

Sample-Level Weighting for Multi-Task Learning with Auxiliary Tasks

Grégoire, Emilie; Chaudhary, Muhammad Hafeez; Verboven, Sam

Publication date:
2023

Document Version:
Submitted manuscript

[Link to publication](#)

Citation for published version (APA):
Grégoire, E., Chaudhary, M. H., & Verboven, S. (2023). *Sample-Level Weighting for Multi-Task Learning with Auxiliary Tasks*. ArXiv. <https://arxiv.org/pdf/2306.04519.pdf>

Copyright

No part of this publication may be reproduced or transmitted in any form, without the prior written permission of the author(s) or other rights holders to whom publication rights have been transferred, unless permitted by a license attached to the publication (a Creative Commons license or other), or unless exceptions to copyright law apply.

Take down policy

If you believe that this document infringes your copyright or other rights, please contact openaccess@vub.be, with details of the nature of the infringement. We will investigate the claim and if justified, we will take the appropriate steps.

Sample-Level Weighting for Multi-Task Learning with Auxiliary Tasks

Emilie Grégoire^a, Hafeez Chaudhary^b and Sam Verboven^a

^aVrije Universiteit Brussel

^bRoyal Military Academy

Abstract. Multi-task learning (MTL) can improve the generalization performance of neural networks by sharing representations with related tasks. Nonetheless, MTL can also degrade performance through harmful interference between tasks. Recent work has pursued task-specific loss weighting as a solution for this interference. However, existing algorithms treat tasks as atomic, lacking the ability to explicitly separate harmful and helpful signals beyond the task level. To this end, we propose SLGrad, a sample-level weighting algorithm for multi-task learning with auxiliary tasks. Through sample-specific task weights, SLGrad reshapes the task distributions during training to eliminate harmful auxiliary signals and augment useful task signals. Substantial generalization performance gains are observed on (semi-) synthetic datasets and common supervised multi-task problems.

1 Introduction

Joint optimization with auxiliary tasks can lead to a better generalizable solution for the main task of interest while requiring less data. Recently, multi-task learning (MTL) has seen various successful applications in Deep Learning (DL), e.g., in computer vision [25, 48], reinforcement learning [31], and recommender systems [27].

Nonetheless, the optimization of MTL networks remains challenging in practice. During training, negative interference between tasks can prevent the network from reaching a good optimum. Negative inter-task interference mainly occurs for two reasons: (i) task updates providing harmful gradient directions or (ii) gradient norms being imbalanced [33, 17, 6, 10]. Existing approaches prevent negative gradient interference *at the task level* by finding appropriate task weights in the multi-task loss function [43, 8] or by aligning the different task gradients [50, 19]. At the same time, it is widely recognized in machine learning literature that some samples are more useful than others, with some even harming training, e.g., noisy data [15, 35, 37, 53].

In this paper, we abandon the widespread assertion of MTL literature of a task as an atomic unit. We disentangle the composite task signal into the sum of its parts on a sample-level. Instead of focusing on task-level weighting, this paper adopts a novel unified perspective that jointly optimizes task- and sample-level weights.

To fill the corresponding research gap, we propose SLGrad, a sample-level dynamic task-weighting algorithm for learning with

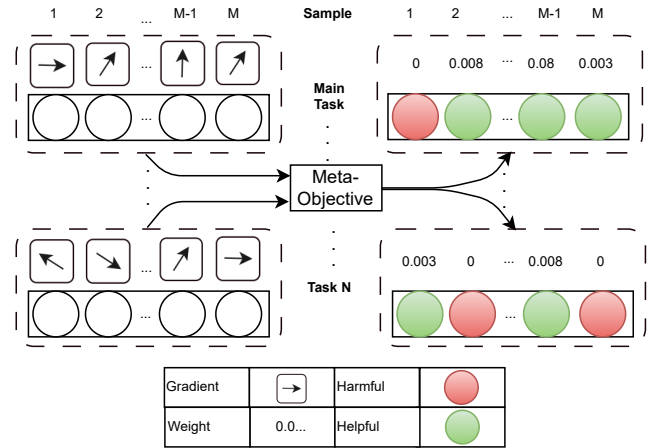


Figure 1. Pictorial representation of the SLGrad task-sample weighting algorithm. SLGrad combines the ability to focus on helpful samples with the ability to prioritize helpful auxiliary tasks for learning a task of interest.

auxiliary tasks. Different from existing task weighting algorithms, SLGrad evaluates the usefulness of samples from each task distribution with respect to the generalization performance of the main task. As such, the proposed algorithm is able to systematically avoid negative interference between different task distributions while providing meaningful sample weighting. SLGrad thus joins the benefits of an auxiliary learning approach with those of sample weighting. A pictorial representation of SLGrad is provided in Figure 1. Experiments on synthetic, semi-synthetic, and real data show that SLGrad is flexible to adjust to the variability in task and sample signals, providing state-of-the-art performance.

The contributions of this paper are summarized below:

- We bridge the existing gap between sample weighting in single-task literature and dynamic task weighting in multi-task learning.
- We propose a new dynamic task weighting algorithm for MTL with auxiliary tasks: SLGrad.
- SLGrad’s sample and task weight distributions are exhaustively studied.
- Extensive benchmark experiments on (semi-)synthetic and real-world datasets are performed.

Table 1. SLGrad innovations as compared to baseline dynamic task weighting algorithms. The checkmarks indicate which algorithms contain specific innovations discussed in the literature review.

Algorithm	Auxiliary Learning	Highly-Dynamic	Sample-Level	Meta-Objective Learning	Look-Ahead	Paper
OL-AUX	✓				✓	[31]
PCGrad		✓				[50]
CAGrad		✓				[32]
CosSim	✓	✓				[10]
GradNorm		✓				[5]
Random		✓				[29]
SLGrad	✓	✓	✓	✓	✓	Ours

2 Related Literature

In MTL literature, one either cares about optimizing all the tasks or only a single task of interest. This paper is positioned in the latter setting, which we will refer to as auxiliary learning. Here, auxiliary tasks only matter insofar as they help in improving learning the task of interest, denoted as the main task. This work draws from recent developments in dynamic task weighting, sample (i.e., importance) weighting and meta-learning literature. As such, we review these relevant strands of the literature and summarize them in Table 1.

2.1 Multi-task Learning with Auxiliary Tasks

Auxiliary MTL has been successfully applied to many complex, high-dimensional domains, e.g., natural language processing [36], computer vision [2, 34], and reinforcement learning [31]. A key challenge for MTL, and by extension auxiliary learning setting is finding the optimal balance between each task in training [16, 43].

2.2 Dynamic Task Weighting

To address the challenge of balancing task contributions, dynamic task weighting algorithms have been proposed. In deep MTL, a neural network is typically trained using a linearly weighted combination of individual task losses, i.e., the composite loss. Setting a single uniform weight for each task loss in the composite loss throughout the training of the network has been shown to lead to poor performance [38]. Furthermore, a particular task may help learn a useful base representation initially but interfere with fine-tuning later on. Consequently, task weights are often designed to be dynamic, i.e., change during training, or highly dynamic, i.e., set independently of the previous update [46, 1]. Various highly dynamic task weighting strategies exist, each adapting the loss weights or the gradients of the different tasks. Strategies include weighting based on inverse validation loss [1], geometric loss [7], and uncertainty [8]. For a complete overview, we refer to [16] and [43].

Negative Task Interference. A task is said to be negatively interfering when its learning signal negatively impacts learning the main task. Reasons could be that the gradient magnitudes are imbalanced [5, 17, 33] or that the gradient directions of the updates simply do not align with a direction increasing generalization performance [10]. To avoid task gradients with low gradient magnitudes being overshadowed by tasks with large gradients, gradient magnitudes can be normalized throughout training with respect to one main or several tasks

[5, 17]. Many different approaches to deal with harmful or conflicting gradient directions exist, e.g., dropping contributions with negative cosine similarity with the average direction [6], projecting gradients to either the normal plane of the other task [50] or an optimized vector [32, 19]. All of these strategies, however, treat tasks as atomic units, with the exception of [45]. An important research question is thus: “how to account for differences in task usefulness beyond the task level in MTL?”, as we know from the literature that some samples matter more than others.

2.3 Some Samples Matter More than Others

In single-task learning, it has been well-established that some samples matter more than others for learning well-generalizable solutions [22]. Some samples can be removed from the dataset at the beginning or during training without any loss of generalization performance [35]. Others, through data set biases such as noisy or class-imbalanced data, can explicitly harm performance [12]. Introducing sample-level weighting in the loss function is a common strategy to mitigate these biases.

