

AUXILIARY TASK UPDATE DECOMPOSITION: THE GOOD, THE BAD AND THE NEUTRAL

Anonymous authors

Paper under double-blind review

ABSTRACT

While deep learning has been very beneficial in data-rich settings, tasks with smaller training set often resort to pre-training or multitask learning to leverage data from other tasks. In this case, careful consideration is needed to select tasks and model parameterizations such that updates from the auxiliary tasks actually help the primary task. We seek to alleviate this burden by formulating a model-agnostic framework that performs fine-grained manipulation of the auxiliary task gradients. We propose to decompose auxiliary updates into directions which help, damage or leave the primary task loss unchanged. This allows weighting the update directions differently depending on their impact on the problem of interest. We present a novel and efficient algorithm for that purpose and show its advantage in practice. Our method leverages efficient automatic differentiation procedures and randomized singular value decomposition for scalability. We show that our framework is generic and encompasses some prior work as particular cases. Our approach consistently outperforms strong and widely used baselines when leveraging out-of-distribution data for Text and Image classification tasks. In particular, with only 1% of Imagenet, our approach improves AUC by 1.9 points over an Imagenet pre-trained model on the ChexPert medical imaging transfer task.

1 INTRODUCTION

Multitask learning (Caruana, 1997) and pretraining (Devlin et al., 2018; Caron et al., 2019) have transformed machine learning by allowing downstream tasks with small training sets to benefit from statistical regularities from data-rich related tasks (Collobert & Weston, 2008; Zhang et al., 2014; Liu et al., 2019; Kornblith et al., 2019). Despite these advances, leveraging the mixing of tasks is still an art left to the practitioner. When one is interested in a *primary* task, it is unclear how to select helpful auxiliary tasks, an appropriate parameter sharing architecture and a good way to filter out auxiliary data which might be detrimental to the primary tasks. Without careful choices, pre-training might hurt end-task performance (Gururangan et al., 2020) or have limited impact (Raghu et al., 2019).

Prior work has examined these problems and proposed solutions, either to choose auxiliary tasks depending on their impact on the primary task (Du et al., 2018; Lin et al., 2019) or to equalize the impact of updates across tasks (Sener & Koltun, 2018; Chen et al., 2018; Hessel et al., 2019). Recently, several approaches (Sinha et al., 2018; Suteu & Guo, 2019; Yu et al., 2020) have been proposed that attempt to minimize interference between the updates across tasks. Our work builds on this direction, but unlike these previous approaches, we do not consider a symmetric view of multi-task learning in the sense that our goal is not to train a model performing well on all tasks. Instead, we focus on improving generalization for a single task, the primary task, and the other tasks, the auxiliary tasks are considered only through their impact on the problem of interest.

For that purpose, we introduce a framework which decomposes the gradient updates from the auxiliary tasks according to their impact on the primary task. We analyze the auxiliary task gradients in the subspace spanned by the primary task per-example gradients. This allows us to decompose auxiliary gradients into three components : components that help, interfere or have no impact on the primary task according to the Taylor expansion of the expected primary loss. This decomposition allows us to re-weight each component differently prior to the update. Our framework enables us to treat each auxiliary update differently depending on its impact on the task of interest and it

encompasses prior methods such as classical multitask learning (Caruana, 1997) or more novel gradient surgery techniques (Yu et al., 2020). To achieve a tractable approach, we introduce an efficient, robust algorithm (ATTITTUD, Auxiliary Task Training with Influence from Target Task Update Direction) to estimate the subspace spanned by the primary task gradients in an online manner and decompose the auxiliary updates appropriately. As a result, we can integrate our approach with the stochastic training of large neural networks in various contexts.

The contribution of our work is four-fold. To our knowledge, this paper proposes the first approach to adapt auxiliary gradients using a decomposition built from the span of the primary task Jacobian. In order to scale this approach to deep neural nets, we contribute a tractable and efficient algorithm called ATTITTUD that leverages insights from randomized linear algebra and automatic differentiation such as the R-operator (Pearlmutter, 1994). As our third contribution, we show that the fine-grained manipulation of the auxiliary task gradients under ATTITTUD, represents a unified framework that encompasses several previous approaches to asymmetrical task learning as special cases. Finally, we demonstrate the efficacy of our approach in both data-rich and data-starved primary tasks, over both images and textual data.

2 RELATED WORK

Methods to leverage data outside of the task of interest are popular in machine learning since the inception of multitask learning (Caruana, 1997; Ruder, 2017; Vandenhende et al., 2020). These methods address multiple task simultaneously and have been successful in various application domains (Collobert & Weston, 2008; Zhang et al., 2014; Misra et al., 2016). The optimization problem induced by multitask learning is difficult and solutions have been proposed for the various difficulties, including dealing with task gradients of different magnitude (Sener & Koltun, 2018; Chen et al., 2018; Hessel et al., 2019), or interfering with each others (Sinha et al., 2018; Suteu & Guo, 2019; Yu et al., 2020). The specific problem of interference has been studied extensively in the context of continual learning. Continual learning visits task in sequence and update interference is particularly problematic as it yields newer tasks to damage previously mastered tasks. In particular, a family of methods to project the gradient of the new tasks to be orthogonal to the gradient of the previous tasks has been proposed (Lopez-Paz & Ranzato, 2017; Chaudhry et al., 2018; Farajtabar et al., 2019).

