

Independent Component Alignment for Multi-Task Learning

Dmitry Senushkin Nikolay Patakin Arseny Kuznetsov Anton Konushin
Samsung Research

{d.senushkin, n.patakin, a.konushin}@samsung.com

Abstract

In a multi-task learning (MTL) setting, a single model is trained to tackle a diverse set of tasks jointly. Despite rapid progress in the field, MTL remains challenging due to optimization issues such as conflicting and dominating gradients. In this work, we propose using a condition number of a linear system of gradients as a stability criterion of an MTL optimization. We theoretically demonstrate that a condition number reflects the aforementioned optimization issues. Accordingly, we present Aligned-MTL, a novel MTL optimization approach based on the proposed criterion, that eliminates instability in the training process by aligning the orthogonal components of the linear system of gradients. While many recent MTL approaches guarantee convergence to a minimum, task trade-offs cannot be specified in advance. In contrast, Aligned-MTL provably converges to an optimal point with pre-defined task-specific weights, which provides more control over the optimization result. Through experiments, we show that the proposed approach consistently improves performance on a diverse set of MTL benchmarks, including semantic and instance segmentation, depth estimation, surface normal estimation, and reinforcement learning. The source code is publicly available at <https://github.com/SamsungLabs/MTL>.

1. Introduction

In a multi-task learning (MTL), several tasks are solved jointly by a single model [2, 10]. In such a scenario, information can be shared across tasks, which may improve the generalization and boost the performance for all objectives. Moreover, MTL can be extremely useful when computational resources are constrained, so it is crucial to have a single model capable of solving various tasks [17, 19, 30]. In reinforcement learning [39, 50], MTL setting arises naturally, when a single agent is trained to perform multiple tasks.

Several MTL approaches [15, 24, 28, 29, 31, 35, 42] focus on designing specific network architectures and elaborate strategies of sharing parameters and representations across tasks for a given set of tasks. Yet, such complicated and

powerful models are extremely challenging to train.

Direct optimization of an objective averaged across tasks might experience issues [54] related to conflicting and dominating gradients. Such gradients destabilize the training process and degrade the overall performance. Accordingly, some other MTL approaches address these issues with multi-task gradient descent: either using gradient altering [9, 27, 48, 54] or task balancing [16, 25, 28]. Many recent MTL methods [27, 37, 48] guarantee convergence to a minimum, yet task trade-offs cannot be specified in advance. Unfortunately, the lack of control over relative task importance may cause some tasks to be compromised in favor of others [37].

In this work, we analyze the multi-task optimization challenges from the perspective of stability of a linear system of gradients. Specifically, we propose using a *condition number* of a linear system of gradients as a stability criterion of an MTL optimization. According to our thorough theoretical analysis, there is a strong relation between the condition number and conflicting and dominating gradients issues. We exploit this feature to create Aligned-MTL, a novel gradient manipulation approach, which is the major contribution of this work. Our approach resolves gradient conflicts and eliminates dominating gradients by aligning principal components of a gradient matrix, which makes the training process more stable. In contrast to other existing methods (e.g. [27, 37, 48, 54]), Aligned-MTL has a provable guarantee of convergence to an optimum with pre-defined task weights.

We provide an in-depth theoretical analysis of the proposed method and extensively verify its effectiveness. Aligned-MTL consistently outperforms previous methods on various benchmarks. First, we evaluate the proposed approach on the problem of scene understanding; specifically, we perform joint instance segmentation, semantic segmentation, depth and surface normal estimation on two challenging datasets – Cityscapes [6] and NYUv2 [36]. Second, we apply our method to multi-task reinforcement learning and conduct experiments with the MT10 dataset [55]. Lastly, in order to analyze generalization performance, Aligned-MTL has been applied to two different network architectures, namely PSPNet [48] and MTAN [28], in the scene understanding experiments.

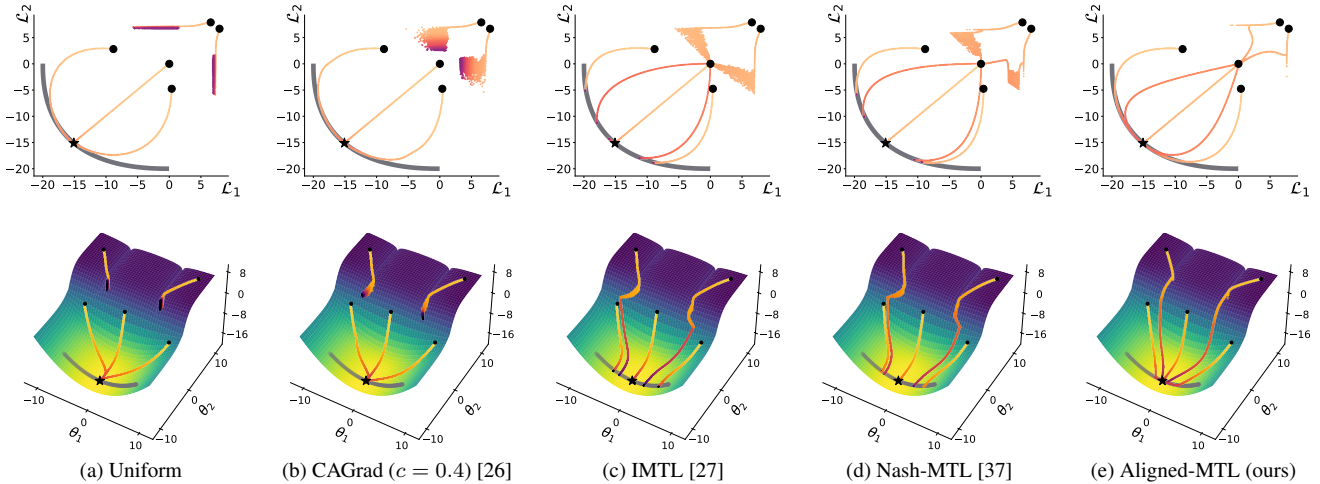


Figure 1. Comparison of MTL approaches on a challenging synthetic two-task benchmark [26, 37]. We visualize optimization trajectories w.r.t. objectives value (\mathcal{L}_1 and \mathcal{L}_2 , top row), and cumulative objective w.r.t. parameters (θ_1 and θ_2 , bottom row). Initialization points are marked with \bullet , the Pareto front (Def. 1) is denoted as --- . Other MTL approaches produce noisy optimization trajectories (Figs. 1a to 1d) inside areas with conflicting and dominating gradients (Fig. 2). In contrast, our approach converges to the global optimum (\star) robustly. Approaches aiming to find a Pareto-stationary solution (such as Fig. 1c and Fig. 1d) terminate once the Pareto front is first reached, as a result, they might provide a suboptimal solution. Differently, Aligned-MTL drifts along the Pareto front and provably converges to the optimum w.r.t. pre-defined tasks weights.

