \cdot g_b + (\frac{\eta^2\,\beta}{2} - \eta)\,\|g_a\|^2 + \frac{\eta^2\,\beta}{2}\,\|g_b\|^2 \end{aligned}$$

We would like to show that the last line is less than or equal to a lower-bound on the loss obtained using the combined gradient $g_a + g_c$

$$L_a(a) - \eta\, g_a\,(a_a + g_c) + \frac{\eta^2\,\alpha}{2}\,\|g_a + g_c\|^2 \leq L_a(\theta - \eta\,(g_a + g_c))$$

Eliminating the common terms, we would like to show the following inequality holds

$$(\eta\beta - 1)\, g_a \cdot g_b + \frac{\eta\,\beta}{2}(\|g_a\|^2 + \|g_b\|^2) \leq (\eta\alpha - 1)\, g_a \cdot g_c + \frac{\eta\alpha}{2}(\|g_a\|^2 + \|g_c\|^2)\,.$$

Using $\eta \leq \frac{1}{\beta}$, the first term becomes negative. Replacing for $g_a \cdot g_b$ using Eq. (2), it suffices to show the following inequality

$$\begin{aligned} (\eta\beta - 1)\,(g_a \cdot g_c + \frac{\eta\alpha}{2}(\|g_b\|^2) - \frac{\eta\beta}{2}(\|g_c\|^2)) + \frac{\eta\beta}{2}(\|g_a\|^2 + \|g_b\|^2) \\ \leq (\eta\alpha - 1)\, g_a \cdot g_c + \frac{\eta\alpha}{2}(\|g_a\|^2 + \|g_c\|^2)\,. \end{aligned}$$

Rearranging the terms, it suffices to show

$$0 \leq \|g_a\|^2(\frac{\eta}{2}(\alpha-\beta)) + \|g_b\|^2(-\frac{\eta}{2}\alpha(\eta\beta-1)) + \|g_c\|^2(\frac{\eta}{2}(\alpha-\beta)+\frac{\eta^2}{2}\beta^2) + \eta(\alpha-\beta)\, g_a \cdot g_c\,.$$

Under the mild assumption that $\|g_a\| = \|g_b\| = \|g_c\|$, we can rearrange the terms to obtain $\cos(g_a, g_c) \leq \frac{\eta}{4}\frac{\beta/\alpha}{\beta/\alpha-1} - 1$, where $\cos(u,v) := \frac{u \cdot v}{\|u\|\,\|v\|}$ is the cosine similarity between u and v. □

Proposition 1 intuitively implies the grouping chosen by maximizing per-task inter-task affinity is guaranteed to make more progress than any other group. Moreover, the assumptions for this condition to hold in the convex setting are fairly mild. Specifically, our first assumption relies on the gradients to have equal norms; but our proof can be generalized to when the gradient norms are approximately equal. This is often the case during training when tasks use similar loss functions and/or compute related quantities. The second assumption relies on the cosine similarity between the lower-affinity task c and the primary task a being smaller than a constant. This is a mild assumption since the ratio $\frac{\beta/\alpha}{\beta/\alpha-1}$ becomes sufficiently large for the assumption to hold trivially when the Hessian of the loss has a sufficiently small condition number.

5 Experiments

We evaluate the capacity of TAG to select task groupings on CelebA, a large-scale face attributes dataset [35] and Taskonomy, a massive computer vision dataset of indoor scenes [51]. Following this analysis, we direct our focus to answering the following questions with ablation experiments on CelebA:

- Does our measure of inter-task affinity align with identifying which tasks should train together?
- Should inter-task affinity be measured at every step of training to determine task groupings?
- Is measuring the change in train loss comparable with the change in validation loss when computing inter-task affinity?
- Do changes in a model's hyperparamters change which tasks should be trained together?

As described in Section 3, we constrain all networks to have the same number of parameters to adhere to a fixed inference-time latency constraint. We provide experimental results removing this constraint, as well as additional experimental results and detail relating to experimental design, in the Appendix.