Different from many previous approaches, we are not interested in addressing multiple tasks per se. In our setting, only the *primary* task matters and the other *auxiliary* task have the sole role of improving generalization on the primary task. This is the setting considered by Du et al. (2018); Lin et al. (2019), who favor auxiliary tasks whose gradient directions are helpful to the primary task. Unlike these works that use coarse properties like the cosine similarity between averaged gradients, our approach allows fine-grained gradient manipulation within a subspace. Also, in our case, we do not distinguish between the different auxiliary tasks. Instead, we aim at correcting every auxiliary gradient in the same manner to improve the loss on the primary task. This type of gradient correction is related to Yu et al. (2020), which considers projecting multi-task gradients such that the directions of disagreement are removed. This method is actually a special case of our framework.

Our work also shares some similarities with data selection and domain adaptation approaches. In this case, the training data comes from a single task but its distribution is different from the validation/test distribution (Moore & Lewis, 2010; Axelrod et al., 2011; Ngiam et al., 2018). This classical problem has recently been addressed by sampling training points whose gradient aligns well with the expected validation gradient (Wang et al., 2020b,a). Instead of sampling individual points based on an estimated distribution of how helpful they will be to the primary task, our work avoids the use (and inherent challenges) of this reinforcement learning approach by operating on batch gradients of groups of points.

Our primary task/auxiliary task setting is also related to the pre-training then fine-tuning paradigm in which the auxiliary tasks are visited first (pre-training) to give an initialization for training on the primary task (fine-tuning). These methods have been very successful in settings where primary task data are rare. In particular, it is common to first rely on an unsupervised task over very large datasets prior to fine tuning over a supervised task (Devlin et al., 2018; Liu et al., 2019; Kornblith et al., 2019; Yang et al., 2019; Song et al., 2019; Caron et al., 2018).

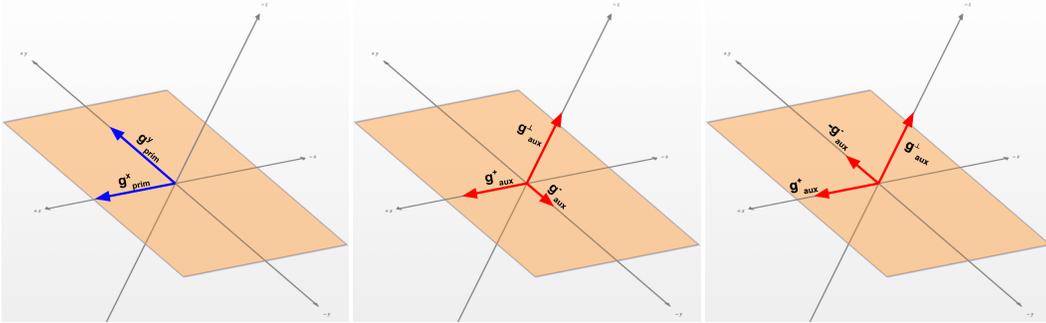


Figure 1: Example gradient manipulation in the 2-D $x - y$ plane with ATTITUD. ATTITUD can operate in any n -dimensional subspace. **Left:** Primary task gradient \mathbf{g}_{prim} decomposed along the 3 Dimensions x, y and z . **Mid:** Decomposed Auxiliary task gradient \mathbf{g}_{aux} . We label the x component of \mathbf{g}_{aux} as positive since it agrees (in direction) with the x component of \mathbf{g}_{prim} . Since the y component of \mathbf{g}_{aux} is in the opposite direction as that of \mathbf{g}_{prim} , this is assigned a negative label. **Right:** Corresponds to $\tilde{\mathbf{g}}_{aux}$ obtained by applying $\eta_{aux} = (1.0, 1.0, -1.0)$. We flip the conflicting gradient direction to agree with our primary task. This is just one configuration achievable under our framework.

3 AUXILIARY TASK UPDATE DECOMPOSITION

This section introduces a new method to improve generalization on a primary task T^* using training data from auxiliary tasks $\mathbb{T} = \{T_1, \dots, T_n\}$, where $\theta \in \mathbb{R}^D$ denote the parameters shared by all tasks. Our approach leverages gradient updates from the auxiliary tasks, but unlike the traditional approach, we decompose these gradients to maximize their usefulness to T^* . Precisely, we decompose the auxiliary task gradients into directions which decrease a first-order approximation of the primary task loss, increase it or have no effect. This decomposition allows weighting these three directions differently when learning from the auxiliary tasks.

In order to decompose the auxiliary gradient, we must collect more fine-grained statistics about the primary task. At each training step, we collect the gradient of the loss with respect to θ for individual examples from the primary task, $\{\nabla_{\theta} \mathcal{L}_i^{prim}, \forall i\}$. The span of these vectors,

$$\mathcal{S} = \text{Span}\{\nabla_{\theta} \mathcal{L}_i^{prim}, \forall i\}$$

defines a subspace in which any linear combination of primary task gradients lies, including the gradient of the expected primary task loss, i.e. $\mathbf{g}_{prim} = \mathbb{E}(\nabla_{\theta} \mathcal{L}_i^{prim}) \in \mathcal{S}$. If we define the orthogonal complement of \mathcal{S} as \mathcal{S}^{\perp} , any vector $v \in \mathcal{S}^{\perp}$, is therefore orthogonal to \mathbf{g}_{prim} , i.e. $v \cdot \mathbf{g}_{prim} = 0$. This means that adding such a vector to the parameters has no impact on the expected primary task loss, according the order-1 Taylor expansion of \mathcal{L}^{prim} , i.e.