2. Related Work

A multi-task setting [2, 8, 44] is leveraged in computer vision [1, 16, 19, 38, 56], natural language processing [5, 11, 32], speech processing [47], and robotics [23, 52] applications. Prior MTL approaches formulate the total objective as a weighted sum of task-specific objectives, with weights being manually tuned [17, 21, 34]. However, finding optimal weights via grid search is computationally inefficient. Kendall et al. [16] overcame this limitation, assigning task weights according to the homoscedastic uncertainty of each task. Other recent methods, such as GradNorm [3] and DWA [28], optimize weights based on task-specific learning rates or by random weighting [25].

The most similar to Aligned-MTL approaches (e.g. [9, 26, 27, 37, 54]) aim to mitigate effects of conflicting or dominating gradients. Conflicting gradients having opposing directions often induce a negative transfer (e.g. [22]). Among all approaches tackling this problem, the best results are obtained by those based on an explicit gradient modulation [26, 27, 54] where a gradient of a task which conflicts with a gradient of some other task is replaced with a modified, non-conflicting, gradient. Specifically, PCGrad [54] proposes a "gradient surgery" which decorrelates a system of vectors, while CAGrad [26] aims at finding a conflict-averse direction to minimize overall conflicts. GradDrop [4] forces task gradients sign consistency. Other methods also address an issue of dominating gradients. Nash-MTL [37] leverages advances of game theory [37], while IMTL [27] searches

for a gradient direction where all the cosine similarities are equal.

Several recent works [41, 43] investigate a multiple-gradient descent algorithm (MGDA [9, 12, 45]) for MTL: these methods search for a direction that decreases all objectives according to multi-objective Karush–Kuhn–Tucker (KKT) conditions [20]. Sener and Koltun [48] propose extending the classical MGDA [9] so it scales well to high-dimensional problems for a specific use case. However, all the described approaches converge to an arbitrary Pareto-stationary solution, leading to a risk of imbalanced task performance.

3. Multi-Task Learning

Multi-task learning implies optimizing a single model with respect to multiple objectives. The recent works [9, 16, 39, 54] have found that this learning problem is difficult to solve by reducing it to a standard single-task approach. In this section, we introduce a general notation and describe frequent challenges arising in gradient optimization in MTL.

3.1. Notation

In MTL, there are $T > 1$ tasks. Each task is associated with an objective $\mathcal{L}_i(\theta)$ depending on a set of model parameters θ shared between all tasks. The goal of MTL training




Figure 2. Synthetic two-task MTL benchmark [26, 37]. Loss landscapes w.r.t. individual objectives are depicted on the right side. The cumulative loss landscape (on the left side) contains areas with conflicting and dominating gradients.

is to find a parameter θ that minimizes an average loss:

$$\theta^* = \arg \min_{\theta \in \mathbb{R}^m} \left\{ \mathcal{L}_0(\theta) \stackrel{\text{def}}{=} \sum_{i=1}^T \frac{1}{T} \mathcal{L}_i(\theta) \right\} \quad (1)$$

We introduce the following notation: $\mathbf{g}_i = \nabla \mathcal{L}_i(\theta)$ – individual task gradients; $\mathcal{L}_0(\theta)$ – a cumulative objective; $\mathbf{G} = \{\mathbf{g}_1, \dots, \mathbf{g}_T\}$ – a gradient matrix; $w_i \stackrel{\text{def}}{=} \frac{1}{T}$ – predefined task weights. The task weights are supposed to be fixed. We omit task-specific parameters in our notation, since they are independent and not supposed to be balanced.

3.2. Challenges

In practice, directly solving a multi-objective optimization problem via gradient descent may significantly compromise the optimization of individual objectives [54]. Simple averaging of gradients across tasks makes a cumulative gradient biased towards the gradient with the largest magnitude, which might cause overfitting for a subset of tasks. Conflicting gradients with negative cosine distance complicate the training process as well; along with dominating gradients, they increase inter-step direction volatility that decreases overall performance (Fig. 1). To mitigate the undesired effects of conflicting and dominating gradients in MTL, we propose a criterion that is strongly correlated with the presence of such optimization issues. This measure is a *condition number* of a linear system of gradients.

4. Stability

The prevailing challenges in MTL are arguably task dominance and conflicting gradients, accordingly, various criteria for indicating and measuring these issues have been formulated. For instance, a gradient dominance can be measured with a gradient magnitude similarity ([54] Def. 2). Similarly,

gradient conflicts can be estimated as a cosine distance between vectors ([54] Def. 1, [26]). However, each of these metrics describes a specific characteristic of a linear system of gradients, and cannot provide a comprehensive assessment if taken separately. We show that our stability criterion indicates the presence of both MTL challenges (Fig. 4); importantly, it describes a whole linear system and can be trivially measured on any set of gradients. Together with magnitude similarity and cosine distance, this criterion accurately describes the training process.

4.1. Condition Number

Generally, the stability of an algorithm is its sensitivity to an input perturbation, or, in other words, how much the output changes if an input gets perturbed. In numerical analysis, the stability of a linear system is measured by a *condition number* of its matrix. In a multi-task optimization, a cumulative gradient is a linear combination of task gradients: $\mathbf{g} = \mathbf{G}\mathbf{w}$. Thus, the stability of a linear system of gradients can be measured as the condition number of a gradient matrix \mathbf{G} . The value of this stability criterion is equal to the ratio of the maximum and minimum singular values of the corresponding matrix:

$$\kappa(\mathbf{G}) = \frac{\sigma_{max}}{\sigma_{min}}. \quad (2)$$

Remark. A linear system is well-defined if its condition number is equal to one, and ill-posed if it is non-finite. A standard assumption for multi-task optimization is that a gradient system is not ill-posed, *i.e.* task gradients are linearly independent. In this work, we suppose that the linear independence assumption holds unless otherwise stated.