Sample Weighting. Sample-level weighting for deep learning is used across many domains such as causal inference and off-policy reinforcement learning. In causal inference, inverse propensity weighting [21] balances the distributions of the control and treatment groups in observational data. Similarly, sample weighting can reduce training set bias caused by class imbalance [28, 23] or label noise [20, 37, 41]. Finally, sample weights can take into account prioritization based on domain knowledge e.g., in cost-sensitive learning [51, 52]. In effect, across these domains, weighting samples of the original source distribution leads an algorithm to learn from a semi-synthetic target distribution favourable to the task of interest. This mechanic is explicitly exploited in the importance weighting literature [4].

2.4 Generalization as a Meta-Objective

Meta-learning [42] uses the learning performance of machine learning approaches to improve the learning of the same or other tasks [44]. Recent work has instantiated hold-out (i.e., validation set) performance on the task of interest as an optimization-based meta-learning objective [37]. The incorporation of hold-out performance in the outer training loop is commonly used in the field of meta-learning to prioritize solutions that generalize better [18, 14].

3 SLGrad: Sample-Level Weighting for MTL With Auxiliary Tasks

To unify sample-level weighting with MTL, we propose SLGrad, a sample-level weighting algorithm for auxiliary learning that dynamically assigns higher weights to samples that positively contribute to the generalization performance of the main task, while discarding negatively contributing samples. This section is organized as follows: first, the problem setting is defined in subsection 3.1. Next, the main components of SLGrad are introduced and motivated in subsection 3.2. Third, the algorithm and its implementation details are formalized in subsection 3.3. Finally, the main theoretical properties of the algorithm are presented in subsection 3.4.

3.1 Problem Formulation

The objective is to optimize the generalization performance of a deep neural network on a main task \mathcal{T}_m , measured by \mathcal{M}_m , through joint training with a set of auxiliary tasks \mathcal{T}_a .

The standard multi-task loss $\mathcal{L}_{MTL}^{(t)}$ consists of a linear combination of task losses \mathcal{L}_i ,

$$\mathcal{L}_{MTL}^{(t)} = \sum_{i=1}^{N_T} w_i \mathcal{L}_i(\theta_s^{(t)}, \theta_i^{(t)}), \quad (1)$$

where N_T denotes to the number of different tasks, $\theta_s^{(t)}$ the model parameters of the shared layers, and $\theta_i^{(t)}$ the parameters of the task- i -specific layers at training step t . Finally, the coefficients w_i correspond to the task-specific weights.

The model parameters $\theta^{(t)}$, both shared and task-specific, are updated based on a mini-batch of data, using standard gradient descent algorithms on $\mathcal{L}_{MTL}^{(t)}$:

$$\theta^{(t+1)} = \theta^{(t)} - \eta \nabla_{\theta^{(t)}} \mathcal{L}_{MTL}^{(t)}, \quad (2)$$

where η is the learning rate.

This standard setup in multi-task learning only facilitates weighting at the task-level.

3.2 Proposed Solution: SLGrad

In order to separate harmful from useful signals *beyond the task level*, the standard MTL problem must be extended to facilitate the weighting of individual task-sample combinations. As such, we extend the standard multi-task loss in equation (1) to:

$$\mathcal{L}_{SLGrad}^{(t)} = \sum_{i=1}^{N_T} \sum_{j=1}^{N_B} w_{ij}^{(t)} \mathcal{L}_{ij}(y_{ij}, \hat{y}_{ij}, \theta_s^{(t)}, \theta_i^{(t)}), \quad (3)$$

where N_B represents the size of the mini-batch. In this paper, the goal is to find the task-sample-specific weights $w_{ij}^{(t)}$ such that the generalization metric \mathcal{M}_m^{val} is minimized. As such, we propose SLGrad, an algorithm that jointly solves the task and sample weighting problems. We thus tackle the MTL weighting problem at a more granular level compared to current solutions. In effect, such a solution jointly tackles the multi-task weighting problem from MTL literature, and the sample weighting problem from general machine learning literature.

Specifically, SLGrad dynamically optimizes task-sample-specific importance weights at each step of the training according to a simple and intuitive update rule. The sample-level weights $w_{ij}^{(t)}$ are computed as

$$\tilde{w}_{ij}^{(t)} = \nabla_{\theta^{(t)}} \mathcal{L}_{ij}^{train} \cdot \nabla_{\theta^{(t)}} \mathcal{M}_m^{val}, \quad (4)$$

where \mathcal{L}_{ij}^{train} represents the task-sample specific training loss and \mathcal{M}_m^{val} represents the generalization metric. Next, the sample-level weights are normalized according to

$$w_{ij}^{(t)} = \frac{\max(\tilde{w}_{ij}^{(t)}, 0)}{\sum_{i=1}^{N_T} \sum_{j=1}^{N_B} \max(\tilde{w}_{ij}^{(t)}, 0)}. \quad (5)$$

In the following section, we explain how the SLGrad weighting rule defined by equations (3) to (5) incorporates the following features:

- Prioritization of the main task
- Highly dynamic weighting
- Sample-level weighting
- Generalization as a meta-objective
- Look-ahead updating

A feature-wise comparison of SLGrad and competing algorithms is depicted in Table 1.

Prioritization of the main task. In contrast to the majority of previous dynamic weighting algorithms, which focus on a multi-objective setting, equation (4) is tuned to the auxiliary learning set-up, explicitly prioritizing learning the main task m . As implied by Theorem 2, the decrease in the generalization loss of the main task is proportional to the size of the weights $w_{i,j}$ (defined by equation 4), i.e., samples that help to decrease \mathcal{M}_m^{val} receive a higher weight.

Highly dynamic updates. We speak of highly dynamic weight updates when $w_{ij}^{(t+1)} \perp\!\!\!\perp w_{ij}^{(t)}$. As each mini-batch contains different samples, this is a necessary condition for sample-level weighting. As equation (4) does not depend on the previous weights, the condition naturally holds for SLGrad. Moreover, previous work has argued that highly dynamic updates can better adjust to variability in task usefulness on a mini-batch basis as well [46]. To assure full independency between two parameter updates, SLGrad initializes the sample-level weights to uniform weights before performing the first forward and backward passes used to compute the sample-level gradients.

Sample-level weighting. Cosine similarity is commonly used to quantify the similarity between gradient directions [43, 16]. Equation (4) reflects the cosine similarity between the gradient of a training sample (x_{ij}, y_{ij}) and the gradient of the main task generalization metric \mathcal{M}_m^{val} . Then, using equation (5), negative interference between samples and the main task is avoided by assigning zero weight to samples whose gradient is misaligned with the main task gradient ($\tilde{w}_{ij} < 0$). Samples that are well aligned with respect to this objective ($\tilde{w}_{ij} > 0$) are weighted proportionally to their similarity with the gradient direction of \mathcal{M}_m^{val} . The denominator in equation (5) ensures that the weights are normalized, avoiding fluctuations in the total learning rate [33]. Furthermore, this normalization factor is model-agnostic, meaning that we could replace equation (5) with any

other gradient normalization rule, e.g., [5, 17]. In summary, the sample weighting mechanism of SLGrad handles both sources of negative task interference: harmful gradient directions and imbalanced gradient magnitudes.

Generalization as a meta-objective. Previous algorithms that focus on main task training loss \mathcal{L}_{ij}^{train} [50, 32] are prone to overfitting. Equation (4) uses the gradient of the validation (minimization) metric $\nabla_{\theta^{(t)}} \mathcal{M}_m^{val}$ as a meta-objective that provides a proxy for generalization performance. In doing so, we draw from meta-learning literature to incorporate an explicit generalization objective. This approach contrasts with algorithms that use gradient directions computed on the main task training loss as an optimization guide. A second benefit of this approach, which we will validate experimentally, is that if this meta-objective is unbiased, the weighting mechanism will also have a debiasing effect on the learned representation.

Look-ahead update. Look-ahead approaches locally scout the loss surface to inform optimization [31, 46, 13]. The update rule in equation (4) represents a computationally efficient version of a look-ahead update [10, 46, 13], i.e., information from a throw-away update is used in $\mathcal{M}_m^{val}(\theta^{(t+1)})$. Note that the extra update to arrive at $\mathcal{M}_m^{val}(\theta^{(t+1)})$ is not explicitly calculated in our algorithm; its difference with $\mathcal{M}_m^{val}(\theta^{(t)})$ is approximated using Theorem 2. Furthermore, we justify the use of the approximation in the supplemental material (section C1).

3.3 The Algorithm

Based on equations (3) to (5), we can formulate the SLGrad algorithm:

1. Collect sample-task-specific gradients $\nabla_{\theta^{(t)}} \mathcal{L}_{ij}^{train}$ by performing a forward and backward pass on a mini-batch sampled from the training set. The backpropagated total loss is constructed with uniform task weights $w_{ij}^{(t)} = \frac{1}{N_B * N_T}$ based on the mini-batch size N_B and the number of tasks N_T .
2. Determine the gradient of the main task metric $\nabla_{\theta^{(t)}} \mathcal{M}_m^{val}$ computed by performing a forward and backward pass on a batch of the validation set.
3. Compute the cosine similarity $\nabla_{\theta^{(t)}} \mathcal{L}_{ij}^{train} \cdot \nabla_{\theta^{(t)}} \mathcal{M}_m^{val}$ between all gradients obtained in steps (1) and (2).
4. Reweight all samples of each task by following the SLGrad update rule defined in equations (4) and (5).
5. Use the computed weights $w_{ij}^{(t)}$ from step (4) to perform a final backward pass with the resulting total loss $\mathcal{L}_{SLGrad}^{(t)}$ and update the model parameters through gradient descent step $\theta^{(t)} \rightarrow \theta^{(t+1)}$ (equation 2).