$$\mathcal{L}^{prim}(\theta + v) \simeq \mathcal{L}^{prim}(\theta) + v \cdot \mathbf{g}_{prim} = \mathcal{L}^{prim}(\theta).$$

We propose to project auxiliary task gradients onto \mathcal{S} and \mathcal{S}^{\perp} . This allow to distinguish between the directions of the auxiliary task updates which impact the primary task loss and those which do not. If we denote the averaged auxiliary task gradient as $\mathbf{g}_{aux} = \mathbb{E}(\nabla_{\theta} \mathcal{L}_i^{aux})$, we can decompose this gradient as $\mathbf{g}_{aux} = \mathbf{g}_{aux}^{\pm} + \mathbf{g}_{aux}^{\perp}$, where $\mathbf{g}_{aux}^{\pm} \in \mathcal{S}$ is the portion of the gradient that lies in the span of the primary task example gradients and $\mathbf{g}_{aux}^{\perp} \in \mathcal{S}^{\perp}$ is the portion that lies outside of it. Since $\mathbf{g}_{aux}^{\perp} \in \mathcal{S}^{\perp}$, it is orthogonal to the average primary task gradient and parameter updates along the direction of \mathbf{g}_{aux}^{\perp} are expected to have limited impact on the primary task loss. On the other hand, updates along the direction of \mathbf{g}_{aux}^{\pm} can potentially improve or damage the averaged primary task loss. This component deserves a more careful treatment.

For that purpose, we introduce $\{u_i, i = 1, \dots, k\}$ an orthonormal basis of \mathcal{S} . In this basis, we can measure if the components of \mathbf{g}_{aux}^{\pm} agree or disagree with \mathbf{g}_{prim} . We say that the two gradients agree along u_i iff $\text{sign}(\mathbf{g}_{aux}^{\pm} \cdot u_i) = \text{sign}(\mathbf{g}_{prim} \cdot u_i)$. This means that we can decompose $\mathbf{g}_{aux}^{\pm} = \mathbf{g}_{aux}^+ + \mathbf{g}_{aux}^-$ where \mathbf{g}_{aux}^+ refers to the projection of \mathbf{g}_{aux}^{\pm} onto the basis vectors where \mathbf{g}_{aux}^{\pm} and \mathbf{g}_{prim} agree. By this

decomposition, \mathbf{g}_{aux}^+ helps the primary task, $\mathbf{g}_{aux}^+ \cdot \mathbf{g}_{prim} > 0$, while \mathbf{g}_{aux}^- interfere with the primary task, $\mathbf{g}_{aux}^- \cdot \mathbf{g}_{prim} < 0$.

Guided by the primary task, we can therefore decompose the auxiliary task gradient as

$$\mathbf{g}_{aux} = \mathbf{g}_{aux}^\perp + \mathbf{g}_{aux}^+ + \mathbf{g}_{aux}^- \quad (1)$$

which is described on Fig 1. Our approach proposes to re-weight differently the components of \mathbf{g}_{aux} , i.e.

$$\tilde{\mathbf{g}}_{aux} = \eta_\perp \mathbf{g}_{aux}^\perp + \eta_+ \mathbf{g}_{aux}^+ + \eta_- \mathbf{g}_{aux}^- \quad (2)$$

where $\boldsymbol{\eta}_{aux} = (\eta_\perp, \eta_+, \eta_-)$ are hyper-parameters adjusting the auxiliary gradient according to the impact on the main task. If we also wish to include the primary task gradient in descent, as with multitasking, we can introduce $\boldsymbol{\eta}_{prim}$ as a scalar control variable to control its weighting.

A consequence of introducing $\boldsymbol{\eta}_{aux}$ is that specific configurations lead us to gradient updates that are guaranteed to do no harm to both tasks. This is captured by Theorem 1 below.

Theorem 1. *Let $\mathcal{L}^{aux}(\theta_t)$ and $\mathcal{L}^{prim}(\theta_t)$ represent the full batch losses of the auxiliary tasks and primary task respectively at step t . We assume the gradients of \mathcal{L}^{aux} and \mathcal{L}^{prim} are Lipschitz continuous with constant $L > 0$. Following the update rule : $\theta_{t+1} = \theta_t - \alpha \cdot \tilde{\mathbf{g}}_{aux}$, where $\alpha \leq \frac{1}{L}$ is the learning rate, we are guaranteed :*

$$\begin{aligned} \mathcal{L}^{aux}(\theta_{t+1}) &\leq \mathcal{L}^{aux}(\theta_t) \\ \mathcal{L}^{prim}(\theta_{t+1}) &\leq \mathcal{L}^{prim}(\theta_t) \end{aligned}$$

If $\eta_- = 0$ and $\eta_\perp, \eta_+ \geq 0$

Proof. See Appendix A □

This theorem focuses on a single update and guarantees progress on both auxiliary and primary tasks. However, our asymmetric scenario is not interested in improving the auxiliary tasks per se and is amenable to more aggressive settings. Ideally we want gradient updates during pre-training with \mathbb{T} to not only do-no-harm to T^* when applied downstream but also to be along descent directions that are maximally beneficial to T^* . We can consider $\eta_- < 0$ as in Fig 1. Reversing the direction of \mathbf{g}_{aux}^- by setting $\eta_- < 0$ preserves the descent guarantee on $\mathcal{L}_{prim}(\theta_{t+1})$ but no longer ensures descent on $\mathcal{L}_{aux}(\theta_{t+1})$. There are other interesting settings for our control parameters. One can recover the original gradient \mathbf{g}_{aux} with $\eta_\perp = \eta_- = \eta_+ = 1.0$. One can choose to drop gradients orthogonal to the primary task gradient span with $\eta_\perp = 0.0$, or ignore those which conflict with the main task by setting $\eta_- = 0.0$.

