

Learning to Teach Fairness-aware Deep Multi-task Learning

Arjun Roy^{1,2} and Eirini Ntoutsi²

¹ L3S Research Center, Leibniz University Hannover

² Institute of Computer Science, Free University Berlin
{arjun.roy,eirini.ntoutsi}@fu-berlin.de

Abstract. Fairness-aware learning mainly focuses on single task learning (STL). The fairness implications of multi-task learning (MTL) have only recently been considered and a seminal approach has been proposed that considers the fairness-accuracy trade-off for each task and the performance trade-off among different tasks. Instead of a rigid fairness-accuracy trade-off formulation, we propose a flexible approach that learns how to be fair in a MTL setting by selecting which objective (accuracy or fairness) to optimize at each step. We introduce the *L2T-FMT* algorithm that is a teacher-student network trained collaboratively; the student learns to solve the fair MTL problem while the teacher instructs the student to learn from either accuracy or fairness, depending on what is harder to learn for each task. Moreover, this dynamic selection of which objective to use at each step for each task reduces the number of trade-off weights from $2T$ to T , where T is the number of tasks. Our experiments on three real datasets show that *L2T-FMT* improves on both fairness (12-19%) and accuracy (up to 2%) over state-of-the-art approaches.

1 Introduction

Multi-Task Learning (MTL) [21] aims to leverage useful information contained in multiple tasks to help improve the generalization performance over all tasks. It is inspired by human’s ability to learn multiple tasks and it has been already successfully applied in a variety of applications from natural language processing [12] to vision [3]. Many methods and deep neural network architectures [3, 9, 12] for MTL have been proposed, still the basic optimization problem is formulated as minimizing a weighted sum of task-specific losses, where the weights are inter-task trade-off hyperparameters used to avoid inter-task loss dominance [21].

Despite its popularity, the fairness implications of MTL have only recently come into focus [22]. The area of fairness-aware learning for Single Task Learning (STL) has received a lot of attention in the last years [15] and methods have been proposed that aim to learn correct predictions without discriminating on the basis of some protected attribute, like gender or race. Methods in this category reformulate the classification problem by explicitly incorporating the model’s discrimination behavior in the objective function through e.g., regularization or constraints. The fair-MTL problem was only recently introduced and a solution,

MTA-F, has been proposed [22] that considers *inter-task* fairness-accuracy trade-offs (as is typical in MTL) and *intra-task* performance trade-offs (as is typical in fairness-aware STL). For T tasks, such an approach requires $2T$ trade-off weights. The current practice to find the correct trade-off weights is by hyperparameter tuning [22]. The performance of MTLs are highly dependent on the correct weighting between each task’s loss [21]. Tuning these weights by hand is a difficult and expensive process [13], which further gets leveraged in the fair-MTL problem. Moreover, the trade-off weights in *MTA-F* are fixed throughout the training process. However, the inter-task trade-offs and the intra-task fairness-accuracy balance may change over the training due to e.g., factors like data batch [1].

In this paper, instead of a fixed fairness-accuracy trade-off formulation, we propose to *dynamically* select one among fairness and accuracy objectives at each training step for each task. To this end, we formulate the fair-MTL problem as a student-teacher problem and propose the Learning to Teach Fair Multi-tasking (referred to as *L2T-FMT*) algorithm. Our design inspiration comes from recent learning to teach (L2T) algorithms [6, 23]. The student in our proposed algorithm is the desired MTL model, which follows the instruction of the teacher to learn from the available accuracy or fairness objectives for each task, and updates its parameters accordingly. The student sends feedback about its progress on fairness and accuracy in each task to the teacher. The teacher learns from the feedback and updates its model. This way, both student and teacher networks are trained collaboratively. Except for the dynamic intra-task loss selection, we also propose to set the inter-task parameters *dynamically* at each training step using GradNorm [1], as opposed to fixing them throughout the training process [22].