4.2. Condition Number and MTL Challenges

The dependence between the stability criterion and MTL challenges is two-fold. Let us consider a gradient system having a minimal condition number. According to the singular value decomposition theorem, its gradient matrix $\hat{\mathbf{G}}$ with $\kappa(\hat{\mathbf{G}}) = 1$ must be orthogonal with equal singular values:

$$\hat{\mathbf{G}} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top, \quad \text{where } \mathbf{\Sigma} = \sigma\mathbf{I} \quad (3)$$

Moreover, since \mathbf{U} , \mathbf{V} matrices are orthonormal, individual task gradients norms are equal to σ . Thus, minimizing the condition number of the linear system of gradients leads to mitigating dominance and conflicts within this system.

On the other hand, if an initial linear system of gradients is not well-defined, reducing neither gradient conflict nor dominance only does not guarantee minimizing a condition number. The stability criterion reaches its minimum iff both issues are solved jointly and gradients are orthogonal. This restriction eliminates positive task gradients interference (co-directed gradients may produce $\kappa > 1$), but it can guarantee




Figure 3. Geometric interpretation of our approach on a two-task MTL. Here, individual task gradients g_1 and g_2 are directed oppositely (*conflict*) and have different magnitude (*dominance*) (Fig. 3a). Aligned-MTL enforces stability via aligning principal components u_1, u_2 of an initial linear system of gradients. This can be interpreted as re-scaling axes of a coordinate system set by principal components, so that singular values of gradient matrix σ_1 and σ_2 are rescaled to be equal to the minimal singular value (σ_2 , in this case). The aligned gradients \hat{g}_1, \hat{g}_2 are orthogonal (non-conflicting) and of equal magnitude (non-dominant) (Fig. 3b). Finally, the aligned gradients are summed up with pre-defined tasks weights w and $1 - w$, resulting in a cumulative gradient \hat{g}_0 (Fig. 3c).

the absence of negative interaction, which is essential for a stable training. Noisy convergence trajectories w.r.t. objectives values (Fig. 1, top row) indicate instability of the training process.

To demonstrate the relation between MTL challenges and our stability criterion, we conduct a synthetic experiment as proposed in [26, 37]. There are two objectives to be optimized, and the optimization landscape contains areas with conflicting and dominating gradients. We compare our approach against recent approaches that do not handle stability issues, yielding noisy trajectories in problematic areas. By enforcing stability, our method performs well on the synthetic benchmark.

5. Aligned-MTL

We suppose that multi-task gradient optimization should successfully resolve the main MTL challenges: conflicts and dominance in gradient system. Unlike existing approaches [51, 54] that focus on directly resolving the optimization problems, we develop an algorithm that handles issues related to the stability of a linear system of gradients and accordingly addresses both gradient conflicts and dominance.

Specifically, we aim to find a cumulative gradient \hat{g}_0 , so that $\|\mathbf{g}_0 - \hat{g}_0\|_2^2$ is minimal, while a linear system of gradients is stable ($\kappa(\hat{\mathbf{G}}) = 1$). This constraint is defined up to an arbitrary positive scaling coefficient. Here, we assume $\sigma = 1$ for simplicity. By applying a triangle inequality to the initial problem, we derive $\|\mathbf{g}_0 - \hat{g}_0\|_2^2 \leq \|\mathbf{G} - \hat{\mathbf{G}}\|_F^2 \|\mathbf{w}\|_2^2$. Thereby, we consider the following optimization task:

$$\min_{\hat{\mathbf{G}}} \|\mathbf{G} - \hat{\mathbf{G}}\|_F^2 \quad \text{s.t.} \quad \hat{\mathbf{G}}^\top \hat{\mathbf{G}} = \mathbf{I} \quad (4)$$

The stability criterion, a condition number, defines a linear system up to an arbitrary positive scale. To alleviate this ambiguity, we choose the largest scale that guarantees convergence to the optimum of an original problem (Eq. (1)): this is a minimal singular value of an initial gradient matrix

$\sigma = \sigma_{\min}(\mathbf{G}) > 0$. The final linear system of gradients defined by $\hat{\mathbf{G}}$ satisfies the optimality condition in terms of a condition number.

5.1. Gradient Matrix Alignment

The problem Eq. (4) can be treated as a special kind of Procrustes problem [46]. Fortunately, there exists a closed-form solution of this task. To obtain such a solution, we perform a singular value decomposition (SVD) and rescale singular values corresponding to principal components so that they are equal to the smallest singular value.

Technically, the matrix alignment can be performed in the *parameter* space or in the *task* space; being equivalent, these options have different computational costs. This duality is caused by SVD providing two different eigen decompositions of Gram matrices $\mathbf{G}^\top \mathbf{G}$ and $\mathbf{G}\mathbf{G}^\top$:

$$\hat{\mathbf{G}} = \sigma \mathbf{U}\mathbf{V}^\top = \sigma \underbrace{\mathbf{U}\boldsymbol{\Sigma}^{-1}\mathbf{U}^\top}_{\text{Parameter space}} \mathbf{G} = \sigma \mathbf{G} \underbrace{\mathbf{V}\boldsymbol{\Sigma}^{-1}\mathbf{V}^\top}_{\text{Task space}} \quad (5)$$

We perform the gradient matrix alignment at each optimization step. Since the number of tasks T is relatively small compared to the number of parameters, we operate in a task space. This makes a gradient matrix alignment more computationally efficient as we need to perform an eigen decomposition of a small $T \times T$ matrix. Fig. 3 provides a geometric interpretation of our approach, while pseudo-code is given in Alg. 1.

Remark. If an initial matrix \mathbf{G} is singular (gradients are linear dependent), then the smallest singular value is zero. Fortunately, the singular value decomposition provides a unique solution even in this case; yet, we need to choose the smallest singular value greater than zero as a global scale.

5.2. Aligned-MTL: Upper Bound Approximation

The major limitation of our approach is the need to run multiple backward passes through the shared part of the model to calculate the gradient matrix. The backward passes

Algorithm 1 Gradient matrix alignment

Require: $G \in \mathbb{R}^{|\theta| \times T}$ – gradient matrix,
 $w \in \mathbb{R}^T$ – task importance
/* Compute task space Gram matrix */
 $M \leftarrow G^\top G$
/* Compute eigenvalues and eigenvectors of M */
 $(\lambda, V) \leftarrow \text{eigh}(M)$
 $\Sigma^{-1} \leftarrow \text{diag}\left(\sqrt{\frac{1}{\lambda_1}}, \dots, \sqrt{\frac{1}{\lambda_R}}\right)$
/* Compute balance transformation */
 $B \leftarrow \sqrt{\lambda_R} V \Sigma^{-1} V^\top$
 $\alpha \leftarrow Bw$
return $G\alpha$

are computationally demanding, and the training time depends linearly on the number of tasks: if it is large, our approach may be non-applicable in practice.