5.1 Supervised Task Grouping Evaluation

For our task grouping evaluation, we compare two classes of approaches: approaches that determine task groupings, and approaches that train on all tasks together but alter the optimization. In the first class, we consider simply training all tasks together in the same network (MTL), training every task by itself (STL), the expected value from randomly selecting task groupings (RG), grouping tasks by maximizing inter-task cosine similarity between pairs of gradients (CS), our method (TAG), and HOA [43] which approximates higher-order task groupings from pair-wise task performance. For the latter class, we consider Uncertainty Weights (UW) [25], GradNorm (GN) [9], and PCGrad [50]. In principle, these two classes of approaches are complementary and can be combined.

Our empirical findings are summarized in Figure 2. For task grouping methods, we report the time to determine task groupings, not determine task groupings and train the resultant multi-task networks. This is to facilitate comparison between the efficiency of different task grouping methods and provide a high-level overview of how long it takes to select task groupings compared to popular multi-task learning benchmarks.

CelebA. We select a subset of 9 attributes {a1, a2, a3, a4, a5, a6, a7, a8, a9} from the 40 possible attributes in CelebA and optimize the baseline MTL model by tuning architecture, batch size, and learning rate to maximize the performance of training all tasks together on the validation set. We do not tune other methods with the exception of GradNorm, for which we search over {0.1, 0.5, 1.0, 1.5, 2.0, 3.0, 5.0} for alpha. For task grouping algorithms, we evaluate the set of {2-splits, 3-splits, 4-splits} inference-time budgets. Our findings are summarized in Figure 2 (left).

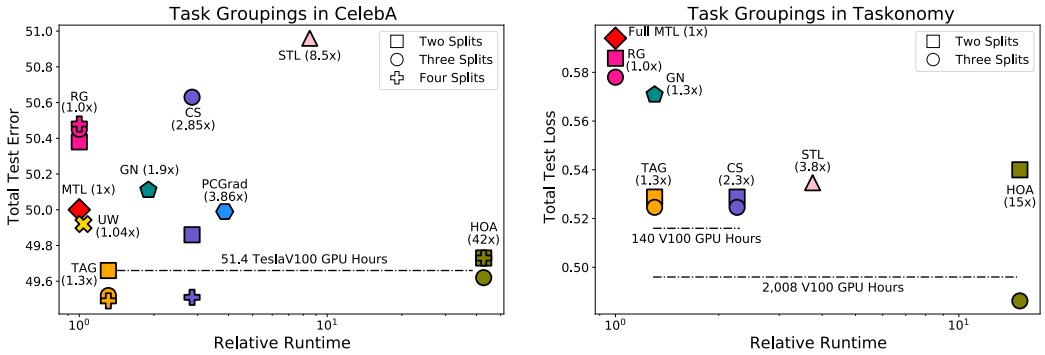

Figure 2: (Left) average classification error for 2, 3, and 4-split task groupings for the subset of 9 tasks in CelebA. (Right) total test loss for 2 and 3-split task groupings for the subset of 5 tasks in Taskonomy. All models were run on a TeslaV100 instance with the time to train the full MTL model being approximately 83 minutes in CelebA and 146 hours in Taskonomy. Note the x-axis is in log scale, and the relative runtime for each task grouping method only considers the time to find groups.

We find the performance of TAG surpasses that of the HOA, RG, and CS task grouping methods, while operating 22 times faster than HOA. We also find UW, GN, and PCGrad to perform worse than the groups found by either HOA or TAG, suggesting the improvement from identifying tasks which train well together cannot be replaced with current multi-task training augmentations. Nevertheless, we find multi-task training augmentations can be complementary with task grouping methods. For example, augmenting the groups found by TAG with PCGrad improves 2-splits performance by 0.85%, 3-splits performance by 0.18%, but changes 4-splits performance by -0.08%.

Similar to the results from [43], we find GradNorm [9] can sometimes perform worse than training all tasks together. We reason this difference is due to our common experimental design that loads the weights from the epoch with lowest validation loss to reduce overfitting, and differs from prior work which uses model weights after training for 100 epochs [42]. Reformulating our design to mirror [42], we find training the model to 100 epochs without early stopping results in worse performance on MTL, GradNorm, UW, and PCGrad; however with this change, each training augmentation method now significantly outperforms the baseline MTL method.