Our contributions can be summarized as follows: i) We introduce the dynamic objective selection paradigm for fair and accurate MTL. ii) We propose a new algorithm, *L2T-FMT*, based on a student-teacher framework that executes the dynamic objective selection paradigm and efficiently solves fairness-aware MTL. iii) Our dynamic objective selection results in a reduction of parameters to be learned per training step from $2T$ to T . iv) We eliminate the dependency of searching for the correct balance of inter-task trade-off weights by automatically learning the weights at each training step. v) *L2T-FMT* outperforms state-of-the-art methods by improving on both fairness and accuracy as demonstrated on real-world datasets of varying characteristics and number of tasks.

The rest of the paper is organized as follows: In Section 2 we review the related work. Necessary basic concepts are provided in Section 3. Our method is introduced in Section 4, followed by an experimental evaluation in Section 5. Conclusions and outlook are summarized in Section 6.

2 Related work

Fairness-aware learning A growing body of work has been proposed over the last years to address the problem of fairness and algorithmic discrimination [15] against demographic groups defined on the basis of protected attributes like gender or race. In parallel to bias mitigation methods, a plethora of fairness notions

have been proposed; the interested reader is referred to [15] for a taxonomy of various fairness definitions. Statistical parity [5], equal opportunity and equalized odds [10] are among the most popular measures for measuring discrimination. In this work, we adopt equalized odds as our notion of fairness (Section 3.1).

Multi-task learning In MTL, multiple learning tasks are solved simultaneously, while exploiting commonalities and differences across the tasks [21]. There are two main categories of parameter sharing: hard vs. soft. In hard parameter sharing [1, 3, 9, 22], model weights are shared between multiple tasks, while output layers are kept task-specific. In soft parameter sharing [12], different tasks have individual task-specific models with separate weights, but the distance between the model parameters is regularized in order to encourage the parameters to be similar. In this work, we follow the most popular hard parameter sharing approach where model weights are shared between multiple tasks.

Fairness-aware multi-task learning The fairness implications of MTL have been only recently considered: [24] studies multi-task regression to improve fairness in ranking; [18] proposes an MTL formulation to solve multi-attribute fairness on a single task. The closest to our work is the seminal work [22], which formulates the fair-MTL problem as a weighted sum of task-specific accuracy-fairness trade-offs. This formulation results in duplication of parameters, which are learned via hyperparameter tuning. Moreover, [22] introduced the concept of task-exclusive labels signifying examples that are only positive (negative) for the concerned task and negative (positive) for all the other tasks. Then, they proposed the *MTA-F* algorithm that updates the task-specific layer with the summed loss of accuracy and task-specific fairness (computed with exclusive examples) and shared layers with the summed loss of accuracy and shared fairness (computed with non-exclusive examples). In our method, we do not use the task-exclusive concept, as in the presence of a large number of tasks, the exclusive set of instances may reduce to null. Moreover, contrary to [22] that assumes a fixed intra-task accuracy-fairness trade-off, we rather *learn to choose at each step* of the training process whether the accuracy loss or the fairness loss should be used for model training. Also, instead of fixing the inter-task weights, we propose to learn the right trade-off parameters at each step using GradNorm [1].

3 Problem setting and basic concepts

We assume a set of tasks $T = \{1, \dots, T\}$ sharing an input space $X = U \times S$, where U is the subspace of *non-protected attributes* and S is the subspace of *protected attributes*. Each task t has its own label space Y_t . A dataset D of i.i.d. instances from the input space $X = U \times S$ and task spaces $\{Y_t\}_{t=1}^T$ is given as: $D = \{(u_i, s_i, y_i^1, \dots, y_i^T)\}_{i=1}^n$ where (u_i, s_i) is the description of instance i and $(y_i^1 \dots y_i^T)$ are the associated class labels for tasks $1 \dots T$. For simplicity, we assume binary tasks, i.e., $\forall t \in T : Y_t \in \{1, 0\}$, with 1 representing a positive (e.g., “granted”) and 0 representing a negative (e.g., “rejected”) class. We also assume the protected subspace S to be a binary protected attribute: $S \in \{g, \bar{g}\}$, where g and \bar{g} represent the *protected* and *non-protected group*, respectively.