The full SLGrad procedure is described in Algorithm 1.

3.4 Theoretical Analysis

In this subsection, we theoretically motivate our update rule (equations (4) and (5)) through two theorems. First, we prove that, under general assumptions, SLGrad ensures the meta-minimization objective \mathcal{M}_m^{val} monotonically decreases between every two updates.

Algorithm 1 SLGrad Algorithm

Input: Deep MTL Network $\mathcal{N}_{\mathcal{M}\mathcal{T}\mathcal{L}}$, N_T tasks, Main task metric \mathcal{M}_m^{val} , Batch size N_B , $N_T * N_B$ task losses \mathcal{L}_{ij} , learning rate η , Number of training steps N_s , Initial model parameters $\theta^{(0)}$

- 1: **For** $t = 0$ **to** $t = N_s - 1$:
 - 2: Sample batch from training set $\leftarrow (X, Y)_{train}$
 - 3: Sample batch from validation set $\leftarrow (X, Y)_{val}$
 - 4: Initialize model with model parameters $\theta^{(t)}$ and uniform sample weights $w_{ij} = (N_T * N_B)^{-1}$
 - 5: **Compute gradient validation metric**
 $\nabla_{\theta^{(t)}} \mathcal{M}_m^{val} \leftarrow \mathcal{N}_{MTL}[(X, Y)_{val}, \theta^{(t)}]$
 - 6: **Compute sample-level task gradients**
For $i \in \{0, \dots, N_T - 1\}$:
For $j \in \{0, \dots, N_B - 1\}$:
 $\nabla_{\theta^{(t)}} \mathcal{L}_{ij}^{(t)} \leftarrow \mathcal{N}_{MTL}[(x_{ij}, y_{ij})_{train}, \theta^{(t)}, w_{ij}]$
 - 7: **Determine cosine similarity**
 $\tilde{w}_{ij}^{(t)} \leftarrow \nabla_{\theta^{(t)}} \mathcal{M}_m^{val} \cdot \nabla_{\theta^{(t)}} \mathcal{L}_{ij}^{(t)}$
 - 8: **Clamp and normalize all weights**
 $w_{ij}^{(t)} = \frac{\max(\tilde{w}_{ij}^{(t)}, 0)}{\sum_i \sum_j \max(\tilde{w}_{ij}^{(t)}, 0)}$
 - 9: **Update model parameters with new weights**
 $\theta^{(t+1)} \leftarrow \theta^{(t)} - \eta \nabla_{\theta^{(t)}} \mathcal{L}_{SLGrad}$
 - 10: **End For**
-

Second, we show that the cosine similarity between sample level gradients $\nabla_{\theta^{(t)}} \mathcal{L}_{ij}^{train}$ and the gradient of the generalization metric $\nabla_{\theta^{(t)}} \mathcal{M}_m^{val}$ is approximately proportional to the decrease in the main task metric \mathcal{M}_m^{val} . Finally, we discuss SLGrad’s computational efficiency.

Theorem 1. Assume \mathcal{M}_m^{val} is a differentiable function representing the meta-(minimization)- objective as measured on the main task validation set, \mathcal{L}_{SLGrad} (as defined in equation (3)) is the total training loss and $\mathcal{L}_{ij}(\theta^{(t)}) \geq 0$ is the sample-level loss corresponding to sample i of task j . If the sample-level weights w_{ij} are computed by the SLGrad update rule (4) and (5), and the model parameter vector $\theta^{(t)}$ is updated through gradient descent (2), then the following statement holds:

$$\mathcal{M}_m^{val}(\theta^{(t+1)}) - \mathcal{M}_m^{val}(\theta^{(t)}) \leq 0 \quad (6)$$

$\forall t \in [0, \#trainingsteps]$.

The proof of this Theorem is provided in the Supplemental Material (section A).

If the validation metric \mathcal{M}_m^{val} in Theorem 1 is also bounded by below an important implication is the convergence of the main task validation metric to a minimum. Additionally, assuming \mathcal{M}_m^{val} is convex, this will correspond to the global minimum. The equality in equation (6) holds only when an optimum of \mathcal{L}_{SLGrad} or \mathcal{M}_m^{val} is reached.

Next, extending upon the proofs by [46, 10] we show that the sample-task level weights resulting from equation (4) are proportional to the decrease of the main task objective metric \mathcal{M}_m^{val} .

Table 2. Comparison of the test set performance (Mean Squared Error) for the toy set-up with different noise levels. SLGrad significantly outperforms the other benchmarks. The hyperparameters are optimized for each algorithm. We compute the mean performance over different initializations.

Noise	Static	OL-AUX	PCGrad	CAGrad	CosSim	GradNorm	Random	SLGrad
40 %	0.86 ± 0.03	0.89 ± 0.02	0.90 ± 0.02	0.89 ± 0.02	0.73 ± 0.38	0.88 ± 0.01	0.59 ± 0.35	$0.03 \pm 3E^{-4}$
70 %	$1.07 \pm 3E^{-3}$	1.27 ± 0.02	1.37 ± 0.05	1.13 ± 0.10	0.94 ± 0.18	1.04 ± 0.02	$1.05 \pm 6E^{-4}$	$0.04 \pm 9E^{-4}$

Theorem 2. Assume \mathcal{M}_m^{val} is a differentiable function, representing the meta-objective as measured on the main task validation set and \mathcal{L}_{SLGrad} the total training loss as defined in equation (3). Consider $\theta^{(t)}$ as the model parameter vector at time t . If the sample-level weights w_{ij} are computed by the SLGrad update rule (defined by equations (4) and (5)), and the model parameter vector $\theta^{(t)}$ is updated through gradient descent (equation (2)), then the decrease in the validation metric can be approximated as

$$\mathcal{M}_m^{val}(\theta^{(t+1)}) - \mathcal{M}_m^{val}(\theta^{(t)}) \approx -\eta[\nabla_{\theta^{(t)}} \mathcal{M}_m^{val}]^T \nabla_{\theta^{(t)}} \mathcal{L}_{SLGrad}, \quad (7)$$

where η is the learning rate.

The proof is provided in the Supplemental Material (section A). Theorem 2 shows that (i) the higher the sample weight $w_{ij}^{(t)}$ (equation 4), the greater the decrease of the validation metric $\mathcal{M}_m(\theta^{(t+1)})$ the sample induces, and (ii) that no explicit computation of $\mathcal{M}_m(\theta^{(t+1)})$ is required. As such, no explicit computation of the look-ahead update is needed, resulting in greater computational efficiency for the algorithm. We provide experimental justification for the use of the approximation in the Supplemental Material (section C).

The computational efficiency of SLGrad. SLGrad has a total computational overhead, compared to static MTL, with one forward and two backward passes. However, two main efficiency measures are taken: compared to other algorithms that apply a similar validation meta-objective [37], we save another forward pass and parameter update by using the approximation in equation (S.28). Furthermore, the recently introduced Python library *Opacus* [49] facilitates a dramatical speed-up in the calculation of sample-level gradients (i.e., line 6 in Algorithm 1).

4 Simulation Study

Using a toy set-up, we explore how SLGrad handles (i) the sample-level weighting problem, and (ii) the task-level weighting problem.

Aside from a comparison of SLGrad to state-of-the-art dynamic weighting algorithms, we investigate:

- the sample-level weighting distribution throughout training.
- the influence of sample-level weighting on aggregate task prioritization.

4.1 Set-up

Following [46, 5], a multi-task toy regression set-up is generated based on the following function:

$$y_{out} = f_i(x_{in}) = \sigma_i \tanh[(\mathbf{B} + \epsilon_i)\mathbf{x}], \quad (8)$$

where x_{in} is the ten-dimensional input vector. The common basis \mathbf{B} and task-dependent ϵ_i represent constant matrices generated IDD from $\mathcal{N}(0, \sqrt{1})$ and $\mathcal{N}(0, \sqrt{3.5})$ respectively. Additionally, σ_i is a task-dependent scalar that determines the scale of each task. To ensure the beneficial potential of effective sample-level weighting in this scenario, we corrupt part of the training data by adding noise as follows $y_{out} = y_{out} + \mathcal{N}(0, \sqrt{2})$. Samples from the original distribution are labeled as *clean*, corrupted samples as *noisy*. We use the validation set loss as the meta-objective. The full implementation details, including the model backbone used for all models, are included in supplement B.