Taskonomy. Following the experimental setup of [43], we evaluate the capacity of TAG to select task groupings on the "Segmentic Segmentation", "Depth Estimation", "Keypoint Detection", "Edge Detection", and "Surface Normal Prediction" objectives in Taskonomy. Unlike [43], our evaluation uses an augmented version of the medium Taskonomy split (2.4 TB) as opposed to the Full+ version (12 TB) to reduce computational overhead and increase reproducibility.

Our results are summarized in Figure 2 (right). Similar to our findings on CelebA, TAG continues to outperform MTL by 10.0%, GN by 7.7%, STL by 1.5%, and RG by 9.5%. Comparing TAG to HOA, we find the 2-split task grouping found by TAG to surpass the performance of that found by HOA by 2.5%, but HOA's 3-split task grouping performance is superior to that of TAG. In terms of compute, TAG is significantly more efficient than HOA, with HOA demanding an additional 2,008 TeslaV100 GPU hours to find task groupings. To put this cost into perspective, even on an 8-GPU, on-demand p3.16xlarge AWS instance, the difference in monetary expenditure between TAG and HOA would be $6,144.48. On a similar note, the performance of TAG and CS on Taskonomy are equivalent, but TAG is more efficient, requiring 140 fewer TeslaV100 GPU hours to compute task groupings.

Tasks	Improvement in Test Accuracy Relative to Random Grouping		
	optimal	(ours)	worst
a1	2.60%	2.6%	-3.03%
a2	1.53%	1.29%	-1.95%
a3	2.37%	1.72%	-3.04%
a4	2.67%	0.72%	-4.14%
a5	2.75%	2.29%	-3.87%
a6	2.04%	0.00%	-2.08%
a7	2.40%	0.74%	-2.43%
a8	1.61%	-1.59%	-1.59%
a9	8.38%	8.38%	-6.21%

Table 1: Performance of each task when trained with it's partner task *relative* to the expected performance from random groupings.

5.2 Multi-Task Ablation Studies

Does our measure of inter-task affinity correlate with optimal task groupings? To further evaluate if TAG can be used to select which tasks should train together in multi-task learning models, we evaluate its capacity to select the best co-training partner for a given task. We compare with the

optimal (best) and worst auxiliary task computed from the test set and normalize scores with respect to the expected performance of selecting an auxiliary task at random.

Our findings are summarized in Table 1 and indicates the performance of auxiliary tasks found by TAG correlates with performance of the optimal auxiliary task (Pearson's Correlation of 0.93%). However, a notable exception occurs with attribute a8 where TAG actually selects the worst partner for co-training. In this instance, and unlike every other objective in our dataset, no task manifests especially high or low inter-task affinity onto a8. The difference in normalized inter-task affinity for a8 between the best and worst partner predicted by inter-task affinity is 0.04, while the next smallest difference is 4x larger at 0.16 for a3. A table showing these differences is included in the Appendix. This case represents a limitation of TAG, and it will struggle to find the best auxiliary task for objectives like a8.

However, this weakness does not seem to significantly affect the capacity of TAG to select strong task groupings. As no other task exhibits especially high or low inter-task affinity onto a8, it is often slotted into a larger task group rather than being one of the tasks with a large difference in inter-task affinity which precipitates the formation of a new group.

Should Inter-Task Affinity Be Computed at Every Step?

It is likely that the inter-task affinity of consecutive steps are similar, and it could be the case that inter-task affinity signals at the beginning, middle, or end of training are sufficient for selecting which tasks should train together. In particular, if task relationships crystallize early in training, computing inter-task affinities during the initial stages of training may be sufficient. We evaluate both hypotheses on CelebA and our results are summarized in Table 2.

Method	Relative Performance	Relative Speedup
Every 1 Step	5.13%	1.0x
Every 5 Steps	5.13%	2.56x
Every 10 Steps	5.13%	3.19x
Every 25 Steps	4.84%	3.73x
Every 50 Steps	3.73%	3.96x
Every 100 Steps	2.06%	4.08x
First 25%	3.67%	4.00x
Middle 25%	4.31%	2.95x
Final 25%	4.02%	2.34x

Table 2: Change in total test accuracy across inference-time budget of {2-groups, 3-groups, 4-groups} relative to the expected performance from random groupings. Speedup is relative to computing inter-task affinity in each step.