3.1 Fairness Definition and Metric

As our fairness measure, we use **Equalized odds** (EO) [10], introduced for STL. For a task t , EO states that a classifier’s prediction \hat{Y}_t conditioned on ground truth Y_t must be independent from the protected attribute S . Based on [10], fairness is preserved when: $P(\hat{Y}_t = 1|S = g, Y_t = y) = P(\hat{Y}_t = 1|S = \bar{g}, Y_t = y)$, where $y \in \{1, 0\}$. A classifier satisfies the EO definition if the protected and non-protected groups have equal true positive rate (TPR) and false positive rate (FPR). For a task t , the FPR w.r.t. the protected group g is given by: $FPR_t(g): P(\hat{Y}_t = 1|Y_t = 0, S = g)$, whereas the TPR is given by: $TPR_t(g): P(\hat{Y}_t = 1|Y_t = 1, S = g)$. Similarly, we can define $FPR_t(\bar{g})$ and $TPR_t(\bar{g})$ for the non-protected group \bar{g} . The (absolute) differences in TPR and FPR define the violation of EO w.r.t. S , denoted by EO_{viol} . Given $TPR = 1 - FNR$, the violation can be also expressed in terms of FNR and FPR differences:

$$\begin{aligned} EO_{viol} &= \left| TPR_t(g) - TPR_t(\bar{g}) \right| + \left| FPR_t(g) - FPR_t(\bar{g}) \right| \\ &= \left| FNR_t(g) - FNR_t(\bar{g}) \right| + \left| FPR_t(g) - FPR_t(\bar{g}) \right| \end{aligned} \quad (1)$$

3.2 Vanilla Multi-task Learning (MTL)

Let \mathcal{M} be an MTL model parameterized by the set of parameters $\theta \in \Theta$, which includes: shared parameters θ_{sh} (i.e., weights of layers shared by all tasks) and task-specific parameters θ_t (i.e., weights of the task specific layers), i.e.: $\theta = \theta_{sh} \times \theta_1 \times \dots \times \theta_T$. An overview is given in Figure 1.

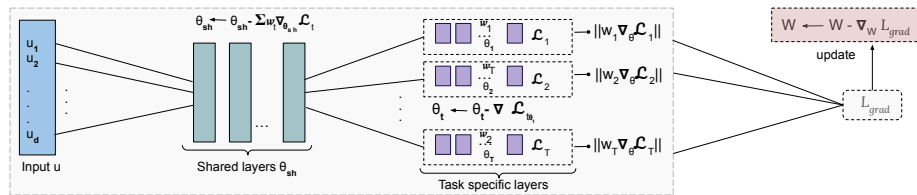


Fig. 1: An overview of (vanilla) MTL with GradNorm update

The goal of vanilla MTL training is to jointly minimise multiple loss functions, one for each task: $\operatorname{argmin}_{\theta} (\mathcal{L}_1(\theta, U), \dots, \mathcal{L}_T(\theta, U))$. Finding a model θ that minimizes all T tasks simultaneously is hard and typically, a scalarization approach is followed [1, 9, 12] which turns the multi-tasking function into a single function using inter-task trade-off weights w_t , as follows:

$$\operatorname{argmin}_{\theta} \sum_t w_t \mathcal{L}_t(\theta, U) \quad (2)$$

The trade-off weights w_t signify the importance of each task; usually selected by hyperparameter tuning over a validation set. This, however, means *fixing* the weights over the whole training process. A better approach, which we follow in this work, is to vary the weights at each training step to balance the contribution of each task for optimal model training. In this direction, GradNorm [1] proposes to balance the training rates of different tasks; if a task is training relatively quickly, then its weight should decrease relative to other tasks' weights to allow other tasks to influence training. GradNorm is implemented as an $L1$ loss function between the actual and target gradient norms at each training step for each task, summed over all tasks:

$$L_{grad} = \sum_t |G_t(\theta, U) - \overline{G(\theta, U)} \times \rho_t^\alpha| \quad (3)$$

where $G_t(\theta, U)$ is the L2 norm of the gradient of the loss for task t w.r.t. the chosen weights θ , defined as $G_t(\theta, U) = \|\nabla_\theta w_t \mathcal{L}_t(\theta, U)\|_2$; $\overline{G(\theta, U)}$ is the average gradient norm. Finally, ρ_t is the inverse training rate for task t and α is a hyperparameter of strength to pull any task back to the average training rate.