This limitation can be mitigated for encoder-decoder networks, where each task prediction is computed using the same shared representation. We can employ the chain rule trick [48] to upper-bound an original objective (Eq. (4)):

$$\|G - \hat{G}\|_F^2 \leq \left\| \frac{\partial H}{\partial \theta} \right\|_F^2 \|Z - \hat{Z}\|_F^2 \quad (6)$$

Here, H stands for a hidden shared representation, and Z and \hat{Z} are gradients of objective w.r.t. a shared representation of the initial and aligned linear system of gradients, respectively. Thus, the gradient alignment can be performed for a shared representation:

$$\min_Z \|Z - \hat{Z}\|_F^2 \quad \text{s.t.} \quad \hat{Z}^\top \hat{Z} = I \quad (7)$$

Aligning gradients of shared representation does not require additional backward passes, since matrix Z is computed during a conventional backward pass. We refer to such an approximation of Aligned-MTL as to Aligned-MTL-UB. With $O(1)$ time complexity w.r.t. the number of tasks, this approximation tends to be significantly more efficient than the original Aligned-MTL having $O(T)$ time complexity.

5.3. Convergence Analysis

In this section, we formulate a theorem regarding the convergence of our approach. Similar to single-task optimization converging to a stationary point, our MTL approach converges to a *Pareto-stationary* solution.

Definition 1 A solution $\theta^* \in \Theta$ is called **Pareto-stationary** iff there exists a convex combination of the gradient-vectors that is equal to zero. All possible Pareto-stationary solutions form a **Pareto set** (or **Pareto front**).

The overall model performance may vary significantly within points of the Pareto front. Recent MTL approaches [27, 37] that provably converge to an arbitrary

Pareto-stationary solution, tend to overfit to a subset of tasks. In contrast, our approach converges to a Pareto-stationary point with pre-defined tasks weights, thus providing more control over an optimization result Eq. (1).

Theorem 1 Assume $\mathcal{L}_0(\theta), \dots, \mathcal{L}_T(\theta)$ are lower-bounded continuously differentiable functions with Lipschitz continuous gradients with $\Lambda > 0$. A gradient descent with aligned gradient and step size $\alpha \leq \frac{1}{\Lambda}$ converges linearly to a Pareto stationary point where $\nabla \mathcal{L}_0(\theta) = 0$.

A similar theorem is valid for aligning gradients in the shared representation space (Aligned-MTL upper-bound approximation is described in Sec. 5.2). Mathematical proofs of both versions of this theorem versions are provided in supplementary materials.

6. Experiments

We empirically demonstrate the effectiveness of the proposed approach on various multi-task learning benchmarks, including scene understanding, multi-target regression, and reinforcement learning.

Competitors. We consider the following MTL approaches: (1) **Linear Scalarization (LS, Uniform baseline)**: optimizing a uniformly weighted sum of individual task objectives, i.e. $\frac{1}{T} \sum_t \mathcal{L}_t$; (2) **Dymanic Weight Average (DWA)** [28]: adjusting task weights based on the rates of loss changes over time; (3) **Uncertainty** [16] weighting; (4) **MGDA** [9]: a multi-objective optimization with KKT [20] conditions; (5) **MGDA-UB** [48]: optimizing an upper bound for the MGDA optimization objective; (6) **GradNorm** [3]: normalizing the gradients to balance the learning of multiple tasks; (7) **GradDrop** [4]: forcing the sign consistency between task gradients; (8) **PCGrad** [54]: performing gradient projection to avoid the negative interactions between tasks gradients; (9) **GradVac** [51]: leveraging task relatedness to set gradient similarity objectives and adaptively align task gradients, (10) **CAGrad** [26]: finding a conflict-averse gradients; (11) **IMTL** [27]: aligning projections to task gradients; (12) **Nash-MTL** [37]: utilizing a bargaining games for gradient computation, and (13) **Random loss weighting (RLW)** [25] with normal distribution. The proposed approach and the baseline methods are implemented using the PyTorch framework [40]. The technical details on the training schedules and a complete listing of hyperparameters are provided in supplementary materials.

Evaluation metrics. Besides task specific metrics we follow Maninis et al. [33] and report a model performance drop relative to a single task baseline averaged over tasks: $\Delta m_{task} = \frac{1}{T} \sum_{t=1}^T \sum_{k=1}^{n_t} (-1)^{\sigma_{tk}} (M_{m,tk} - M_{b,tk}) / M_{b,tk}$ – or over metrics: $\Delta m_{metric} = \frac{1}{T} \sum_{t=1}^T (-1)^{\sigma_t} (M_{m,t} - M_{b,t}) / M_{b,t}$. Here, $M_{m,tk}$ denotes the performance of a model m on a




Figure 4. Empirical evaluation of a stability criterion. We plot a condition number (Eq. (2)), gradient magnitude similarity [54], and minimal cosine distance during training on the CITYSCAPES three-task benchmark. This benchmark suffers from high dominance since instance segmentation loss is of much larger scale than the others. The most intuitive way to define the dominance is the maximum ratio of task gradients magnitudes. The condition number coincides with this definition in a nearly orthogonal case, as in this benchmark Fig. 4c. However, gradient magnitude similarity measure Fig. 4b proposed in [54] does not reveal much correlation with a condition number (and, accordingly, with a maximal gradients magnitude ratio) Fig. 4a, so we assume it does not represent dominance issues comprehensively. From the empirical point of view Table 1, the value of a target metric is more correlated with the condition number, than with the gradient magnitude similarity.

task t , measured with a metric k . Similarly, $M_{b,tk}$ is a performance of a single-task t baseline; n_t denotes number of metrics per task t . $\sigma_{tk} = 1$ if higher values of metric is better, and $\sigma_{tk} = 0$ otherwise. We mostly rely on the task-weighted measure since the metric-weighted criterion tends to be biased to a task with high number of metrics.