Our findings indicate significant redundancy can be eliminated without degrading performance by computing inter-task affinities every 10 steps rather than every step. This modification increases training-time efficiency by 319% and is used in Section 5.1 to compute task groupings. After this threshold, signal strength decreases and leads to higher task grouping error. We also find that computing inter-task affinities in the first 25%, middle 25%, or final 25% of training degrades task-grouping performance. This result suggest the relationships among tasks change throughout training as measured by inter-task affinity. As a result, we choose to average inter-task affinity scores throughout the entirety of training to determine which tasks should train together.

Is Change in Train Loss Comparable to Change in Validation Loss? Given multi-task learning's capacity to improve generalization, computing the change in validation loss after a gradient step may capture a more informative signal as to which tasks should train together. On the other hand, certain datasets may not contain a validation split and loading a batch from the validation set every 10 steps of training would decrease efficiency.

To our surprise, the inter-task affinity scores computed on the validation set are very similar to the inter-task affinity scores computed on the training set (Pearson's Coefficient: 0.9804). For context, this similarity with computing inter-task affinities every step is greater than any other ablation in Table 2 with the exception of "Every 5 Steps" as measured by Pearson's Coefficient. Moreover, the performance of groupings found by both methods are similar: 49.574 average total error across our inference time budget of {2-groups, 3-groups, 4-groups} compared with 49.576 for groupings found on the validation set.

More formally, let X_{tr} and X_{val} be independent and identically distributed random variables for training and validation, respectively. Given the updated shared parameter $\theta_{s|i}^{t+1}$ using the gradient of task i calculated on X_{tr}, the loss of task j (s.t. $j \neq i$) yields similar expectations with respect to X_{tr} and X_{val}. That is,

$$\mathbb{E}_{\text{val}}[\mathbb{E}_{\text{tr}}[\mathcal{L}_j(X_{\text{val}}^t, \theta_{s|i}^{t+1}, \theta_j^t)]] \approx \mathbb{E}_{\text{val}}[\mathcal{L}_j(X_{\text{val}}^t, \mathbb{E}_{\text{tr}}[\theta_{s|i}^{t+1}], \theta_j^t)] = \mathbb{E}_{\text{tr}}[\mathcal{L}_j(X_{\text{tr}}^t, \mathbb{E}_{\text{tr}}[\theta_{s|i}^{t+1}], \theta_j^t)] \,.$$

The leftmost term is a joint expectation w.r.t. both X_{tr} and X_{val}, in which X_{tr} is only used to calculate the updated shared parameters $\theta^{t+1}_{s|i}$. Assuming both tasks have distinct loss functions that do not directly depend on each other during training, we can move the second expectation inside the loss function to obtain the middle term. By using the fact that X_{tr} and X_{val} are identically distributed, thus yielding the same expectation, we can move from the middle term to the rightmost term. Hence, the inter-task affinity computed on the training dataset would approximately equal the inter-task affinity computed on the validation set.

Do Changes in Hyperparameters Affect Task Groupings? We define three settings: (i) typical, the base setting, (ii) b=0.5x, or halving the batch size, and (iii) lr=2x, or increasing the learning rate by a factor of 2. To determine the ground-truth task groupings, we train all 511 combinations of task groupings from our 9-task subset of CelebA and select the groups in each setting with lowest total test error. Our aim is to assess the extent to which task groupings chosen in setting b=0.5x or lr=2x generalize to the typical setting.

Our results are summarized in Table 3. They indicate that simply changing the batch size or learning rate of a model changes which tasks should be trained together, with the groupings found by lr=2x exhibiting worse generalization than those found by b=0.5x. This result suggests that how tasks should be trained together does not simply depend on the relationships among tasks, but also on detailed aspects of the model and training. It is notably difficult to build intuition for the latter, illustrating the need to develop automated methods that can take into account these nuances.