4.2 Results

Our results on the toy example show that SLGrad is able to jointly solve the sample- and task-level weighting problems.

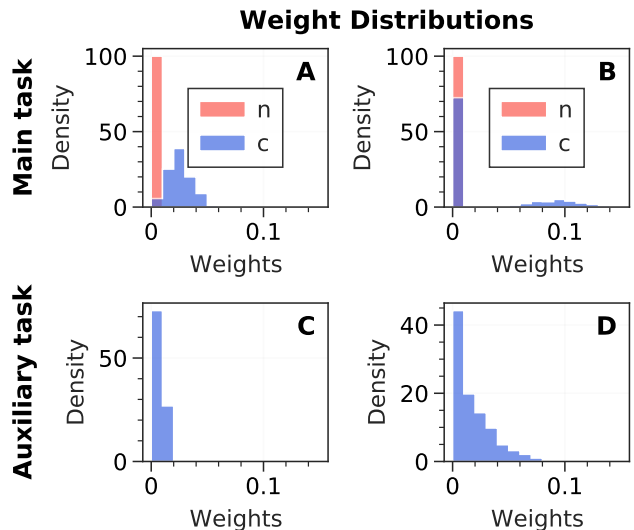


Figure 2. Weight distributions of the main task, for clean and noisy samples (top), and for an auxiliary task (bottom). Panel A is extracted in epoch 5, B in epoch 30. Panels C and D depict the auxiliary task weight distributions for 40 and 70 percent noise, respectively. SLGrad shifts weight from the full clean distribution to hard examples, and to relatively helpful tasks. We provided the learning curves corresponding to the experiments corresponding to Figure 2 in the Supplemental Material (section C).

Sample-level weighting. Figure 2 shows that during the initial stages of training (panel A), SLGrad systematically cancels out noisy samples by assigning zero weight. At the same time, clean samples are effectively exploited (i.e., positive weights) to learn the objective distribution. As training progresses, fewer clean samples remain

useful for improving generalization performance (panel C). The progression from panel A to B illustrates how the algorithm first learns the general distribution and then focuses on the harder examples (cfr. hard example mining), while still ignoring the noisy samples. These dynamics were observed for the main task and the auxiliary tasks alike.

Task-level weighting. Going from panel C to D on Figure 2, we observe that SLGrad increasingly assigns positive weights to samples belonging to \mathcal{T}_a to learn a good representation for \mathcal{T}_m . On aggregate, the auxiliary task distribution thus receives more weight as the main task becomes more noisy. This finding appeals to the general intuition that the weight distribution should adapt to a shift of signal across tasks. Finally, we note that SLGrad significantly outperforms the MTL benchmarks (Table 2) in terms of generalization performance. This performance discrepancy is the direct result of classic MTL algorithms being unable to adapt weights beyond the task level. As a result, the benchmark algorithms cannot effectively differentiate between the two types of samples during model training.

5 Real-World Experiments

In this section, the generalization performance of SLGrad is benchmarked in a series of experiments based on real-world and semi-synthetic data. First, the empirical set-up and baseline algorithms are introduced in subsection 5.1 and subsection 5.2 respectively. Second, the resulting empirical analysis is presented in subsection 5.3.

5.1 Set-up

We evaluate the generalization performance and robustness of SLGrad on two real-world datasets commonly used in MTL benchmarks: CIFAR-10 [24] and Multi-MNIST [39]. For all experiments, we adapt the LeNet-5 architecture [26] as the backbone for the shared layers. Binary cross-entropy and cross-entropy are used as the task-specific loss functions for CIFAR-10 and Multi-MNIST, respectively. All hyperparameters are optimized through grid search for each algorithm while averaging several initializations. The full configurations for all the experiments are in the supplemental material (section B). For reproducibility purposes, the Python code used to conduct the experiments in this paper will be released.

Experiments. We use the following set-ups for Multi-MNIST, CIFAR-10:

- *Clean*: the original dataset without any semi-synthetic interventions. For Multi-MNIST and CIFAR-10, the leftmost digit and the airplanes class are chosen as the respective main tasks.

Additionally, the effectiveness of SLGrad is evaluated for the following common semi-synthetic extensions [37] of CIFAR-10:

- *Uniform Label Flips*: samples from all classes are assigned a different label with 40% and 70% probability, respectively.
- *Background Label Flips*: all flipped labels (20 percent) are flipped to the same class, i.e., the background class. This type of flip also induces class imbalance as the background class will be likely to dominate.

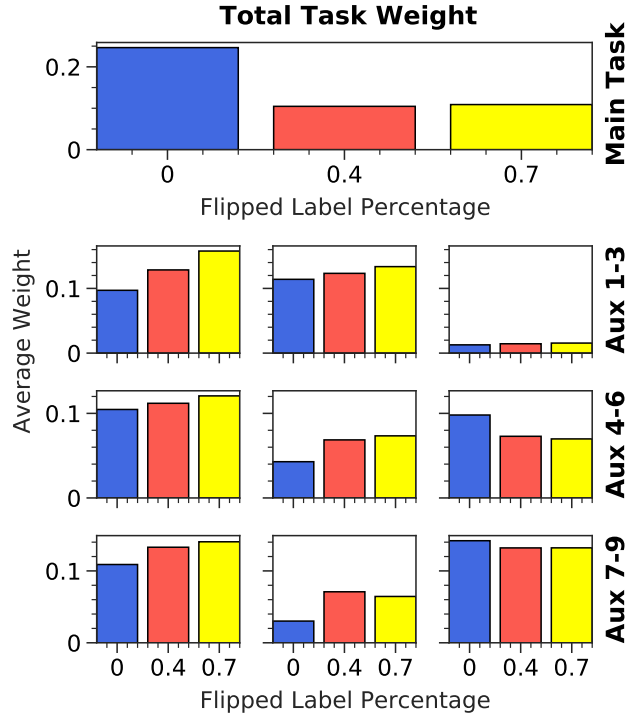


Figure 3. The total average task weights for the main task and the nine auxiliary tasks assigned by SLGrad on CIFAR-10 for 0 (blue), 40 (red) and 70 (yellow) percent of flipped labels. The average task weights are computed by summing over the sample-level weights for all tasks and then taking the time average. While the weights of the main task decrease for an increasing amount of noise, the auxiliary task weights increase.

More information about the experimental setups is provided in the supplemental material, section B.

5.2 Baseline Algorithms

We compare SLGrad with seven baseline algorithms: *Static Weighting*: A standard baseline with static and uniformly assigned task weights. *Random Weighting (RW)*: A simple weighting method where a model is trained with randomly assigned weights [29]. *Cosine Similarity (CosSim)*: This algorithm only considers auxiliary task signals when the corresponding gradient aligns with the main task gradient direction [10]. *OL-AUX*: An online learning algorithm that uses auxiliary task gradient directions, computed on previous batches to update auxiliary task weights. The gradient directions are computed on past batches and updated with gradient descent every few steps [31]. *PCGrad*: An algorithm that projects conflicting gradient directions of a task to the normal plane of any other task it negatively interferes with [50]. *CAGrad*: Conflict-Averse Gradient descent that aims to simultaneously minimize the average loss and leverages the worst local improvement of individual tasks [32]. *GradNorm*: Gradient normalization algorithm that automatically tunes gradient magnitudes to dynamically balance training in deep multi-task models [5]. Selection of the baseline algorithms was performed based on the width of adoption and current state-of-the-art performance in dynamic MTL.

Table 3. Comparison with state-of-the-art MTL weighting algorithms on original and label flipped CIFAR-10 set-ups and clean Multi-MNIST. We compare the performance in terms of (Binary) Cross-Entropy for CIFAR-10 and Multi-MNIST. The best results for each experiment are marked in bold. The hyperparameters are optimized for each algorithm. We compute the mean performance over different initializations.

Model	Clean	Uniform Flip 40	Uniform Flip 70	Background Flip	Clean Multi-MNIST
Static	0.336 ± 0.0038	0.586 ± 0.0018	0.86 ± 0.03	0.41 ± 0.09	1.95 ± 0.38
OL-AUX	0.237 ± 0.0024	0.591 ± 0.0272	0.96 ± 0.03	0.34 ± 0.03	1.16 ± 0.01
PC Grad	0.328 ± 0.0002	0.585 ± 0.0012	0.84 ± 0.05	0.36 ± 0.02	1.22 ± 0.03
CA Grad	0.235 ± 0.0031	0.569 ± 0.0038	0.94 ± 2E ⁻³	0.34 ± 0.04	1.19 ± 0.02
CosSim	0.288 ± 0.0027	0.571 ± 0.0035	0.67 ± 0.01	0.69 ± 0.14	1.38 ± 0.10
GradNorm	0.238 ± 0.0302	0.557 ± 0.0017	1.09 ± 0.02	0.33 ± 0.05	1.29 ± 0.33
Random	0.251 ± 0.0012	0.572 ± 0.0027	0.96 ± 0.02	0.29 ± 0.03	1.20 ± 0.03
SLGrad	0.249 ± 0.0072	0.288 ± 0.0263	0.33 ± 0.04	0.28 ± 0.02	0.77 ± 0.03

5.3 Results

In this subsection, we summarize the results obtained from the experiments introduced in the previous sections.