6.1. Synthetic Example

To illustrate the proposed approach, we consider a synthetic MTL task (Fig. 2) introduced in [26] (a formal definition is provided in the supplementary material). As shown in Fig. 1, we perform optimization from five initial points tagged with \bullet . IMTL [27], and Nash-MTL [37] aims at finding Pareto-stationary solution (Def. 1). As a result, they terminate optimization once they reach a solution in the Pareto front. Accordingly, the final result strongly depends on an initialization point, and the optimization may not converge to the global optimum \star in some cases (Fig. 1c and Fig. 1d). Meanwhile, Aligned-MTL provides a less noisy and more stable trajectory, and provably converges to an optimum.

6.2. Scene Understanding

The evaluation is performed on NYUv2 [36] and CITYSCAPES [6, 7] datasets. We leverage two network architectures: Multi-Task Attention Network (MTAN) [28] and Pyramid Scene Parsing Network (PSPNet) [48, 57] on scene understanding benchmarks. MTAN applies a multi-task specific attention mechanism built upon MTL SegNet [16]. PSPNet features a dilated ResNet [14] backbone and multiple decoders with pyramid parsing modules [57]. Both networks were previously used in MTL benchmarks [48].

Table 1. **Scene understanding** (CITYSCAPES: three tasks). We report PSPNet [48, 57] model performance averaged over 3 random seeds. The best scores are provided in gray.

Method	Segmentation \uparrow mIoU [%]	Instance \downarrow L1 [px]	Disparity \downarrow MSE	$\Delta m\%$ \downarrow
Single task baselines	66.73	10.55	0.33	–
Baseline:Uniform	52.98	10.89	0.39	14.30
RLW [25]	51.26	10.25	0.41	15.58
DWA [28]	53.15	10.22	0.40	13.20
Uncertainty [16]	60.12	9.87	0.33	1.53
MGDA [9]	66.72	17.02	0.33	20.62
MGDA-UB [48]	66.37	18.63	0.32	25.05
GradNorm [3]	57.24	10.29	0.35	6.55
GradDrop [4]	52.98	10.09	0.40	12.50
PCGrad [54]	54.06	9.91	0.38	10.00
GradVac [51]	54.07	10.39	0.40	12.99
CAGrad [26]	64.33	10.15	0.34	1.46
IMTL [27]	65.13	11.58	0.32	3.10
Nash-MTL [37]	64.84	11.90	0.37	9.38
Aligned-MTL (ours)	67.06	10.63	0.33	−0.02
Aligned-MTL-UB (ours)	66.07	10.54	0.32	−0.35

NYUv2. Following Liu *et al.* [26, 28, 37], we evaluate the performance of our approach on the NYUv2 [36] dataset by jointly solving semantic segmentation, depth estimation, and surface normal estimation tasks. We use both MTAN [28] and PSPNet [48] model architectures.

For MTAN, we strictly follow the training procedure described in [26, 37]: training at 384×288 resolution for 200 epochs with Adam [18] optimizer and 10^{-4} initial learning rate, halved after 100 epochs. The evaluation results are presented in Table 2. We report metric values averaged across three random initializations as in previous works. We calculate both metric-weighted measure to compare with

Table 2. **Scene understanding** (NYUv2, three tasks). We report MTAN [28] model performance averaged over 3 random seeds. The best scores are provided in **gray**.

Method	Segmentation \uparrow		Depth \downarrow		Surface normals \downarrow					$\Delta m\%$ \downarrow	$\Delta m\%$ \downarrow	
	mIoU	Pix Acc	Abs.	Rel.	Angle	Dist. \downarrow	Within t° \uparrow	11.25	22.5	30	Metric-weighted	Task-weighted
Single task baselines	38.30	63.76	0.68	0.28	25.01	19.21	30.14	57.20	69.15		-	-
Baseline: Uniform	39.29	65.33	0.55	0.23	28.15	23.96	22.09	47.50	61.08		5.46	-1.07
RLW [25]	37.17	63.77	0.58	0.24	28.27	24.18	22.26	47.05	60.62		7.67	2.00
DWA [28]	39.11	65.31	0.55	0.23	27.61	23.18	24.17	50.18	62.39		3.49	-2.06
Uncertainty [16]	36.87	63.17	0.54	0.23	27.04	22.61	23.54	49.05	63.65		4.01	-0.97
MGDA [48]	30.47	59.90	0.61	0.26	24.88	19.45	29.18	56.88	69.36		1.47	1.79
GradNorm [3]	20.09	52.06	0.72	0.28	24.83	18.86	30.81	57.94	69.73		7.22	11.51
GradDrop [4]	39.39	65.12	0.55	0.23	27.48	22.96	23.38	49.44	62.87		3.61	-2.03
PCGrad [54]	38.06	64.64	0.56	0.23	27.41	22.80	23.86	49.83	63.14		3.83	-1.33
GradVac [51]	37.53	64.35	0.56	0.24	27.66	23.38	22.83	48.66	62.21		5.44	0.01
CAGrad [26]	39.79	65.49	0.55	0.23	26.31	21.58	25.61	52.36	65.58		0.29	-4.18
IMTL [27]	39.35	65.60	0.54	0.23	26.02	21.19	26.20	53.13	66.24		-0.59	-4.76
Nash-MTL [37]	40.13	65.93	0.53	0.22	25.26	20.08	28.40	55.47	68.15		-4.04	-7.56
Aligned-MTL (ours)	40.82	66.33	0.53	0.22	25.19	19.71	28.88	56.23	68.54		-4.93	-8.40
Aligned-MTL-UB (ours)	43.11	67.22	0.55	0.22	25.67	20.57	27.58	54.37	67.12		-3.48	-7.83

Table 3. **Scene understanding** (NYUv2, three tasks). We report PSPNet [48, 57] model performance averaged over 3 random seeds. The best scores are provided in **gray**.