Relationships to other approaches Our framework is generic and encompasses other approaches as a particular case. One can train solely on the primary task by selecting $\boldsymbol{\eta}_{aux} = (0.0, 0.0, 0.0)$ and $\boldsymbol{\eta}_{prim} = 1.0$. Classical multitasking corresponds to $\boldsymbol{\eta}_{aux} = (1.0, 1.0, 1.0)$ and $\boldsymbol{\eta}_{prim} > 0.0$, while classical pre-training corresponds to performing a first phase with $\boldsymbol{\eta}_{aux} = (1.0, 1.0, 1.0)$ and $\boldsymbol{\eta}_{prim} = 0.0$. Interestingly, our formulation introduces novel variants of pre-training, for instance, one can consider pre-training with only auxiliary gradients helpful to the primary task, $\boldsymbol{\eta}_{aux} = (0.0, 1.0, 0.0)$ and $\boldsymbol{\eta}_{prim} = 0.0$, followed by fine-tuning with $\boldsymbol{\eta}_{aux} = (0.0, 0.0, 0.0)$ and $\boldsymbol{\eta}_{prim} = 1.0$.

Our approach also instantiates PCGrad (Yu et al., 2020) as a particular case. This method was introduced to address the issue of conflicting gradients in multitask settings. PCGrad orthogonalizes the gradients of each task and removes conflicting gradients. To recover PCGrad under our approach, note that it is equivalent to a specific choice of our decomposition in the 1-D subspace spanned by the \mathbf{g}_{prim} . PCGrad then removes components of \mathbf{g}_{aux} that conflict with \mathbf{g}_{prim} which is equivalent to $\boldsymbol{\eta}_{aux} = (\alpha_{aux}, \alpha_{aux}, 0.0)$ and $\boldsymbol{\eta}_{aux} = \alpha_{prim}$.

4 IMPLEMENTATION

Equation 2 requires selecting a basis for the span of primary task gradients. Multiple choices are possible to define the basis $\{u_i\}$, to represent the span at each optimization time-step. This choice is important since the components of \mathbf{g}_{aux}^\pm are labeled positive or negative depending on how they

agree with the projection of the averaged primary task gradient onto the same basis. A natural choice is to select the basis as the singular vectors of the matrix of primary task per-example gradients $\mathbf{J}^* \in \mathbb{R}^{m \times D}$, also known as the Jacobian. To improve efficiency and prevent over-fitting on a few examples, we consider the span defined by the k largest principal vectors of \mathbf{J}^* . Using the principal vectors as directions of descent instead of the mean induces a more robust algorithm since the mini-batch average gradient is susceptible to outliers and skew from replicated data-points. To the best of our knowledge, we are the first to propose using the singular vectors of \mathbf{J}^* as directions of descent. We leave the theoretical implications of this algorithm to future work but note that its variance reduction properties may induce generalization benefits (Namkoong & Duchi, 2017).

We also consider alternative choices of bases as baselines, including the canonical parameter basis. This choice will examine the sign of every parameter update to verify whether it agrees with \mathbf{g}_{prim} . Whilst Theorem 1 holds irrespective of the choice of basis, its proof reveals that the amount of progress made on each loss depends on the choice of basis. Specifically, the reduction in $\mathcal{L}^{prim}(\theta_{t+1}), \mathcal{L}^{aux}(\theta_{t+1})$ after a gradient step along $\tilde{\mathbf{g}}_{aux}$ is proportional to the fraction of the norms of \mathbf{g}_{prim} and \mathbf{g}_{aux} captured by the subspace spanned by our choice of basis. We evaluate this fraction for different choice of basis in our experiments (see Appendix C).

We are interested in applying our approach to the training of large neural networks and must consider a scalable algorithmic solution. As stochastic optimization is prevalent in this setting, we construct subspace \mathbf{S} from a mini-batch of primary task data. Similarly, the expected gradients \mathbf{g}_{prim} and \mathbf{g}_{aux} are defined over a mini-batch. Instead of computing the singular value decomposition (SVD) of $\{\nabla_{\theta} \mathcal{L}_i^{prim}, \forall i\}$ exactly, we rely on a randomized approximation (Halko et al., 2011; Rokhlin et al., 2010; Nakatsukasa, 2017). This method does not require instantiating the vectors $\{\nabla_{\theta} \mathcal{L}_i^{prim}, \forall i\}$ and only needs a low dimensional projection onto a random subspace. This is advantageous for high dimensional cases, i.e. when the number of model parameters is large. In our case, this method also allows us to benefit from memory-efficient computation of Jacobian Vector product using the R-operator (Pearlmutter, 1994) offered by automatic differentiation packages (Baydin et al., 2015) like Pytorch (Paszke et al., 2017). This means that we can compute SVD with a limited computational and memory burden, albeit without sacrificing approximation accuracy (Nakatsukasa, 2017). Additionally, we do not recompute the basis at every optimization step but at every n steps, which is efficient when training with small updates, e.g. when small learning rates and gradient clipping are used (Pascanu et al., 2013).