Adaptive weighting. In Figure 3, the average task weight, calculated by summing over all sample weights, is depicted for all CIFAR tasks for the three CIFAR-10 set-ups: 0, 40, and 70 percent label flips. Interestingly, we observe how SLGrad shifts more weight to the auxiliary tasks when more labels are flipped. Specifically, the average total weight increases for seven out of nine auxiliary tasks when comparing the 0 to the 40 percent scenario. In an MTL context, this implies that the algorithm increasingly exploits auxiliary tasks to learn a meaningful featurization to correctly classify the main task. The data used to generate Figure 3 is provided as a plot in the Supplemental Material (section C).




Figure 4. Learning curves for 0 (blue), 40 (red) and, 70 (yellow) percent of flipped labels for SLGrad and 5 other baseline algorithms on the semi-synthetic CIFAR 10 Uniform Flip dataset. The benchmark algorithms overfit the training set, while SLGrad effectively exploits the meta-objective signal.

Generalization performance. As shown in Table 3, SLGrad outperforms all the baseline algorithms in the flipped label set-ups. In the clean set-ups, SLGrad also performs best for Multi-MNIST, and competitively for CIFAR-10.

Figure 4 depicts the learning curves for a selection of algorithms as measured on the test set. While all baseline algorithms overfit the inserted training noise, the learning trajectory of SLGrad remains able to generalize well.

Finally, these results confirm that the benefit of sample-level weighting depends on the specific dataset and task at hand. This was previously only noted in a single-task context [47].

6 Conclusions and Future Work

In this paper, we have investigated sample-level weighting as a solution to the auxiliary learning problem. To this end, SLGrad was proposed, a simple yet flexible MTL weighting algorithm. SLGrad performs dynamic effective weighting at the sample level and can filter out data set biases by exploiting a hold out meta-objective. From a MTL perspective, this paper can be seen as an extension towards more granular, i.e., sample-level task-specific weighting (i in w_{ij}). From a sample weighting perspective, importance weighting is extended to multiple dimensions (i.e., tasks, or j in w_{ij}). The resulting unified perspective offers multiple promising avenues for future research.

Regarding applied research, depending on the meta-objective formulation, SLGrad can be applied across all domains where importance weighting used, such as causality, cost-sensitive learning, and learning with class imbalance. Additionally, SLGrad provides novel perspectives for practitioners and researchers alike to incorporate auxiliary tasks to solve the problem at hand.

Regarding fundamental research, deeper understanding of the complex interactions on a granular level between auxiliary tasks and the main task across various types of network architectures and data modalities is needed. Moving beyond the task level marks a paramount step in this process.

References

- [1] Waseem Abbas and Murtaza Tap, ‘Adaptively weighted multi-task learning using inverse validation loss’, in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1408–1412, (2019).
- [2] Noha Radwan Abhinav Valada and Wolfram Burgard, ‘Deep auxiliary learning for visual localization and odometry’, in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Brisbane, Australia, (2018).
- [3] Lukas Biewald. Experiment tracking with weights and biases, 2020. Software available from wandb.com.
- [4] Jonathon Byrd and Zachary Lipton, ‘What is the effect of importance weighting in deep learning?’, in *International Conference on Machine Learning*, pp. 872–881. PMLR, (2019).

- [5] Zhao Chen, Vijay Badrinarayanan, Chen-Yu Lee, and Andrew Rabinovich, 'GradNorm: Gradient normalization for adaptive loss balancing in deep multitask networks', in *ICML*, pp. 793–802, (2018).
- [6] 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*, eds., H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, volume 33, pp. 2039–2050. Curran Associates, Inc., (2020).
- [7] Sumanth Chennupati, Ganesh Sistu, Senthil Kumar Yogamani, and Samir A. Rawashdeh, 'Multinet++: Multi-stream feature aggregation and geometric loss strategy for multi-task learning', *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 1200–1210, (2019).
- [8] Roberto Cipolla, Yarin Gal, and Alex Kendall, 'Multi-task learning using uncertainty to weigh losses for scene geometry and semantics', in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7482–7491, (2018).
- [9] Li Deng, 'The mnist database of handwritten digit images for machine learning research', *IEEE Signal Processing Magazine*, **29**(6), 141–142, (2012).
- [10] Yunshu Du, Wojciech M. Czarnecki, Siddhant M. Jayakumar, Mehrdad Farajtabar, Razvan Pascanu, and Balaji Lakshminarayanan, 'Adapting auxiliary losses using gradient similarity', 2018.
- [11] Yunshu Du, Wojciech M. Czarnecki, Siddhant M. Jayakumar, Razvan Pascanu, and Balaji Lakshminarayanan, 'Adapting auxiliary losses using gradient similarity', *ArXiv*, **abs/1812.02224**, (2018).
- [12] Tongtong Fang, Nan Lu, Gang Niu, and Masashi Sugiyama, 'Rethinking importance weighting for deep learning under distribution shift', *Advances in Neural Information Processing Systems*, **33**, 11996–12007, (2020).
- [13] Chris Fifty, Ehsan Amid, Zhe Zhao, Tianhe Yu, Rohan Anil, and Chelsea Finn, 'Efficiently identifying task groupings for multi-task learning', *Advances in Neural Information Processing Systems*, **34**, 27503–27516, (2021).
- [14] Chelsea Finn, Pieter Abbeel, and Sergey Levine, 'Model-agnostic meta-learning for fast adaptation of deep networks', in *International conference on machine learning*, pp. 1126–1135. PMLR, (2017).
- [15] Yoav Freund and Robert E Schapire, 'A decision-theoretic generalization of on-line learning and an application to boosting', *Journal of Computer and System Sciences*, **55**(1), 119–139, (1997).
- [16] Ting Gong, Tyler Lee, Cory Stephenson, Venkata Renduchintala, Suchismita Padhy, Anthony Ndirango, Gokce Keskin, and Oguz H. Elilbol, 'A comparison of loss weighting strategies for multi task learning in deep neural networks', *IEEE Access*, **7**, 141627–141632, (2019).
- [17] Yun He, Xue Feng, Cheng Cheng, Geng Ji, Yunsong Guo, and James Caverlee, 'Metabalance: Improving multi-task recommendations via adapting gradient magnitudes of auxiliary tasks', in *Proceedings of the ACM Web Conference 2022*, WWW '22, p. 2205–2215, New York, NY, USA, (2022). Association for Computing Machinery.
- [18] T. Hospedales, A. Antoniou, P. Micaelli, and A. Storkey, 'Meta-learning in neural networks: A survey', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **44**(09), 5149–5169, (sep 2022).
- [19] Adrián Javaloy and Isabel Valera, 'RotoGrad: Gradient Homogenization in Multitask Learning', *arXiv e-prints*, arXiv:2103.02631, (March 2021).
- [20] Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei, 'Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels', in *ICML*, (2018).
- [21] Yonghan Jung, Jin Tian, and Elias Bareinboim, 'Estimating causal effects using weighting-based estimators', *Proceedings of the AAAI Conference on Artificial Intelligence*, **34**(06), 10186–10193, (Apr. 2020).
- [22] Angelos Katharopoulos and Francois Fleuret, 'Not all samples are created equal: Deep learning with importance sampling', in *Proceedings of the 35th International Conference on Machine Learning*, eds., Jennifer Dy and Andreas Krause, volume 80 of *Proceedings of Machine Learning Research*, pp. 2525–2534. PMLR, (10–15 Jul 2018).
- [23] Salman Hameed Khan, Mohammed Bennamoun, Ferdous Ahmed Sohel, and Roberto Togneri, 'Cost sensitive learning of deep feature representations from imbalanced data', *CoRR*, **abs/1508.03422**, (2015).
- [24] Alex Krizhevsky, 'Learning multiple layers of features from tiny images', Technical report, (2009).
- [25] Isabelle Leang, Ganesh Sistu, Fabian Bürger, Andrei Bursuc, and Senthil Yogamani, 'Dynamic task weighting methods for multi-task networks in autonomous driving systems', in *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, pp. 1–8, (2020).
- [26] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner, 'Gradient-based learning applied to document recognition', *Proceedings of the IEEE*, **86**(11), 2278–2324, (1998).
- [27] Hui Li, Yanlin Wang, Ziyu Lyu, and Jieming Shi, 'Multi-task learning for recommendation over heterogeneous information network', *IEEE Transactions on Knowledge and Data Engineering*, **34**(2), 789–802, (2022).
- [28] Kuan Li, Xiangfei Kong, Zhi Lu, Liu Wenyin, and Jianping Yin, 'Boosting weighted elm for imbalanced learning', *Neurocomput.*, **128**, 15–21, (mar 2014).
- [29] Baijiong Lin, Feiyang Ye, Yu Zhang, and Ivor W. Tsang, 'Reasonable Effectiveness of Random Weighting: A Litmus Test for Multi-Task Learning', *arXiv e-prints*, arXiv:2111.10603, (November 2021).
- [30] Baijiong Lin and Yu Zhang, 'LibMTL: A python library for multi-task learning', *arXiv preprint arXiv:2203.14338*, (2022).
- [31] Xingyu Lin, Harjatin Baweja, George Kantor, and David Held, 'Adaptive auxiliary task weighting for reinforcement learning', in *Advances in Neural Information Processing Systems*, eds., H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, volume 32. Curran Associates, Inc., (2019).
- [32] 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*, eds., M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, volume 34, pp. 18878–18890. Curran Associates, Inc., (2021).
- [33] Liyang Liu, Yi Li, Zhanghui Kuang, Jing-Hao Xue, Yimin Chen, Wenming Yang, Qingmin Liao, and Wayne Zhang, 'Towards impartial multi-task learning', in *International Conference on Learning Representations*, (2021).
- [34] Xianpeng Liu, Nan Xue, and Tianfu Wu, 'Learning auxiliary monocular contexts helps monocular 3d object detection', *Proceedings of the AAAI Conference on Artificial Intelligence*, **36**(2), 1810–1818, (Jun. 2022).
- [35] Mansheej Paul, Surya Ganguli, and Gintare Karolina Dziugaite, 'Deep learning on a data diet: Finding important examples early in training', in *Advances in Neural Information Processing Systems*, eds., M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, volume 34, pp. 20596–20607. Curran Associates, Inc., (2021).
- [36] Bowen Qin, Min Yang, Lidong Bing, Qingshan Jiang, Chengming Li, and Ruifeng Xu, 'Exploring auxiliary reasoning tasks for task-oriented dialog systems with meta cooperative learning', *Proceedings of the AAAI Conference on Artificial Intelligence*, **35**(15), 13701–13708, (May 2021).
- [37] Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun, 'Learning to reweight examples for robust deep learning', in *Proceedings of the 35th International Conference on Machine Learning*, eds., Jennifer Dy and Andreas Krause, volume 80 of *Proceedings of Machine Learning Research*, pp. 4334–4343. PMLR, (10–15 Jul 2018).
- [38] Sebastian Ruder, 'An overview of multi-task learning in deep neural networks', (06 2017).
- [39] Sara Sabour, Nicholas Frosst, and Geoffrey E Hinton, 'Dynamic routing between capsules', *Advances in neural information processing systems*, **30**, (2017).
- [40] Ozan Sener and Vladlen Koltun, 'Multi-task learning as multi-objective optimization', *Advances in neural information processing systems*, **31**, (2018).
- [41] Jun Shu, Qi Xie, Lixuan Yi, Qian Zhao, Sanping Zhou, Zongben Xu, and Deyu Meng, 'Meta-weight-net: Learning an explicit mapping for sample weighting', in *NeurIPS*, (2019).
- [42] Sebastian Thrun, 'Lifelong learning algorithms', in *Learning to learn*, 181–209, Springer, (1998).
- [43] Simon Vandenhende, Stamatios Georgoulis, Wouter Van Gansbeke, Marc Proesmans, Dengxin Dai, and Luc Van Gool, 'Multi-task learning for dense prediction tasks: A survey', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **44**(7), 3614–3633, (2022).
- [44] Joaquin Vanschoren, 'Meta-learning', in *Automated machine learning*, 35–61, Springer, Cham, (2019).
- [45] Pavan Kumar Anasosalu Vasu, Shreyas Saxena, and Oncel Tuzel, 'Instance-level task parameters: A robust multi-task weighting framework', 2021.
- [46] Sam Verboven, Hafeez Chaudhary, Jeroen Berrevoets, Vincent Ginis, and Wouter Verbeke, 'Hydlearn', *Applied Intelligence*, 1–15, (07