Method	Segmentation \uparrow		Depth \downarrow		Surface normals \downarrow					$\Delta m\%$ \downarrow	$\Delta m\%$ \downarrow	
	mIoU	Pix Acc	Abs.	Rel.	Angle	Dist. \downarrow	Within t° \uparrow	11.25	22.5	30	Metric-weighted	Task-weighted
Single task baselines	49.37	72.03	0.52	0.24	22.97	16.94	0.34	0.62	0.73		-	-
Baseline: Uniform	45.21	69.70	0.49	0.21	26.10	21.08	0.26	0.52	0.66		8.97	4.72
RLW [25]	46.19	69.71	0.46	0.19	26.09	21.09	0.27	0.53	0.66		6.67	1.73
DWA [28]	45.83	69.65	0.50	0.22	26.10	21.27	0.26	0.52	0.66		9.64	5.61
MGDA [48]	40.96	65.80	0.54	0.22	23.36	17.45	0.33	0.61	0.72		3.54	4.24
MGDA-UB [48]	41.15	65.10	0.53	0.22	23.42	17.60	0.32	0.60	0.72		4.02	4.40
GradNorm [3]	45.63	69.64	0.48	0.20	25.46	20.06	0.28	0.55	0.67		5.88	2.18
GradDrop [4]	45.69	70.13	0.49	0.20	26.16	21.21	0.26	0.52	0.65		8.60	3.92
PCGrad [54]	46.37	69.69	0.48	0.20	26.00	21.05	0.26	0.53	0.66		7.78	3.17
GradVac [51]	46.65	69.97	0.49	0.21	25.95	20.88	0.27	0.53	0.66		7.89	3.75
CAGrad [26]	45.46	69.35	0.47	0.20	24.28	18.73	0.30	0.58	0.70		2.66	0.13
IMTL [27]	44.02	68.56	0.47	0.19	23.69	18.03	0.32	0.59	0.72		0.76	-1.02
Nash-MTL [37]	47.25	70.38	0.46	0.20	23.95	18.83	0.31	0.59	0.71		1.13	-1.48
Aligned-MTL (ours)	46.70	69.97	0.46	0.19	24.19	18.77	0.30	0.58	0.71		1.44	-1.55
Aligned-MTL-UB (ours)	46.47	69.92	0.48	0.20	24.37	18.88	0.30	0.58	0.70		2.70	0.07

previous works alongside a task-weighted Δm modification. We claim the latter measure to be more important, as it is not biased towards surface normal estimation, thereby assessing overall performance more fairly. Accordingly, it exposes inconsistent task performance of GradNorm [3] and MGDA [48], which are biased towards surface normal estimation task and perform poorly on semantic segmentation. Although the MTAN model is not encoder-decoder architecture, our Aligned-MTL-UB approach outperforms all previous MTL optimization methods according to task-weighted Δm . Our original Aligned-MTL approach improves model performance even further in terms of both metrics.

We report results of PSPNet (Table 3), trained on NYUv2 [36] following the same experimental setup. PSPNet architecture establishes much stronger baselines for all three tasks than vanilla SegNet. As a result, most of MTL approaches fail to outperform single-task models. According to the task-weighted metric, only two previous approaches

provide solutions better than single-task baselines, while our Aligned-MTL approach demonstrates the best results.

CITYSCAPES: two-task. We follow Liu *et al.* [26] experimental setup for Cityscapes [6], which implies jointly addressing semantic segmentation and depth estimation with a single MTAN [28] model. According to it, the original 19 semantic segmentation categories are classified into 7 categories. Our Aligned-MTL approach demonstrates the best results according to semantic segmentation and overall Δm metric. Our upper bound approximation of our Aligned-MTL again achieves a competitive performance, although MTAN does not satisfy architectural requirements.

CITYSCAPES: three-task. We adopt a more challenging experimental setup [16, 48], and address MTL with disparity estimation and instance and semantic segmentation tasks. The instance segmentation is reformulated as a centroid regression [16], so that the instance objective has a much larger scale than others. In this benchmark, we utilize the training

setup proposed by Sener and Koltun [48]: 100 epochs, Adam optimizer with learning rate 10^{-4} . Input images are rescaled to 256×512 , and a full set of labels is used for semantic segmentation. While many recent approaches experience a considerable performance drop (Table 1), our method performs robustly even in this challenging scenario.

6.3. Multi-task Reinforcement Learning

Following [26,37,54], we consider an MTL reinforcement learning benchmark MT10 in a MetaWorld [55] environment. In this benchmark, a robot is being trained to perform actions, *e.g.* pressing a button and opening a window. Each action is treated as a task, and the primary goal is to successfully perform a total of 10 diverse manipulation tasks. In this experiment, we compare against the optimization-based baseline Soft Actor-Critic (SAC) [13] trained with various gradient altering methods [26, 37, 54]. We also consider MTL-RL [49]-based approaches: specifically, MTL SAC with a shared model, Multi-task SAC with task encoder (MTL SAC + TE) [55], Multi-headed SAC (MH SAC) with task-specific heads [55], Soft Modularization (SM) [53] and CARE [49]. The Aligned-MTL method has higher success rates, superseding competitors by a notable margin.

6.4. Empirical Analysis of Stability Criterion

In this section, we analyze gradient magnitude similarity, cosine distance, and condition number empirically. We use CITYSCAPES three-task benchmark for this purpose, which suffers from the dominating gradients. According to the gradient magnitude similarity measure, PCGrad [54], Uniform, and CAGrad [26] tend to suffer from gradient dominance. For PCGrad and Uniform baseline, imbalanced convergence rates for different tasks result in a suboptimal solution (Table 1). Differently, a well-performing CAGrad is misleadingly marked as problematic by gradient magnitude similarity. In contrast, the stability criterion – *condition number* – reveals domination issues for PCGrad and Uniform baselines and indicates a sufficient balance of different tasks for CAGrad ($\kappa \approx 5$). Thus, the condition number exposes the training issues more evidently (Fig. 4a). The experimental evaluation shows that with $\kappa \leq 10$, model tends to converge to an optimum with better overall performance.

7. Discussion

The main limitation of Aligned-MTL is its computational optimization cost which scales linearly with the number of tasks. The upper-bound approximation of the Aligned-MTL method can be efficiently applied for encoder-decoder architectures using the same Jacobian over the shared representation. This approximation reduces instability, yet, it does not eliminate it since the Jacobian cannot be aligned. For non-encoder-decoder networks, upper-bound approximation

Table 4. **Scene understanding** (CITYSCAPES: two tasks). MTAN [28] model performance is reported as average over 3 random seeds. The best scores are provided in **gray**.