Budget	Improvement in Test Accuracy Relative to Optimal Groupings	
	b=128x	lr=2x
2-groups	-0.61%	-1.22%
3-groups	-0.25%	-2.90%
4-groups	-1.87%	-3.21%

Table 3: Accuracy of task groupings found by b=128x and lr=2x relative to the typical setting.

6 Conclusion

In this work, we present an approach to quantify inter-task affinity in a single training run and show how this quantity can be used to systematically determine which tasks should train together in multi-task networks. Our empirical findings indicate our approach is highly competitive. It outperforms multi-task training augmentations like Uncertainty Weights, GradNorm, and PCGrad, and performs competitively with state-of-the-art task grouping methods like HOA, while improving computational efficiency by over an order of magnitude. Further, our findings are supported by extensive analysis that suggests inter-task affinity scores can find close to optimal auxiliary tasks, and in fact, implicitly measure generalization capability among tasks.

A plethora of research has been undertaken to design better multi-task learning architectures, or improve the optimization dynamics within multi-task learning systems, with relatively little work addressing the question of **which** tasks should train together in the first place. It is our hope this work renews interest in this domain, and given the sensitivity of task groupings to even small changes in hyperparameters, encourages the development of efficient and automatic methods to identify which tasks should train together in multi-task learning networks.

7 Broader Impact

Efficiently identifying task groupings in multi-task learning has the potential to save significant time and computational resources in both academic and industry environments. Despite this benefit, there are several risks associated with this work. In particular, inter-task affinities can be mistakenly interpreted as "task similarity", and incorrectly create an association and/or causation relationship among tasks with high mutual inter-task affinity scores. This association would be especially problematic for datasets involving sensitive prediction quantities related to race, gender, religion, age, status, physical traits, etc., where inter-task affinities could be mistakenly used to support an unfounded conclusion that attempts to posit similarity among tasks. That said, we believe acknowledging these risks mitigates their potential for abuse, and the benefit from this work — most notably decreasing computational resources by over an order of magnitude compared with a state-of-the-art task grouping method while performing competitively in terms of accuracy — merits its dissemination.

References

[1] 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.

[2] Alessandro Achille, Giovanni Paolini, Glen Mbeng, and Stefano Soatto. The information complexity of learning tasks, their structure and their distance. *arXiv preprint arXiv:1904.03292*, 2019.

[3] Andreas Argyriou, Theodoros Evgeniou, and Massimiliano Pontil. Convex multi-task feature learning. *Machine learning*, 73(3):243–272, 2008.

[4] Jonathan Baxter. A model of inductive bias learning. *Journal of artificial intelligence research*, 12:149–198, 2000.

[5] Shai Ben-David and Reba Schuller. Exploiting task relatedness for multiple task learning. In *Learning Theory and Kernel Machines*, pages 567–580. Springer, 2003.

[6] Lukas Brinkmeyer, Rafael Rego Drumond, Randolf Scholz, Josif Grabocka, and Lars Schmidt-Thieme. Chameleon: Learning model initializations across tasks with different schemas. *arXiv preprint arXiv:1909.13576*, 2019.

[7] Rich Caruana. Multitask learning. *Machine Learning, 28*, pages 41–75, 1997.

[8] Richard Caruana. Multitask learning: A knowledge-based source of inductive bias. In *Proceedings of the Tenth International Conference on Machine Learning*, pages 41–48. Morgan Kaufmann, 1993.

[9] Zhao Chen, Vijay Badrinarayanan, Chen-Yu Lee, and Andrew Rabinovich. Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. *arXiv preprint arXiv:1711.02257*, 2017.

[10] Zhao Chen, Jiquan Ngiam, Yanping Huang, Thang Luong, Henrik Kretzschmar, Yuning Chai, and Dragomir Anguelov. Just pick a sign: Optimizing deep multitask models with gradient sign dropout. *Advances in Neural Information Processing Systems*, 33, 2020.

[11] Paul Covington, Jay Adams, and Emre Sargin. Deep neural networks for youtube recommendations. In *Proceedings of the 10th ACM conference on recommender systems*, pages 191–198, 2016.

[12] Long Duong, Trevor Cohn, Steven Bird, and Paul Cook. Low resource dependency parsing: Cross-lingual parameter sharing in a neural network parser. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 845–850, 2015.