3.3 Fairness-aware Multi-task Learning (FMTL)

Fairness has been extensively studied in the recent years for STL problems; most approaches combine accuracy and fairness losses into a single overall loss [19] as:

$$\operatorname{argmin}_\theta \left(\mathcal{L}(\theta, U) + \lambda \mathcal{F}(\theta, S) \right) \quad (4)$$

where $\mathcal{L}(\theta, U)$ is the typical *accuracy loss*, $\mathcal{F}(\theta, S)$ is the loss associated with the protected attribute S (referred to as *fairness loss*) and λ is a weight parameter that determines the fairness-accuracy trade-off.

Fairness-aware MTL extends traditional MTL and fair-STL setups by considering not only the predictive performance but also the fairness performance on the individual tasks. The fair-MTL problem was formulated [22] as optimizing a weighted sum of accuracy and fairness losses over all tasks:

$$\operatorname{argmin}_\theta \sum_t w_t \left(\mathcal{L}_t(\theta, U) + \lambda_t \mathcal{F}_t(\theta, S) \right) \quad (5)$$

where λ_t determines the fairness-accuracy (intra-task) trade-off within a task t and w_t determines the relative importance of task t (inter-task trade-off). Note that with this formulation the number of parameters required to be learned is doubled from T (for vanilla MTL, c.f., Eq. 2) to $2T$ (for FMTL c.f Eq. 5).

3.4 Deep Q-learning (DQN) and Multi-tasking DQN (MT-DQN)

Reinforcement Learning (RL) is based on learning via interaction with a single environment. At each step j , the environment observes a state $z_j \in Z$. The agent

takes an action $a_j \in A$ in the environment, causing a transition to a new state $z_{j+1} \in Z$. For the transition, the agent receives a reward $R(j) \in \mathbb{R}$. The goal of the agent is to learn a policy $\pi : A \times Z \rightarrow A$ that maximizes the expected future discounted reward. State-action values (Q-values) are often used as an estimator of the expected future return. When the state-action space is large, the Q-values are typically approximated via a DNN, known as Deep Q-Network (DQN). The input to the DQN are the states, whereas the outputs are the state-action values (or, Q-values). A DQN is trained with parameters θ^Q , to minimize the loss between the *predicted Q-values* and the *target future return*:

$$L_Q(Z, A, \theta^Q, R) = E \left[Q(z_j, \alpha_j | \theta^Q) - \left(R(j) + \gamma \max_{\alpha'_j} Q(z'_j, \alpha'_j | \theta^Q) \right) \right] \quad (6)$$

where $Q(z_j, \alpha_j | \theta^Q)$ is the predicted Q-value. The term $R(j) + \gamma \max_{\alpha'_j} Q(z'_j, \alpha'_j | \theta^Q)$ is the target value defined as the sum of the direct reward $R(j)$ for transitioning from state z_j to a successor state z'_j and of the Q-value of the best successor state [16] (as predicted by the DQN).

In *multi-task reinforcement learning*, a single agent must learn to master T different environments/task, so the environments are the different tasks. The DQN for approximating the Q-values of state-action pairs is now a multi-tasking deep network (c.f., Section 3.2) with parameters $\theta^Q = \theta_{sh}^Q \times \theta_1^Q \times \dots \times \theta_T^Q$, where θ_{sh}^Q are shared across all the task learning environments, and θ_t^Q are the task exclusive parameters for learning the optimal policy in the particular environment/task. The network should learn to predict the Q-values under different learning environments/tasks, so the new objective becomes:

$$\operatorname{argmin}_{\theta^Q} \sum_t w_t^Q L_Q(Z(t), A(t), \theta^Q, R(t)) \quad (7)$$

where w_t^Q is the learned weight for the environment/task E_t to overcome the challenge of inter-task loss dominance.

4 Learning to Teach Fairness-aware Multi-tasking

We first present the dynamic objective selection paradigm to formulate the fair-MTL problem (Section. 4.1). Then we introduce our *L2T-FMT* algorithm in Section 4.2, which is a teacher-student network framework. The student network (Section 4.3) learns a fair-MTL model, using advice from the teacher (Section 4.4) regarding the choice of the loss function.