	Segmentation		Depth		$\Delta_m\%$ ↓
	mIoU [%] ↑	Pix. Acc ↑	Abs Err ↓	Rel Err ↓	
Single task baselines	74.01	93.16	0.0125	27.77	–
Baseline: Uniform	75.18	93.49	0.0155	46.77	22.60
RLW [25]	74.57	93.41	0.0158	47.79	24.37
DWA [28]	75.24	93.52	0.0160	44.37	21.43
Uncertainty [16]	72.02	92.85	0.0140	30.13	5.88
MGDA [48]	68.84	91.54	0.0309	33.50	44.14
GradNorm [3]	73.72	93.04	0.0124	34.11	5.63
GradDrop [4]	75.27	93.53	0.0157	47.54	23.67
PCGrad [54]	75.13	93.48	0.0154	42.07	18.21
CAGrad [26]	75.16	93.48	0.0141	37.60	11.58
IMTL [27]	75.33	93.49	0.0135	38.41	11.04
Nash-MTL [37]	75.41	93.66	0.0129	35.02	6.72
Aligned-MTL (ours)	75.77	93.69	0.0133	32.66	5.27
A-MTL-UB* (ours)	74.89	93.46	0.0131	33.92	6.37

Table 5. **Reinforcement learning** (MT10). Average success rate on validation over 10 seeds.

	Success ± SEM
STL SAC	0.90 ± 0.032
MTL SAC	0.49 ± 0.073
MTL SAC + TE	0.54 ± 0.047
MH SAC	0.61 ± 0.036
SM	0.73 ± 0.043
CARE	0.84 ± 0.051
PCGrad	0.72 ± 0.022
CAGrad	0.83 ± 0.045
Nash-MTL	0.91 ± 0.031
Ours, Aligned-MTL	0.97 ± 0.045

has no theoretical guarantees but still can be leveraged as a heuristic and even provide a decent performance.

8. Conclusion

In this work, we introduced a stability criterion for multi-task learning, and proposed a novel gradient manipulation approach that optimizes this criterion. Our Aligned-MTL approach stabilize the training procedure by aligning the principal components of the gradient matrix. In contrast to many previous methods, this approach guarantees convergence to the local optimum with pre-defined task weights, providing a better control over the optimization results. Additionally, we presented a computationally efficient approximation of Aligned-MTL. Through extensive evaluation, we proved our approach consistently outperforms previous MTL optimization methods on various benchmarks including scene understanding and multi-task reinforcement learning.

Acknowledgements. We sincerely thank Anna Vorontsova, Iaroslav Melekhov, Mikhail Romanov, Juho Kannala and Arno Solin for their helpful comments, discussions and proposed improvements regarding this paper.

References

- [1] Hakan Bilen and Andrea Vedaldi. Integrated perception with recurrent multi-task neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, volume 29, pages 235–243. Curran Associates, Inc., 2016.
- [2] Richard Caruana. Multitask learning: A knowledge-based source of inductive bias. In *Proceedings of the Tenth International Conference on Machine Learning (ICML)*, pages 41–48. Morgan Kaufmann, 1993.
- [3] Zhao Chen, Vijay Badrinarayanan, Chen-Yu Lee, and Andrew Rabinovich. GradNorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, volume 80 of *Proceedings of Machine Learning Research*, pages 794–803. PMLR, 2018.
- [4] 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. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pages 2039–2050. Curran Associates, Inc., 2020.
- [5] Ronan Collobert and Jason Weston. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *Proceedings of the 25th International Conference on Machine Learning (ICML)*, pages 160–167. ACM, 2008.
- [6] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [7] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Scharwächter, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset. In *CVPR Workshop on The Future of Datasets in Vision*, 2015.
- [8] Michael Crawshaw. Multi-task learning with deep neural networks: A survey. *arXiv preprint arXiv:2009.09796*, 2020.
- [9] Jean-Antoine Désidéri. Multiple-gradient descent algorithm for multiobjective optimization. In *European Congress on Computational Methods in Applied Sciences and Engineering (ECCOMAS)*, 2012.
- [10] Carl Doersch and Andrew Zisserman. Multi-task self-supervised visual learning. In *ICCV*, pages 2051–2060, 2017.
- [11] Daxiang Dong, Hua Wu, Wei He, Dianhai Yu, and Haifeng Wang. Multi-task learning for multiple language translation. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, pages 1723–1732. Association for Computational Linguistics, 2015.
- [12] Jörg Fliege and Benar Fux Svaiter. Steepest descent methods for multicriteria optimization. *Mathematical Methods of Operations Research*, 51:479–494, 2000.
- [13] Tuomas Haarnoja, Aurick Zhou, P. Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *ICML*, 2018.
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.
- [15] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for NLP. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, volume 97 of *Proceedings of Machine Learning Research*, pages 2790–2799. PMLR, 2019.
- [16] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *CVPR*, pages 7482–7491, 2018.
- [17] Alex Kendall, Matthew Grimes, and Roberto Cipolla. PoseNet: A convolutional network for real-time 6-dof camera relocalization. In *ICCV*, pages 2938–2946, 2015.
- [18] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *ICLR*, 2015.
- [19] Iasonas Kokkinos. UberNet: Training a universal convolutional neural network for low-, mid-, and high-level vision using diverse datasets and limited memory. In *CVPR*, pages 6129–6138, 2017.
- [20] Harold W. Kuhn and Albert W. Tucker. Nonlinear programming. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*. University of California Press, 1951.
- [21] Zakaria Laskar, Iaroslav Melekhov, Surya Kalia, and Juho Kannala. Camera relocalization by computing pairwise relative poses using convolutional neural network. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV) Workshops*, pages 920–929, 2017.
- [22] Hae Beom Lee, Eunho Yang, and Sung Ju Hwang. Deep asymmetric multi-task feature learning. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, volume 80 of *Proceedings of Machine Learning Research*, pages 2956–2964. PMLR, 2018.
- [23] Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. End-to-end training of deep visuomotor policies. *The Journal of Machine Learning Research*, 1:1334–1373, 2016.
- [24] Wei-Hong Li, Xialei Liu, and Hakan Bilen. Universal representations: A unified look at multiple task and domain learning. *arXiv preprint arXiv:2204.02744*, 2022.
- [25] Baijiong Lin, Feiyang Ye, Yu Zhang, and Ivor W. Tsang. Reasonable effectiveness of random weighting: A litmus test for multi-task learning. *arXiv preprint arXiv:2111.10603*, 2022.
- [26] Bo Liu, Xingchao Liu, Xiaojie Jin, Peter Stone, and Qiang Liu. Conflict-averse gradient descent for multi-task learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 34, pages 18878–18890. Curran Associates, Inc., 2021.
- [27] Liyang Liu, Yi Li, Zhanghui Kuang, Jing-Hao Xue, Yimin Chen, Wenming Yang, Qingmin Liao, and Wayne Zhang. Towards impartial multi-task learning. In *ICLR*, 2021.