[13] Kshitij Dwivedi and Gemma Roig. Representation similarity analysis for efficient task taxonomy & transfer learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12387–12396, 2019.

[14] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1126–1135. JMLR. org, 2017.

[15] Erin Grant, Chelsea Finn, Sergey Levine, Trevor Darrell, and Thomas Griffiths. Recasting gradient-based meta-learning as hierarchical bayes, 2018.

[16] Michelle Guo, Albert Haque, De-An Huang, Serena Yeung, and Li Fei-Fei. Dynamic task prioritization for multitask learning. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 270–287, 2018.

[17] Pengsheng Guo, Chen-Yu Lee, and Daniel Ulbricht. Learning to branch for multi-task learning. *arXiv preprint arXiv:2006.01895*, 2020.

[18] Geoffrey E Hinton and David C Plaut. Using fast weights to deblur old memories. In *Proceedings of the ninth annual conference of the Cognitive Science Society*, pages 177–186, 1987.

[19] Siyu Huang, Xi Li, Zhi-Qi Cheng, Zhongfei Zhang, and Alexander Hauptmann. GNAS: A greedy neural architecture search method for multi-attribute learning. *2018 ACM Multimedia Conference on Multimedia Conference - MM '18*, 2018.

[20] Pavel Izmailov, Dmitrii Podoprikhin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. *arXiv preprint arXiv:1803.05407*, 2018.

[21] Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in neural information processing systems*, pages 315–323, 2013.

[22] Dmitry Kalashnikov, Jacob Varley, Yevgen Chebotar, Benjamin Swanson, Rico Jonschkowski, Chelsea Finn, Sergey Levine, and Karol Hausman. Mt-opt: Continuous multi-task robotic reinforcement learning at scale. *arXiv preprint arXiv:2104.08212*, 2021.

[23] Zhuoliang Kang, Kristen Grauman, and Fei Sha. Learning with whom to share in multi-task feature learning. In *ICML*, 2011.

[24] Andrej Karpathy. Multi-task learning in the wilderness. ICML, 2019. URL `https://slideslive.com/38917690/multitask-learning-in-the-wilderness`.

[25] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7482–7491, 2018.

[26] Taesup Kim, Jaesik Yoon, Ousmane Dia, Sungwoong Kim, Yoshua Bengio, and Sungjin Ahn. Bayesian model-agnostic meta-learning. *arXiv preprint arXiv:1806.03836*, 2018.

[27] Abhishek Kumar and Hal Daume III. Learning task grouping and overlap in multi-task learning. *arXiv preprint arXiv:1206.6417*, 2012.

[28] Giwoong Lee, Eunho Yang, and Sung Hwang. Asymmetric multi-task learning based on task relatedness and loss. In *International conference on machine learning*, pages 230–238. PMLR, 2016.

[29] Hae Beom Lee, Eunho Yang, and Sung Ju Hwang. Deep asymmetric multi-task feature learning. In *International Conference on Machine Learning*, pages 2956–2964. PMLR, 2018.

[30] Xi Lin, Hui-Ling Zhen, Zhenhua Li, Qing-Fu Zhang, and Sam Kwong. Pareto multi-task learning. In *Advances in Neural Information Processing Systems*, pages 12037–12047, 2019.

[31] Xingyu Lin, Harjatin Singh Baweja, George Kantor, and David Held. Adaptive auxiliary task weighting for reinforcement learning. *Advances in neural information processing systems*, 32, 2019.

[32] 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 *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 19–34, 2018.

[33] Liyang Liu, Yi Li, Zhanghui Kuang, Jing-Hao Xue, Yimin Chen, Wenming Yang, Qingmin Liao, and Wayne Zhang. Towards impartial multi-task learning. In *International Conference on Learning Representations*, 2021. URL `https://openreview.net/forum?id=IMPnRXEWpvr`.

[34] Shikun Liu, Edward Johns, and Andrew J Davison. End-to-end multi-task learning with attention. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1871–1880, 2019.

[35] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. *2015 IEEE International Conference on Computer Vision (ICCV)*, Dec 2015.

[36] 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. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jul 2017.