We study the impact of these choices in practice in Section 6. Putting it all together results in the ATTITTUD algorithm, Auxiliary Task Training with Influence from Target Task Update Direction, shown as Algorithm 1. The sub-procedure `randomized_lowrank_approx` is detailed in Appendix B as Algorithm 2.

Algorithm 1: ATTITTUD : Construct Auxiliary Task Surrogate Gradient

Require : $\mathbf{g}_{aux}, \mathbf{J}^*$: Auxiliary task average gradient, primary task Jacobian

Require : $\boldsymbol{\eta}_{aux} = (\eta_{\perp}, \eta_{+}, \eta_{-})$: Auxiliary task control parameters

Require : k : Size of subspace

$$\mathbf{g}_{prim} = \frac{1}{m} \sum_{i=1}^m \mathbf{J}_{i,:}^*$$

$$\mathbf{V} \leftarrow \text{randomized_lowrank_approx}(\mathbf{J}^*, k)$$

$$\mathbf{p}_{prim}, \mathbf{p}_{aux} = \mathbf{V}_t (\mathbf{g}_{prim})^T, \mathbf{V}_t (\mathbf{g}_{aux})^T$$

// \circ is the hadamard product operator

$$\mathbf{p}_{aux}^+, \mathbf{p}_{aux}^- = \left(\mathbf{1}_{[\mathbf{p}_{prim} \circ \mathbf{p}_{aux} \geq 0]} \right) \circ \mathbf{p}_{aux}, \left(\mathbf{1}_{[\mathbf{p}_{prim} \circ \mathbf{p}_{aux} < 0]} \right) \circ \mathbf{p}_{aux}$$

// Calculate the decomposition components

$$\mathbf{g}_{aux}^+, \mathbf{g}_{aux}^- = (\mathbf{p}_{aux}^+)^T \mathbf{V}, (\mathbf{p}_{aux}^-)^T \mathbf{V}$$

// Calculate the out of span component

$$\mathbf{g}_{aux}^{\perp} = \mathbf{g}_{aux} - (\mathbf{g}_{aux}^+ + \mathbf{g}_{aux}^-)$$

$$\tilde{\mathbf{g}}_{aux} = (\eta_{\perp} \cdot \mathbf{g}_{aux}^{\perp}) + (\eta_{+} \cdot \mathbf{g}_{aux}^+) + (\eta_{-} \cdot \mathbf{g}_{aux}^-)$$

Return : $\tilde{\mathbf{g}}_{aux}$: Auxiliary task surrogate gradient

5 EXPERIMENTAL SETUP

We compare ATTITUDE with previous methods on a variety of tasks and domains. We rely on both text and image classification tasks to conduct our analysis. We also present ablation experiments to explain the impact of hyper-parameter selection.

Text Classification. We apply our method on binary sentiment classification. We consider the Amazon Helpfulness (McAuley et al., 2015) and Imdb Movie Review (Maas et al., 2011) tasks. The Amazon Helpfulness task splits text reviews into 115k/5k/25k documents for train-validation-test split whilst the Imdb Review dataset has a 20k/5k/25k split. The Imdb Review task also has 50k unlabeled reviews as extra data which we utilize.

For our models we build on top of Gururangan et al. (2020)’s work where they introduce Task-Adaptive Pre-training (TAPT). TAPT further pre-trains a generic model, Roberta (Liu et al., 2019), by performing Masked Language Modelling, MLM, (Devlin et al., 2018) on the task specific data (ignoring the labels) before doing supervised learning with the same data. We replicate Gururangan et al. (2020)’s experimental setup and re-use their hyper-parameters for our experiments. We use the TAPT task as our auxiliary task. We extend TAPT to use our method by modifying the TAPT gradient with guidance from the supervised-learning task gradients. As baselines, we compare against TAPT and cross-TAPT: where we swap the masked language modelling pre-training data for the two tasks. Cross-TAPT is a setting where one uses out-of-distribution data for pre-training.

Image Classification. We apply our method to both high-resource and limited-data image classification tasks. We use the Cifar100 dataset (Krizhevsky et al., 2009) to explore the high-resource setting. We follow Rosenbaum et al. (2017) and treat each of the 20 super-classes / coarse labels of Cifar100 as a separate task. In our asymmetrical task setting, each of the 20 tasks is treated as a primary task, whilst the remaining 95 classes are grouped into a single auxiliary task. Thus, for each coarse label, we have an auxiliary 95-way classification task and a 5-way primary classification task. Moving forward, we refer to this setting as MultiCifar100.

We use a down-sampled version of Cifar10 (Krizhevsky et al., 2009) as a low-resource setting. Specifically, we rely on Cat-vs-Dog for the primary task and use the remaining 8 classes for the auxiliary task. Our auxiliary task is therefore an 8-way classification task where each class has 5,000 examples. We restrict the Cat and Dog classes to only 50 training examples from each class. We use the low-resource setting to compare against other methods and for our ablation study.