4.1 Dynamic Loss Selection Formulation

The fair-MTL problem definition (cf. Eq. 5) introduces the λ_t parameters for combining accuracy and fairness losses for each task and therefore, the number of parameters required to be learned is doubled from T (for traditional MTL,

cf. Eq.2) to $2T$ (for fairness-aware MTL, cf. Eq. 5)). We propose to get rid of the λ dependency and rather *learn to choose at each training step* whether the accuracy loss (\mathcal{L}_t) or the fairness loss (\mathcal{F}_t) should be used to train the model on task t . Such an approach reduces the optimization problem of Eq. 5 to that of Eq. 2, which is well studied in the literature [1, 9, 12, 21]. Moreover, it also provides the flexibility to the learner to emphasise on the objective (accuracy or fairness) that is harder to learn for each task at each training step.

The transformed learning problem is formulated as:

$$\operatorname{argmin}_{\theta, m} \sum_t w_t(m) \begin{cases} \mathcal{L}_t(\theta, U, m), & \text{if task } t \text{ is trained for } \textit{accuracy} \text{ in step } m \\ \mathcal{F}_t(\theta, S, m), & \text{if task } t, \text{ is trained for } \textit{fairness} \text{ in step } m \end{cases} \quad (8)$$

where m is the step³ number that sequences until model convergence.

The decision about which loss to use is based on the teacher network, which is trained jointly with the student network (Section 4.2). The accuracy loss and fairness loss functions that we adopt in this work are provided hereafter.

Accuracy loss: We adopt the negative log-likelihood loss $\mathcal{L}_t(\theta, U)$ for a task t :

$$\mathcal{L}_t(\theta, U) = - \sum_i (y_i^t \log \mathcal{M}(u_i, \theta) + (1 - y_i^t) \log(1 - \mathcal{M}(u_i, \theta))) \quad (9)$$

Fairness loss: Several works discuss how to formulate a fairness loss function [8, 7, 4, 17, 20]. To keep the characteristic of the fairness loss function similar to the one used for the accuracy loss, we adopt the robust log-loss [20, 19], which focuses on the worst-case log loss and is given by:

$$\begin{aligned} \mathcal{F}_t(\theta, U, S) = & \max(\mathcal{FN}\mathcal{R}_t^g(\theta, U, S), \mathcal{FN}\mathcal{R}_t^{\bar{g}}(\theta, U, S)) \\ & + \max(\mathcal{FPR}_t^g(\theta, U, S), \mathcal{FPR}_t^{\bar{g}}(\theta, U, S)) \end{aligned} \quad (10)$$

where the negative log-likelihood loss over a group (say g) for a task t :

$$\begin{aligned} \mathcal{FN}\mathcal{R}_t^g(\theta, U, S, Y_t) &= - \sum_i y_i^t \log \mathcal{M}(u_i, \theta | y_i^t = 1, s_i = g) \\ \mathcal{FPR}_t^g(\theta, U, S, Y_t) &= - \sum_i (1 - y_i^t) \log(1 - \mathcal{M}(u_i, \theta | y_i^t = 0, s_i = g)) \end{aligned} \quad (11)$$

4.2 L2T-FMT Algorithm

An overview of our *Learn to Teach Fairness-aware MTL (L2T-FMT)* method is shown in Figure 2. It consists of a student network and a teacher network which are trained collaboratively, as seen in Algorithm 1.

³ To note that the step information though specifically used in Eq. 8 to understand the temporal aspect of the selection, is valid for all the previous and following equations and is omitted for the rest of the paper for ease of reading.