[37] Ishan Misra, Abhinav Shrivastava, Abhinav Gupta, and Martial Hebert. Cross-stitch networks for multi-task learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3994–4003, 2016.

[38] Yu Nesterov. A method of solving a convex programming problem with convergence rate $\mathcal{O}(1/k^2)$. In *Sov. Math. Dokl*, volume 27, pages 372–376, 1983.

[39] Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018.

[40] Sebastian Ruder. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*, 2017.

[41] Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016.

[42] Ozan Sener and Vladlen Koltun. Multi-task learning as multi-objective optimization. In *Advances in Neural Information Processing Systems*, pages 527–538, 2018.

[43] Trevor Standley, Amir R Zamir, Dawn Chen, Leonidas Guibas, Jitendra Malik, and Silvio Savarese. Which tasks should be learned together in multi-task learning? *arXiv preprint arXiv:1905.07553*, 2019.

[44] Ximeng Sun, Rameswar Panda, and Rogerio Feris. Adashare: Learning what to share for efficient deep multi-task learning. *arXiv preprint arXiv:1911.12423*, 2019.

[45] Mihai Suteu and Yike Guo. Regularizing deep multi-task networks using orthogonal gradients. *arXiv preprint arXiv:1912.06844*, 2019.

[46] Simon Vandenhende, Stamatios Georgoulis, Bert De Brabandere, and Luc Van Gool. Branched multi-task networks: deciding what layers to share. *arXiv preprint arXiv:1904.02920*, 2019.

[47] Zirui Wang, Yulia Tsvetkov, Orhan Firat, and Yuan Cao. Gradient vaccine: Investigating and improving multi-task optimization in massively multilingual models. *arXiv preprint arXiv:2010.05874*, 2020.

[48] Sen Wu, Hongyang R Zhang, and Christopher Ré. Understanding and improving information transfer in multi-task learning. *arXiv preprint arXiv:2005.00944*, 2020.

[49] Yongxin Yang and Timothy M. Hospedales. Trace norm regularised deep multi-task learning. *arXiv preprint arXiv:1606.04038*, 2016.

[50] Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. Gradient surgery for multi-task learning. *arXiv preprint arXiv:2001.06782*, 2020.

[51] Amir R. Zamir, Alexander Sax, William Shen, Leonidas Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling task transfer learning. *2018 IEEE/CVF Conference*, Jun 2018.

[52] Michael Zhang, James Lucas, Jimmy Ba, and Geoffrey E Hinton. Lookahead optimizer: k steps forward, 1 step back. In *Advances in Neural Information Processing Systems*, pages 9597–9608, 2019.

[53] Yu Zhang and Qiang Yang. A survey on multi-task learning. *arXiv preprint arXiv:1707.08114*, 2017.

[54] Xiangyun Zhao, Haoxiang Li, Xiaohui Shen, Xiaodan Liang, and Ying Wu. A modulation module for multi-task learning with applications in image retrieval. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 401–416, 2018.

[55] Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing He. A comprehensive survey on transfer learning. *Proceedings of the IEEE*, 2020.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See Section 5.2 and in particular sub-header "Does our measure of inter-task affinity correlate with optimal task groupings."
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] See Section 7.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...

(a) Did you state the full set of assumptions of all theoretical results? [Yes] See Section 4.3 and the Appendix

(b) Did you include complete proofs of all theoretical results? [Yes] See Section 4.3 and the Appendix

3. If you ran experiments...

(a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See the Appendix which details our experimental design. Additionally, we release our code in the supplementary material.

(b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See the Appendix or refer to the released code.

(c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See Appendix where we report the error measurements averaged over multiple runs for CelebA. Given the amount of computational resources to run Tasknomy, we did not compute error measurements for this dataset.

(d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Section 5 and Figure 2.

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

(a) If your work uses existing assets, did you cite the creators? [Yes] See Section 5.

(b) Did you mention the license of the assets? [Yes] See the Appendix.

(c) Did you include any new assets either in the supplemental material or as a URL? [Yes] Our code is released in the supplementary material.

(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A] We only use existing and popular multi-task learning benchmark datasets.

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A] We only use existing and popular multi-task learning benchmark datasets.

5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]