- 2022).
- [47] Da Xu, Yuting Ye, and Chuanwei Ruan, 'Understanding the role of importance weighting for deep learning', in *International Conference on Learning Representations*, (2021).
 - [48] Xi Yin and Xiaoming Liu, 'Multi-task convolutional neural network for pose-invariant face recognition', *IEEE Transactions on Image Processing*, **27**(2), 964–975, (2018).
 - [49] Ashkan Yousefpour, Igor Shilov, Alexandre Sablayrolles, Davide Testuggine, Karthik Prasad, Mani Malek, John Nguyen, Sayan Ghosh, Akash Bharadwaj, Jessica Zhao, et al., 'Opacus: User-friendly differential privacy library in pytorch', *arXiv preprint arXiv:2109.12298*, (2021).
 - [50] 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*, eds., H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, volume 33, pp. 5824–5836. Curran Associates, Inc., (2020).
 - [51] B. Zadrozny, J. Langford, and N. Abe, 'Cost-sensitive learning by cost-proportionate example weighting', in *Third IEEE International Conference on Data Mining*, pp. 435–442, (2003).
 - [52] Huimin Zhao, 'Instance weighting versus threshold adjusting for cost-sensitive classification', *Knowledge and Information Systems*, **15**(3), 321–334, (2008).
 - [53] Zhaowei Zhu, Tongliang Liu, and Yang Liu, 'A second-order approach to learning with instance-dependent label noise', in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10113–10123, (June 2021).

Supplement A: Theoretical proofs

The first supplement consists of the theoretical proofs of the two Theorems proposed in section 3.4 of the main paper.

A1. Proof of Theorem 1

Theorem 1. Assume \mathcal{M}_m^{val} is a differentiable function representing the meta-(minimization)- objective as measured on the main task validation set, \mathcal{L}_{SLGrad} (as defined in equation 3) is the total training loss and $\mathcal{L}_{ij}(\theta^{(t)}) \geq 0$ is the sample-level loss corresponding to sample i of task j . If the sample-level weights w_{ij} are computed by the SLGrad update rule (defined by equations 3 to 5), and the model parameter vector $\theta^{(t)}$ is updated through gradient descent (equation 2), then the following statement holds:

$$\mathcal{M}_m^{val}(\theta^{(t+1)}) - \mathcal{M}_m^{val}(\theta^{(t)}) \leq 0 \quad (\text{S.9})$$

$\forall t \in [0, \#trainingsteps]$.

The proof of Theorem 1 consists of two lemma's, which we will proof first.

Lemma 1. Consider N_B as the batch size and N_T as the number of tasks. If the assumptions made in Theorem 1 hold and if $\nabla_{\theta^{(t)}} \mathcal{L}_{ij}(\theta^{(t)}) \cdot \nabla_{\theta^{(t)}} \mathcal{M}_m^{val}(\theta^{(t)}) = 0 \forall i \in \{0, N_B - 1\}$ and $\forall j \in \{0, N_T - 1\}$ then:

$$\mathcal{M}_m^{val}(\theta^{(t+1)}) \approx \mathcal{M}_m^{val}(\theta^{(t)}) \quad (\text{S.10})$$

holds $\forall t$.

Proof. As \mathcal{M}_m^{val} is a differentiable, scalar-valued function, the first-order Taylor series expansion about a point $\theta^{(t)}$ is

$$\forall t : \mathcal{M}_m^{val}(\theta^{(t+1)}) \quad (\text{S.11})$$

$$= \mathcal{M}_m^{val}(\theta^{(t)}) - \eta \nabla_{\theta^{(t)}} \mathcal{L}_{SLGrad}(\theta^{(t)}) \quad (\text{S.12})$$

$$\approx \mathcal{M}_m^{val}(\theta^{(t)}) - \eta [\nabla_{\theta^{(t)}} \mathcal{M}_m^{val}(\theta^{(t)})]^T \nabla_{\theta^{(t)}} \mathcal{L}_{SLGrad}(\theta^{(t)}). \quad (\text{S.13})$$

Using the SLGrad update rule (defined by equations 3 to 5 of the main paper), we find that

$$\nabla_{\theta^{(t)}} \mathcal{L}_{SLGrad}(\theta^{(t)}) \quad (\text{S.14})$$

$$= \sum_i^{N_B} \sum_j^{N_T} \nabla_{\theta^{(t)}} w_{ij} \mathcal{L}_{ij}(\theta^{(t)}) \quad (\text{S.15})$$

$$= \sum_i^{N_B} \sum_j^{N_T} \max[0, \nabla_{\theta^{(t)}} \mathcal{L}_{ij}(\theta^{(t)}) \cdot \nabla_{\theta^{(t)}} \mathcal{M}_m] \nabla_{\theta^{(t)}} \mathcal{L}_{ij}(\theta^{(t)}) \quad (\text{S.16})$$

$$= 0, \quad (\text{S.17})$$

where we used $\forall i, j : \nabla_{\theta^{(t)}} \mathcal{L}_{ij}(\theta^{(t)}) \cdot \nabla_{\theta^{(t)}} \mathcal{M}_m = 0$ in the last step. Inserting this into equation (S.13), we thus obtain

$$\mathcal{M}_m^{val}(\theta^{(t+1)}) - \mathcal{M}_m^{val}(\theta^{(t)}) \approx 0. \quad (\text{S.18})$$