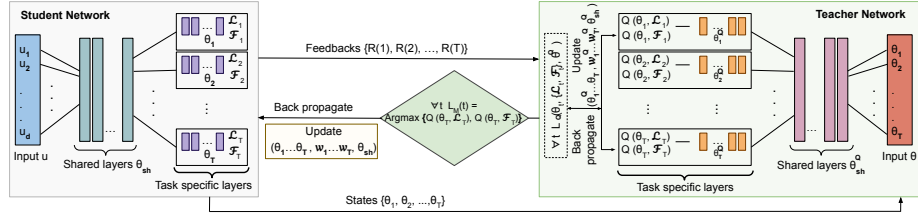


Fig. 2: The L2T-FMT architecture. Steps are in order from bottom to top

The *teacher network* \mathcal{Q} aims at each training round to select (line 2) the best function for each task among the two available options: accuracy loss and fairness loss. The selected loss function is adopted by the student network \mathcal{M} and is used for network update (line 3). After the update, the *student network* provides feedback (line 4) to the *teacher network* about the progress made in each of the tasks following the advice of the *teacher*. Based on the feedback, the *teacher network* updates itself (line 5). So the two networks are trained collaboratively; their functionalities are explained in the following sections.

Algorithm 1: The L2T-FMT algorithm

Input: A MTL dataset $D \in U \times S \times Y$ (see Section 3)

Init: Initialize student model \mathcal{M} parameters: $\theta = \theta_{sh} \times \theta_1 \times \dots \times \theta_T$
and task weights: $w_t = \frac{1}{T}$

Initialise teacher model \mathcal{Q} : parameters- $\theta^Q = \theta_{sh}^Q \times \theta_1^Q \times \dots \times \theta_T^Q$ and
environment weights $w^Q = \{w_t^Q = \frac{1}{T}\}$

1: **Until** convergence **do**

2: $d = \{ \operatorname{argmax}_{\alpha \in \{\mathcal{L}_t, \mathcal{F}_t\}} Q(\theta_t, \alpha), \forall t \}$ // **teacher decides the best loss to learn**

for student upon seeing θ_t

3: Call Algorithm 2 with input (U, S, Y, d) // **Training \mathcal{M}**

4: Evaluate $R = \{[\mathcal{R}(\mathcal{L}_t), \mathcal{R}(\mathcal{F}_t)] \forall t\}$ using Eq. 12 // **generate feedback for the teacher**

5: Call Algorithm 3 with input $(\{\mathcal{M}(t)\}, \{\theta_t\}, R, A)$ // **Training \mathcal{Q}**

Output: \mathcal{M} with learned parameters $\theta = \theta_{sh} \times \theta_1 \times \dots \times \theta_T$

4.3 Student Network

The student network \mathcal{M} is a deep multi-tasking neural network with learning parameters $\theta = \theta_{sh} \times \theta_1 \times \dots \times \theta_T$, as described in Sec. 3.2. It aims to learn to solve the fairness-aware multi-tasking problem by optimizing Eq. 8. The pseudocode is shown in Algorithm 2. For each task t , the decision of the *teacher* \mathcal{Q} (given

Algorithm 2: Student_MTL

Input: $U, S, Y = \{Y_t\}_1^T$, d teacher’s decision for the loss (line 2, Algo. 1)

- 1: **for** $1, \dots, T$ **do**
- 2: **if** $d(t) = \mathcal{L}_t$ **then** $L_{\mathcal{M}}(t) = \mathcal{L}_t(\theta, U)$ per task (as in Eq. 9) // Compute accuracy loss
- 3: **else** $L_{\mathcal{M}}(t) = \mathcal{F}_t(\theta, U, S)$ per task (as in Eq. 10) // Compute fairness loss
- 4: $\theta_t \leftarrow \theta_t - \eta \nabla_{\theta_t} L_{\mathcal{M}}(t)$ // Update task-specific layers params
- 5: $G_t = \|\nabla_{\theta} w_t L_{\mathcal{M}}(t)\|_2$, and $\rho_t = \frac{L_{\mathcal{M}}(t)}{E(L_{\mathcal{M}}(t))}$ // calculate gradient norm, and inverse training rate
- 6: **end for**
- 7: Compute L_{grad} as in Eq. 3 and update $W \forall_{w_t \in W} w_t \leftarrow w_t - \eta \nabla_{w_t} L_{grad}$
- 8: $\theta_{sh} \leftarrow \theta_{sh} - \eta \sum_t w_t \nabla_{\theta_{sh}} L_{\mathcal{M}}(t)$; // update the shared parameters

as input) about which action/loss to use is followed. Based on the decision, the selected loss (accuracy or fairness) is computed (lines 2-3) and used for the update of task-specific network layers (line 4). The task-specific weights w_t are learned using GradNorm (lines 5 and 7). Finally, the shared parameters are updated (line 8) using the updated task-specific weights and the loss function decided by the teacher network. The weight update mechanism using GradNorm (c.f., Section 3.2) is visualized in Fig 1 for reference.