For these vision experiments, we use a WideResNet-22 architecture (Zagoruyko & Komodakis, 2016) with a depth of $k = 4$. We compare our method to 4 different baselines : no pre-training, vanilla pre-training, multitasking and PCGrad (Yu et al., 2020). For MultiCifar100 we do not use the original architecture proposed in Rosenbaum et al. (2017) and used by (Yu et al., 2020) since it relies on Reinforcement Learning.

Medical Imaging Transfer. We apply our method to cross-domain transfer for low-resource medical image classification. Specifically, we use 5k training examples from the ChexPert Dataset (Irvin et al., 2019) as our primary task and seek to identify 5 different thoracic pathologies: atelectasis, cardiomegaly, consolidation, edema and pleural effusion. This setup has been used in several cross-domain pretraining studies (Raghu et al., 2019; Jaiswal et al., 2019). Note that since we do not have access to the test set for this task, we use the validation set (231 images) as a proxy test set, and sample 100 images from the training data as a new validation set. We rely on generic photographs (Imagenet) as an auxiliary task (Deng et al., 2009). We use Tiny Imagenet Dataset (Le & Yang, 2015), a subset of Imagenet which consists of 500 examples each from 200 classes, instead of training on full Imagenet. All approaches are applied to the Resnet18 model (He et al., 2016) trained with Adam (Kingma & Ba, 2014).

For all our experiments, we select the auxiliary task control parameters η_{aux} within $\{(1.0, 1.0, -1.0), (1.0, 1.0, 0.0), (1.0, 0.0, -1.0), (1.0, 0.0, 0.0)\}$. For settings where we compare against multi-tasking, we select η_{prim} within a small subset of the settings that worked best with multitasking baseline experiments. These choices limit the overhead of hyper-parameter search but still allow us to show the empirical advantage of our method. More experimental details are available in Appendix C

	Imdb	Imdb + Amazon MLM	Amazon	Amazon + Imdb MLM
Roberta	95.4 \pm 0.14	-	67.0 \pm 0.50	-
TAPT	96.1 \pm 0.11	95.1 \pm 0.10	70.3 \pm 0.87	67.8 \pm 0.46
Ours	96.1 \pm 0.09	95.4 \pm 0.03	70.1 \pm 1.13	68.5 \pm 1.01

Table 1: Results on Text Classification measured by F1. Experiments are averaged over 5 runs.

6 RESULTS AND DISCUSSION

Text Classification. Table 1 shows the results for text classification. When the same data is used both for the auxiliary task of MLM and the primary classification task, TAPT and ATTITUD both bring a similar improvement over Roberta (Imdb, Amazon columns). When different data is used for the auxiliary task and the primary task (Imdb + Amazon MLM, Amazon + Imdb MLM columns), TAPT does not perform as well as ATTITUD. This highlights the advantage of ATTITUD when the auxiliary task data distribution differ from the primary task distribution.

Image Classification. Our results are presented in Table 2. Both for MultiCifar100 (high resource setting) and Cifar10-Cat-vs-Dog (low resource setting), ATTITUD shows a strong improvement over baselines. In general, we find that primary-task aware pre-training (Multitasking, PCGrad, Ours) is better than vanilla pre-training which also performs better than having no pre-training at all. For MultiCifar100, we find that using $\eta_{aux} = (1.0, 1.0, -1.0)$, $\eta_{prim} = 0.1$ worked best for 11 out of the 20 Cifar100 super-classes tasks. Note that $\eta_{aux} = (1.0, 1.0, -1.0)$ is an aggressive but novel configuration we introduce. Multitask learning and PCGrad produce better models on 6 and 3 tasks respectively. In the low-resource Cat-vs-Dog, setting ATTITUD produces a bigger boost in performance compared to baselines, with the best performing configuration being $\eta_{aux} = (1.0, 0.0, 0.0)$, $\eta_{prim} = 0.01$. We posit that this configuration is successful because removal of the in-span components makes overfitting less likely. Applying the out-of-span components means the model learns features that do not harm the loss of the current mini-batch but could be useful later.

Method	MultiCifar100	Cifar10-Cat-vs-Dogs
No-Pretraining	57.6	53.6 \pm 2.26
Vanilla Pre-training	70.2	64.5 \pm 1.26
PCGrad	75.6	64.2 \pm 1.10
Multitask	75.5	65.3 \pm 1.35
Ours	76.1	67.1 \pm 1.31

Table 2: Average Accuracy on MultiCifar100 and Cat-vs-Dog Cifar10 tasks. Cat-vs-Dog experiments are averaged over 5 runs

Method	Average AUC Across 5 Pathologies
No-Pretraining	78.3 \pm 0.87
Vanilla Pre-training (100% ImageNet)	81.4 \pm 1.34
Ours (0.5% ImageNet)	82.8 \pm 0.70
Ours (1% Imagenet)	83.3 \pm 0.71

Table 3: Results on ChexPert-5k task measured by average AUC (Area Under Roc-Curve). All experiments are averaged over 5 runs

Medical Imaging Transfer. Table 3 shows our results on the CheXPert multi-label classification task. Per-pathology breakdowns are in Appendix C. Doing no pre-training at all performs worst. Our method easily outperforms using a pre-trained Resnet18 model which has seen all of Imagenet. With just 50k examples ($\sim 0.5\%$ of Imagenet) ATTITUD is able to provide improved results, validating the need for the end-task to inform pre-training when known in advance. ATTITUD allows us to effectively leverage more data, as we get a performance improvement by doubling the auxiliary task dataset size.