□

Lemma 2. Consider N_B as the batch size and N_T as the number of tasks. If the assumptions made in Theorem 1 hold and if $\exists i \in \{0, N_B - 1\}$ and $j \in \{0, N_T - 1\}$ such that

$$\nabla_{\theta^{(t)}} \mathcal{L}_{ij}(\theta^{(t)}) \cdot \nabla_{\theta^{(t)}} \mathcal{M}_m^{val}(\theta^{(t)}) \neq 0, \quad (\text{S.19})$$

then:

$$\mathcal{M}_m^{val}(\theta^{(t+1)}) - \mathcal{M}_m^{val}(\theta^{(t)}) < 0. \quad (\text{S.20})$$

Proof. Approximate $\mathcal{M}_m^{val}(\theta^{(t+1)})$ by its first-order Taylor series expansion

$$\mathcal{M}_m^{val}(\theta^{(t+1)}) \approx \mathcal{M}_m^{val}(\theta^{(t)}) - \eta [\nabla_{\theta^{(t)}} \mathcal{M}_m^{val}(\theta^{(t)})]^T \nabla_{\theta^{(t)}} \mathcal{L}_{SLGrad}(\theta^{(t)}). \quad (\text{S.21})$$

We thus need to prove that

$$\eta[\nabla_{\theta^{(t)}} \mathcal{M}_m^{val}(\theta^{(t)})]^T \nabla_{\theta^{(t)}} \mathcal{L}_{SLGrad}(\theta^{(t)}) > 0. \quad (\text{S.22})$$

As $\eta > 0$ and we can ignore contributions for which $w_{ij} = 0$, this reduces to proving that the following sum is positive:

$$\begin{aligned} & \sum_{i:w_{ij} \neq 0}^{N_B} \sum_{j:w_{ij} \neq 0}^{N_T} w_{ij} \nabla_{\theta^{(t)}} [\mathcal{M}_m^{val}(\theta^{(t)})]^T \nabla_{\theta^{(t)}} \mathcal{L}_{ij}(\theta^{(t)}) \\ &= \sum_{i:w_{ij} \neq 0}^{N_B} \sum_{j:w_{ij} \neq 0}^{N_T} w_{ij} \nabla_{\theta^{(t)}} \mathcal{M}_m^{val}(\theta^{(t)}) \cdot \nabla_{\theta^{(t)}} \mathcal{L}_{ij}(\theta^{(t)}). \end{aligned} \quad (\text{S.23})$$

Next, we used the fact that $[\nabla_{\theta^{(t)}} \mathcal{M}_m^{val}(\theta^{(t)})]^T$ and $\nabla_{\theta^{(t)}} \mathcal{L}_{ij}(\theta^{(t)})$ are row and column vectors respectively, such that

$$\begin{aligned} & [\nabla_{\theta^{(t)}} \mathcal{M}_m^{val}(\theta^{(t)})]^T \nabla_{\theta^{(t)}} \mathcal{L}_{ij}(\theta^{(t)}) \\ &= \nabla_{\theta^{(t)}} \mathcal{M}_m^{val}(\theta^{(t)}) \cdot \nabla_{\theta^{(t)}} \mathcal{L}_{ij}(\theta^{(t)}). \end{aligned} \quad (\text{S.24})$$

All the remaining w_{ij} are now positive and non-zero by definition. Therefore, using that $w_{ij} = 0$ if $\nabla_{\theta^{(t)}} \mathcal{M}_m^{val}(\theta^{(t)}) \cdot \nabla_{\theta^{(t)}} \mathcal{L}_{ij}(\theta^{(t)}) \leq 0$, we know that for all the terms occurring in the sum in equation (S.23):

$$\nabla_{\theta^{(t)}} \mathcal{M}_m^{val}(\theta^{(t)}) \cdot \nabla_{\theta^{(t)}} \mathcal{L}_{ij}(\theta^{(t)}) > 0, \quad (\text{S.25})$$

Consequently,

$$\sum_{i:w_{ij} \neq 0}^{N_B} \sum_{j:w_{ij} \neq 0}^{N_T} \eta w_{ij} \nabla_{\theta^{(t)}} \mathcal{M}_m^{val}(\theta^{(t)}) \cdot \nabla_{\theta^{(t)}} \mathcal{L}_{ij}(\theta^{(t)}) > 0, \quad (\text{S.26})$$

such that

$$\mathcal{M}_m^{val}(\theta^{(t+1)}) < \mathcal{M}_m^{val}(\theta^{(t)}), \quad (\text{S.27})$$

which proves the Lemma. \square

We prove Theorem 1 by combining Lemma's 1 and 2.

If we further assume, the generalization metric \mathcal{M}_m^{val} is bounded from below, the non-decreasing behavior implies that the metric convergence to a metric due to the monotone convergence theorem. Additionally, if \mathcal{M}_m^{val} is convex, this minimum corresponds to the global minimum.

A2. Proof of Theorem 2

The proof of Theorem 2 is given below.

Theorem 2. Assume \mathcal{M}_m is a differentiable function, representing the validation metric on the main task and $\mathcal{L}_{SLGrad}^{(t)}$ the total training loss (as defined in equation 3 in the main paper). Consider $\theta^{(t)}$ as the model parameter vector at time t . If the sample-level weights w_{ij} are computed by the SLGrad update rule (defined by equations 3 to 5), and the model parameter vector $\theta^{(t)}$ is updated through gradient descent (equation 2), then the gain in the validation metric, can be approximated as

$$\mathcal{M}_m^{val}(\theta^{(t+1)}) - \mathcal{M}_m^{val}(\theta^{(t)}) \approx -\eta[\nabla_{\theta^{(t)}} \mathcal{M}_m^{val}]^T \nabla_{\theta^{(t)}} \mathcal{L}_{SLGrad}, \quad (\text{S.28})$$

where η is the learning rate.

Proof. $\forall t \in [0, \#trainingsteps]$ we can expand the main task metric $\mathcal{M}_m(\theta^{(t+1)})$ as

$$\mathcal{M}_m^{val}(\theta^{(t+1)}) \approx \mathcal{M}_m^{val}(\theta^{(t)}) - \eta \nabla_{\theta^{(t)}} \mathcal{L}_{SLGrad} \quad (\text{S.29})$$

$$\approx \mathcal{M}_m^{val}(\theta^{(t)}) - \eta[\nabla_{\theta^{(t)}} \mathcal{M}_m^{val}(\theta^{(t)})]^T \nabla_{\theta^{(t)}} \mathcal{L}(\theta^{(t)}), \quad (\text{S.30})$$

where we used the first-order Taylor approximation of a differentiable, scalar-valued function about a point $\theta^{(t)}$. Therefore, the gain in the validation metric, resulting from the parameter update $\theta^{(t)} \rightarrow \theta^{(t+1)}$, can be approximated as

$$\mathcal{M}_m^{val}(\theta^{(t+1)}) - \mathcal{M}_m^{val}(\theta^{(t)}) \approx -\eta[\nabla_{\theta^{(t)}} \mathcal{M}_m^{val}]^T \nabla_{\theta^{(t)}} \mathcal{L}_{SLGrad}. \quad (\text{S.31})$$

\square

Experimental justification for the use of the Taylor approximation is provided in section C1.

Supplement B: Implementation details

The implementation details comprising the data, the baseline neural network and the hyperparameter configuration are discussed for all experimental setups. Accordingly, this section is organized as follows; the simulation study is discussed in subsection B1, CIFAR-10 in subsection B2, and Multi-MNIST in subsection B3. Pytorch was used for the implementation of all the experiments. For the implementation of CA-Grad and PCGrad, we adapted code from LibMTL [30]. Note that we logged all the metrics, losses, and hyperparameter configurations using Weights & Biases [3]. The corresponding projects will be made available.

B1. Simulation Study

This subsection contains detailed information about the simulation study presented in section 4 of the paper. Concretely, we elaborate on how the toy regression tasks were generated in subsection B1.1. The neural network used as a backbone for the experiments and the corresponding hyperparameter configurations are presented in subsections B1.2 and B1.3 respectively.