- [28] Shikun Liu, Edward Johns, and Andrew J. Davison. End-to-end multi-task learning with attention. In *CVPR*, pages 1871–1880, 2019.
- [29] Shikun Liu, Edward Johns, and Andrew J. Davison. End-to-end multi-task learning with attention. In *CVPR*, pages 1871–1880, 2019.
- [30] Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. Multi-task deep neural networks for natural language understanding. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, volume 57. Association for Computational Linguistics, 2019.
- [31] Jiasen Lu, Vedanuj Goswami, Marcus Rohrbach, Devi Parikh, and Stefan Lee. 12-in-1: Multi-task vision and language representation learning. In *CVPR*, pages 10434–10443, 2020.
- [32] Minh-Thang Luong, Quoc Le, Ilya Sutskever, Oriol Vinyals, and Lukasz Kaiser. Multi-task sequence to sequence learning. *ICLR*, 2015.
- [33] Kevis-Kokitsi Maninis, Ilija Radosavovic, and Iasonas Kokkinos. Attentive single-tasking of multiple tasks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [34] Iaroslav Melekhov, Juha Ylioinas, Juho Kannala, and Esa Rahtu. Image-based localization using hourglass networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV) Workshops*, pages 870–877, 2017.
- [35] Ishan Misra, Abhinav Shrivastava, Abhinav Gupta, and Martial Hebert. Cross-stitch networks for multi-task learning. In *CVPR*, pages 3994–4003, 2016.
- [36] Pushmeet Kohli Nathan Silberman, Derek Hoiem and Rob Fergus. Indoor segmentation and support inference from rgbd images. In *ECCV*, 2012.
- [37] Aviv Navon, Aviv Shamsian, Idan Achituve, Haggai Maron, Kenji Kawaguchi, Gal Chechik, and Ethan Fetaya. Multi-task learning as a bargaining game. In *Proceedings of the 39th International Conference on Machine Learning (ICML)*, volume 162 of *Proceedings of Machine Learning Research*, pages 16428–16446. PMLR, 2022.
- [38] Vladimir Nekrasov, Thanuja Dharmasiri, Andrew Spek, Tom Drummond, Chunhua Shen, and Ian Reid. Real-time joint semantic segmentation and depth estimation using asymmetric annotations. In *International Conference on Robotics and Automation (ICRA)*, pages 7101–7107. IEEE, 2019.
- [39] Emilio Parisotto, Lei Jimmy Ba, and Ruslan Salakhutdinov. Actor-mimic: Deep multitask and transfer reinforcement learning. In *ICLR*, 2016.
- [40] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 32, pages 8026–8037. Curran Associates, Inc., 2019.
- [41] Sebastian Peitz and Michael Dellnitz. *Gradient-Based Multiobjective Optimization with Uncertainties*, pages 159–182. Springer International Publishing, 2018.
- [42] Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. AdapterFusion: Non-destructive task composition for transfer learning. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics*, pages 487–503. Association for Computational Linguistics, 2021.
- [43] Fabrice Poirion, Quentin Mercier, and Jean-Antoine Désidéri. Descent algorithm for nonsmooth stochastic multiobjective optimization. *Computational Optimization and Applications*, (2):317–331, 2017.
- [44] Sebastian Ruder. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*, 2017.
- [45] Stefan Schäffler, Richard R. Schultz, and Konstanze Weinzierl. Stochastic Method for the Solution of Unconstrained Vector Optimization Problems. *Journal of Optimization Theory and Applications*, 114:209–222, 2002.
- [46] Peter Schönemann. A generalized solution of the orthogonal procrustes problem. *Psychometrika*, 31(1):1–10, 1966.
- [47] Michael L. Seltzer and Jasha Droppo. Multi-task learning in deep neural networks for improved phoneme recognition. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6965–6969. IEEE, 2013.
- [48] Ozan Sener and Vladlen Koltun. Multi-task learning as multi-objective optimization. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 31, pages 527–538. Curran Associates, Inc., 2018.
- [49] Shagun Sodhani, Amy Zhang, and Joelle Pineau. Multi-task reinforcement learning with context-based representations. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 9767–9779. PMLR, 18–24 Jul 2021.
- [50] Yee Teh, Victor Bapst, Wojciech M. Czarnecki, John Quan, James Kirkpatrick, Raia Hadsell, Nicolas Heess, and Razvan Pascanu. Distral: Robust multitask reinforcement learning. In *Advances in Neural Information Processing Systems (NIPS)*, volume 30, page 4499–4509. Curran Associates, Inc., 2017.
- [51] Zirui Wang, Yulia Tsvetkov, Orhan Firat, and Yuan Cao. Gradient vaccine: Investigating and improving multi-task optimization in massively multilingual models. In *ICLR*, 2021.
- [52] Markus Wulfmeier, Abbas Abdolmaleki, Roland Hafner, Jost Tobias Springenberg, Michael Neunert, Noah Siegel, Tim Hertweck, Thomas Lampe, Nicolas Heess, and Martin Riedmiller. Compositional transfer in hierarchical reinforcement learning. In *Proceedings of Robotics: Science and Systems*, 2020.
- [53] Ruihan Yang, Huazhe Xu, Yi Wu, and Xiaolong Wang. Multi-task reinforcement learning with soft modularization. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS’20*, Red Hook, NY, USA, 2020. Curran Associates Inc.
- [54] Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. Gradient surgery for multi-task learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pages 5824–5836. Curran Associates, Inc., 2020.
- [55] Tianhe Yu, Deirdre Quillen, Zhanpeng He, Ryan Julian, Karol Hausman, Chelsea Finn, and Sergey Levine. Meta-world: A

benchmark and evaluation for multi-task and meta reinforcement learning. In *Conference on Robot Learning (CoRL)*, 2019.

- [56] Amir R. Zamir, Alexander Sax, William Shen, Leonidas J. Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling Task Transfer Learning. In *CVPR*, pages 3712–3722, 2018.
- [57] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *CVPR*, pages 6230–6239, 2017.