Subspace	Canonical	Random	Unit_avg_grad	Randomized_SVD
Average Acc.	51.42 ± 2.09	58.72 ± 2.68	59.13 ± 2.08	62.2±4.00

Table 4: Experiment conducted on Cat-vr-Dog Cifar10 dataset for different choices of subspace basis. We use $k = 5$ for Random and Randomized_SVD.

Ablation Study. Our approach relies on the top- k singular vectors from *randomized_svd* to define the basis to identify the positive and negative component of the auxiliary task gradient, see Section 4. This method is more accurate than several alternatives; see Table 5. Namely, we compare our choice to *random*, the basis spanned by k randomly chosen orthogonal vectors in \mathbb{R}^D , *unit_avg_grad*, the basis spanned by the average primary task gradient, and *canonical*, the per-parameter basis.

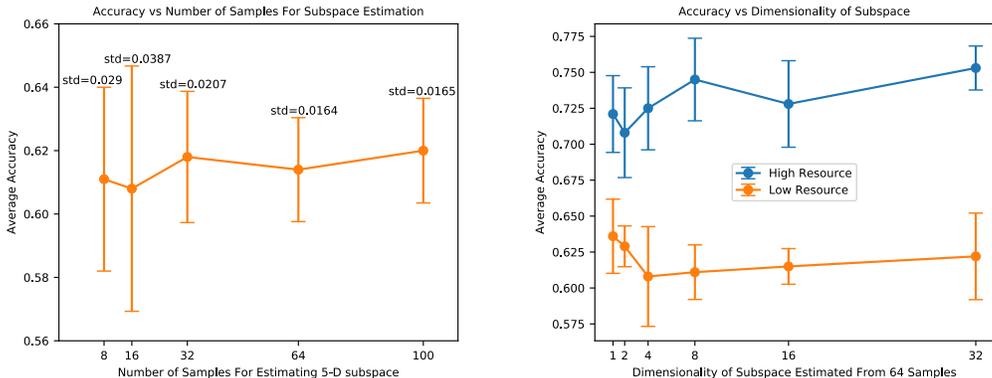


Figure 2: Averaged across 5 random initializations. **Left** We vary the number of samples used to estimate a 5-d subspace up to a maximum of 100 (the total number of training examples in this low-resource setting). **Right.** We compare the effect of the dimensionality of the subspace in the low-resource (50 examples each for Cat, Dog classes) and high-resource (1000 examples each per class).

We also examine the number of samples to estimate the principal directions of the per-example primary task gradient. Larger sample sizes involve more computation but have limited benefit on average accuracy. Large sample sizes however reduce variance, as shown in Figure 2 (left). This is as expected since using more samples gives a higher fidelity estimate of the top- k singular vectors.

Another parameter of our algorithm is the size of our subspace, k . In general, we observe that in low-resource settings, it is better to operate on the auxiliary task gradient in a smaller dimensional subspace. The opposite holds for high-resource settings. This can be seen in Figure 2 (right). Whilst using a larger dimensional subspace captures a richer description of the \mathbf{J}^* , it also creates the risk of over-fitting especially in a limited data setting. This trade-off therefore has to be validated on a per-task basis.

7 CONCLUSIONS

In this work, we propose a new approach to training a model with additional help from an auxiliary task. Our method decomposes the gradients of the auxiliary task according to three directions, with positive, negative and neutral impact on the primary task. This decomposition allows a flexible re-weighting of the auxiliary task components and give rise to a family of training strategies, which encompasses novel and existing approaches. We leverage insights from randomized linear algebra and automatic differentiation to scale the approach to large deep networks. Experiments in multitasking, pretraining and domain transfer over vision and text classification task demonstrate the empirical benefit of our framework.

REFERENCES

- Amittai Axelrod, Xiaodong He, and Jianfeng Gao. Domain adaptation via pseudo in-domain data selection. In *EMNLP*, 2011.
- Atilim Gunes Baydin, Barak A. Pearlmutter, Alexey Andreyevich Radul, and Jeffrey Mark Siskind. Automatic differentiation in machine learning: a survey, 2015.
- Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 132–149, 2018.
- Mathilde Caron, Piotr Bojanowski, Julien Mairal, and Armand Joulin. Unsupervised pre-training of image features on non-curated data. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2959–2968, 2019.
- Rich Caruana. Multitask learning. *Machine Learning*, 1997.
- Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with a-gem. *arXiv:1812.00420*, 2018.
- Zhao Chen, Vijay Badrinarayanan, Chen-Yu Lee, and Andrew Rabinovich. Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In *International Conference on Machine Learning*, 2018.
- Ronan Collobert and Jason Weston. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *International Conference on Machine Learning*, 2008.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Yunshu Du, Wojciech M. Czarnecki, Siddhant M. Jayakumar, Razvan Pascanu, and Balaji Lakshminarayanan. Adapting auxiliary losses using gradient similarity. *CoRR*, abs/1812.02224, 2018.
- Mehrdad Farajtabar, Navid Azizan, Alex Mott, and Ang Li. Orthogonal gradient descent for continual learning. *arXiv preprint arXiv:1910.07104*, 2019.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. Don’t stop pretraining: Adapt language models to domains and tasks. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020. doi: 10.18653/v1/2020.acl-main.740. URL <http://dx.doi.org/10.18653/v1/2020.acl-main.740>.
- Nathan Halko, Per-Gunnar Martinsson, and Joel A Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Matteo Hessel, Hubert Soyer, Lasse Espeholt, Wojciech Czarnecki, Simon Schmitt, and Hado van Hasselt. Multi-task deep reinforcement learning with popart. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 2019.
- Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silvana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghgoo, Robyn Ball, Katie Shpanskaya, et al. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 590–597, 2019.