B1.1 Data

The toy multi-task regression set-up is generated from the following function:

$$y_{out} = f_i(x_{in}) = \sigma_i \tanh[(\mathbf{B} + \epsilon_i)\mathbf{x}], \quad (\text{S.32})$$

where x_{in} is the ten-dimensional input vector. As such, we followed the approach of [5] and [46]. The validation and test sets contain 200 samples while 1000 samples are used to train the network. We generated one auxiliary task with the same function as the main task (defined by equation (S.32)) and made the ϵ_i task-dependent. The common basis \mathbf{B} and task-dependent ϵ_i represent constant matrices generated IDD from $\mathcal{N}(0, \sqrt{1})$ and $\mathcal{N}(0, \sqrt{3.5})$ respectively. As the task distributions have a common basis, it is expected that the auxiliary task to be generally helpful to the main task [5]. In order to test the robustness of the proposed algorithm, different amounts (0, 40, and 70 percent) of Gaussian noise were added to the regression task in equation (S.32):

$$y_{out} = y_{out} + \mathcal{N}(0, \sqrt{2}) \quad (\text{S.33})$$

Note that the validation set remains unaltered. By labeling each observation as (non)-noisy, we kept track of how weights are distributed over the observations. As such, weight distributions as presented on Figure 2 in the main paper were created.

B1.2 Model Architecture

The same model architecture is used for all the different toy setups and baseline algorithms: a standard multi-task neural network with two to four shared layers (64 neurons each) and one to four task-specific layers (32 neurons each) for the two tasks.

B1.3 Hyperparameters

A grid-search was performed to optimize the hyper-parameters (learning rate (LR), batch size (BS), number of shared (SL) and task-specific (TL) layers). For completeness, the grid-search was performed independently for each algorithm. The corresponding optimal hyperparameter configurations for each algorithm are presented in Table 4 below. The parameters are optimized for one and the same initialization after which the optimal configuration is applied to the additional initializations. The values from Table 2 in the main paper were obtained by averaging the best performance for three different initializations (random seeds) for each algorithm.

ALGORITHM	LR	BS	SL	TL
OL-AUX	0.1	64	2	2
PC Grad	0.1	32	3	3
CA Grad	0.1	64	2	2
Gcosim	0.01	64	3	4
Static	0.01	32	4	4
GradNorm	0.1	32	2	2
Random	0.01	64	3	4
SLGrad (ours)	0.1	32	3	4

Table 4. Optimal hyperparameters for toy dataset

B2. CIFAR-10

Implementation details corresponding to the experiments performed on the CIFAR-10 dataset are discussed in this subsection. First, in section B2.1, the standard dataset and its transformations are presented. Then, the neural network used as a backbone for the experiments and the corresponding hyperparameter configurations are provided in subsections B2.2 and B2.3 respectively.

B2.1 Data

CIFAR-10 [24] is a benchmark dataset containing 60000 color images (50000 train and 10000 test images) of size 32x32 that each belong to one of ten classes. For the experiments 20000 images were subsampled from the original training dataset for training. Additionally, the multi-label classification task was transformed into 10 binary classification tasks. One of these classification tasks (i.e., predicting if an airplane appears on an image or not) was chosen to become the main task of interest and the other nine were used as auxiliary tasks. Next to the standard “clean” CIFAR-10 multi-task setup, we also generated two semi-synthetic setups, following [37]:

- *Uniform Label Flips*: samples from all classes are assigned a different label with 40% and 70% probability, respectively.
- *Background Label Flips*: all flipped labels (20 percent) are flipped to the same class, i.e., the background class. This type of flip also induces class imbalance as the background class will be likely to dominate.

Note that, in creating the semi-synthetic setups the validation data set (20 percent of the training set) is clean for all algorithms.

B2.2 Model Architecture

We employ the LeNet-5 architecture [26] as the backbone for the experiments on CIFAR-10. All the layers were used as a shared encoder (except the fully connected layer). On top of this shared encoder, one to four fully connected layers were added as task-specific functions with ReLU activation functions. The Binary Cross Entropy with a sigmoid is used as the task-specific loss function for all tasks.

B2.3 Hyperparameters

For completeness, an independent grid-search (learning rate (LR $\in \{0.1, 0.01, 0.001\}$), number of task-specific layers (TL: 1 to 4) and batch size (BS $\in \{32, 64, 128\}$) was performed for each dynamic weighting algorithm separately. The optimal hyperparameters for each algorithm are presented in Table 2 below. Note that the optimal hyperparameters for one initialization are used on three additional initializations. The values in Table 3 in the main paper are obtained by averaging the test performance over these three different initializations (random seeds) and flipped label experiments.

ALGORITHM	LR	BS	TL
OL-AUX	0.001	64	2
PC Grad	0.01	64	2
CA Grad	0.01	64	2
GCosim	0.001	128	3
GradNorm	0.1	128	1
Static	0.1	32	3
Random	0.1	32	2
SLGrad (ours)	0.1	128	2

Table 5. Hyperparameter optimization on CIFAR-10 clean dataset.

B3. Multi-MNIST

This subsection contains detailed information about the experiments performed on the Multi-MNIST dataset. To start with, subsection B3.1 provides more information about the data. Thereafter, the neural network and its corresponding hyperparameter configurations are discussed in subsections B3.2 and B3.3 respectively.

B3.1 Data

Multi-MNIST [39] is a MTL version of the original MNIST dataset [9] generated by [39] and commonly used for benchmarking in the MTL literature. [39] generated the data by overlaying digits from the MNIST dataset on top of other digits of the same set (training or test) but from another class. As the digits are shifted up to 4 pixels in each direction, the dataset contains 36x36 pixel images. We use this dataset in the MTL setup by defining the classification of the left and right digits on each image as different tasks. The dataset contains 60k training and 20k test images but only 20k training and 5k test samples are used. The first task (identifying the left digit) is identified as the main task for the experiments while the second task (identifying the right digit) is used as an auxiliary task.

B3.2 Model Architecture

The same adaptation of the LeNet-5 architecture as for the experiments on CIFAR-10 is used. Note that the binary cross entropy with the sigmoid function was replaced by the cross entropy function with softmax. For the implementation, we adapted code from [40].

B3.3 Hyperparameters

To optimize each algorithm, a grid-search over the hyperparameters: learning rate ($LR \in \{0.1, 0.01, 0.001, 0.0001\}$) and batch size ($BS \in \{32, 64, 128\}$) was performed for each algorithm separately. The corresponding optimal hyperparameters for one initialization (random seed) are presented in Table 6. The results in Table 4 in the paper are obtained by averaging the test performance for three different initializations.

ALGORITHM	LR	BS
OL-AUX	0.001	128
PC Grad	0.1	128
CA Grad	0.001	32
Gcosim	0.001	128
Static	0.001	128
GradNorm	0.0001	128
Random	0.001	32
SLGrad (ours)	0.001	128

Table 6. Hyperparameter optimization on the Multi-MNIST dataset.

Supplement C: Additional Results

In this section, additional experimental results are presented. In subsection C1, the use of the Taylor approximation (equation S.13) is justified by experimental observations. Next, additional plots supporting the results presented in sections 4 and 5 of the paper are provided.

C1. Taylor Approximation

We compare the learning curve of SLGrad when the algorithm is applied with and without the use of Taylor approximation (S.31). Applying the approximation consists of not explicitly computing the look-ahead update (as discussed in section 3.2 of the main paper and in [11, 13]) to obtain $\mathcal{M}(\theta_{t+1})$, but instead using only $\mathcal{M}(\theta_t)$. In essence, this results in the removal of one forward and parameter update on each training pass. Figure 5 provides the comparison of the learning curves on the experiments with and without approximation. We see that the learning curves remain seemingly unaltered for the toy regression task (left panel) and similar for CIFAR-10 (right panel) where a small improvement is still observed for the explicit computation.




Figure 5. The comparison between the learning curve of SLGrad with and without the use of the Taylor approximation for the Toy simulation (A) and CIFAR-10 experiments (B).

C2. Additional Experimental Results

Additional experimental support for the results discussed in sections 4 and 5 of the main paper is presented in this section. Subsection C2.1 contains additional plots to support the weight distributions provided in section 4 of the paper. Next, subsection C2.2 contain additional plots supporting section 5 in the main paper.

C2.1 Sample Weighting

Figure 6 shows the evolution of the total task weight of the main and auxiliary tasks for the experiments on the clean toy data. The total task weights are obtained by summing over the sample level weights for each task distribution separately. At the beginning of the learning procedure, most of the weight is given to the main task, after which it decreases at the end of the learning procedure.




Figure 6. Evolution of the total task weight of the main task (A) and auxiliary task (B) for a clean toy setup.

C2.2 Task Weighting

The total task weight of the main task assigned by SLGrad when applied on Cifar-10 for different amounts of flipped labels (0, 40, 70) is shown in Figure 7. This plot contains the data used to generate Figure 3 in the main paper. As discussed in section 5 of the main paper, SLGrad decreases the weight assigned to the main task as the amount of flipped labels in the main task distribution increases.




Figure 7. Comparison of total task weight of the main task for SLGrad applied on CIFAR-10 for different amounts of flipped labels (0, 40, 70 percent). This plot was used to generate Figure 3 in the paper. The hyperparameters used for the three experiments are equal.