- Amit Kumar Jaiswal, Prayag Tiwari, Sachin Kumar, Deepak Gupta, Ashish Khanna, and Joel JPC Rodrigues. Identifying pneumonia in chest x-rays: A deep learning approach. *Measurement*, 145: 511–518, 2019.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Simon Kornblith, Jonathon Shlens, and Quoc V. Le. Do better imagenet models transfer better? *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun 2019. doi: 10.1109/cvpr.2019.00277. URL <http://dx.doi.org/10.1109/CVPR.2019.00277>.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7, 2015.
- Xingyu Lin, Harjatin Baweja, George Kantor, and David Held. Adaptive auxiliary task weighting for reinforcement learning. In *Advances in Neural Information Processing Systems*, pp. 4772–4783, 2019.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. In *Advances in Neural Information Processing Systems*, 2017.
- Andrew Maas, Raymond E Daly, Peter T Pham, Dan Huang, Andrew Y Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies*, pp. 142–150, 2011.
- Julian McAuley, Christopher Targett, Qinfeng Shi, and Anton Van Den Hengel. Image-based recommendations on styles and substitutes. In *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval*, pp. 43–52, 2015.
- Ishan Misra, Abhinav Shrivastava, Abhinav Gupta, and Martial Hebert. Cross-stitch networks for multi-task learning. In *Conference on Computer Vision and Pattern Recognition, CVPR*, 2016.
- Robert C Moore and William Lewis. Intelligent selection of language model training data. In *ACL*, 2010.
- Yuji Nakatsukasa. Accuracy of singular vectors obtained by projection-based svd methods. *BIT Numerical Mathematics*, 57(4):1137–1152, 2017.
- Hongseok Namkoong and John C Duchi. Variance-based regularization with convex objectives. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 30*, pp. 2971–2980. Curran Associates, Inc., 2017. URL <http://papers.nips.cc/paper/6890-variance-based-regularization-with-convex-objectives.pdf>.
- Jiquan Ngiam, Daiyi Peng, Vijay Vasudevan, Simon Kornblith, Quoc V. Le, and Ruoming Pang. Domain adaptive transfer learning with specialist models. *CVPR*, 2018.
- Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural networks. In *International conference on machine learning*, pp. 1310–1318, 2013.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- Barak A Pearlmutter. Fast exact multiplication by the hessian. *Neural computation*, 6(1):147–160, 1994.

- Maithra Raghu, Chiyuan Zhang, Jon Kleinberg, and Samy Bengio. Transfusion: Understanding transfer learning for medical imaging. In *Advances in neural information processing systems*, pp. 3347–3357, 2019.
- Vladimir Rokhlin, Arthur Szlam, and Mark Tygert. A randomized algorithm for principal component analysis. *SIAM Journal on Matrix Analysis and Applications*, 31(3):1100–1124, Jan 2010. ISSN 1095-7162. doi: 10.1137/080736417. URL <http://dx.doi.org/10.1137/080736417>.
- Clemens Rosenbaum, Tim Klinger, and Matthew Riemer. Routing networks: Adaptive selection of non-linear functions for multi-task learning. *arXiv preprint arXiv:1711.01239*, 2017.
- Sebastian Ruder. An overview of multi-task learning in deep neural networks. *arXiv:1706.05098*, 2017.
- Ozan Sener and Vladlen Koltun. Multi-task learning as multi-objective optimization. In *Advances in Neural Information Processing Systems*, 2018.
- Ayan Sinha, Zhao Chen, Vijay Badrinarayanan, and Andrew Rabinovich. Gradient adversarial training of neural networks. *arXiv preprint arXiv:1806.08028*, 2018.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. Mass: Masked sequence to sequence pre-training for language generation. *arXiv preprint arXiv:1905.02450*, 2019.
- Mihai Suteu and Yike Guo. Regularizing deep multi-task networks using orthogonal gradients. *arXiv preprint arXiv:1912.06844*, 2019.
- Simon Vandenhende, Stamatios Georgoulis, Marc Proesmans, Dengxin Dai, and Luc Van Gool. Revisiting multi-task learning in the deep learning era, 2020.
- Wei Wang, Ye Tian, Jiquan Ngiam, Yinfei Yang, Isaac Caswell, and Zarana Parekh. Learning a multi-domain curriculum for neural machine translation. In *ACL*, 2020a.
- Xinyi Wang, Hieu Pham, Paul Michel, Antonios Anastasopoulos, Jaime Carbonell, and Graham Neubig. Optimizing data usage via differentiable rewards. In *ACL*, 2020b.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding. In *Advances in neural information processing systems*, pp. 5753–5763, 2019.
- 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.
- Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.
- Zhanpeng Zhang, Ping Luo, Chen Change Loy, and Xiaoou Tang. Facial landmark detection by deep multi-task learning. In *European conference on computer vision*. Springer, 2014.