4.4 Teacher Network

The teacher network \mathcal{Q} is a MT-DQN agent, as described in Sec. 3.4. It aims to learn to decide about which loss function, among the accuracy and fairness, the student network \mathcal{M} should use. In particular, \mathcal{Q} learns to predict the Q-values of accuracy-, fairness-loss/actions; the actual decision is the action with the largest Q-value (line 2 of Algorithm 1).

The pseudocode of \mathcal{Q} is shown in Algorithm 3. For each task t , it estimates the Q-values of the accuracy- and fairness-loss functions/actions based on the current model parameters (line 2). The network is updated (lines 3-9) as described in Section 3.4. Teacher’s training depends on the feedback by the student network, in the form of reward, which is computed by evaluating the student’s output in terms of Acc (accuracy) and EO_{viol} (fairness) (line 4, Algorithm 1). The main intuition for the design of the reward function is to reward positively only if Acc increases and simultaneously EO_{viol} decreases for any action suggested by \mathcal{Q} for a task t . On violation of either of the two conditions, the reward should be negative (the min function ensures this positive/negative property). However, the problem of intra-task dominance arises in estimation of \mathcal{R} as the scales of evaluated *accuracy* and *fairness* for task t might differ. To calculate a scale-invariant reward, we take inspiration from [11], and estimate $\mathcal{R}(\alpha_t) \in \mathbb{R}$ for

Algorithm 3: Teacher_DQN

Input: $\{\mathcal{M}(t)\}, \{\theta_t\}, \{R(t) = [\mathcal{R}(\mathcal{L}_t), \mathcal{R}(\mathcal{F}_t)]\},$
 $\{A(t)|A(t) = \{\mathcal{L}_t(\theta, U), \mathcal{F}_t(\theta, S)\}\},$

- 1: **for** $1, \dots, T$ **do**
- 2: Estimate $Q(\theta_t, \mathcal{L}_t), Q(\theta_t, \mathcal{F}_t)$ // estimating Q-values for each state-action pair
- 3: $\forall \alpha'_t \in A(t)$ take action α'_t in $M(t)$ to produce θ'_t // teacher takes a one look forward into the student environment
- 4: Compute the loss $L_Q(\theta_t, A(t), \theta^Q, R(t))$ (as in Eq. 6)
- 5: $\theta_t^Q \leftarrow \theta_t^Q - \eta \nabla_{\theta_t^Q} L_Q(\theta_t, A(t), \theta^Q)$ // teacher updates the parameters of the task environment specific layers
- 6: $G_t^Q = \|\nabla_{\theta_t^Q} L_Q(\theta_t, A(t), \theta^Q)\|_2$, and $\rho_t^Q = \frac{L_Q(\theta_t, A(t), \theta^Q)}{E(L_Q(\theta_t, A(t), \theta^Q))}$ // calculate gradients norm, and inverse training rate
- 7: **end for**
- 8: Compute L_{grad}^Q as in Eq. 3 and update $w_t^Q \forall w_t^Q \in w^Q w_t^Q \leftarrow w_t^Q - \eta \nabla_{w_t^Q} L_{grad}^Q$
- 9: $\theta_{sh}^Q \leftarrow \theta_{sh}^Q - \eta \sum_t w_t^Q \nabla_{\theta_{sh}^Q} L_Q(\theta_t, A(t), \theta^Q)$ // update the shared parameters

$\alpha_t \in \{\mathcal{L}_t, \mathcal{F}_t\}$, as the transformed evaluated output:

$$\mathcal{R}(\alpha_t) = \min\left(\frac{Acc(\hat{Y}_t) - Acc^{best}(t)}{Acc^{best}(t)}, \frac{EO_{viol}^{best}(t) - EO_{viol}(\hat{Y}_t)}{1 - EO_{viol}^{best}(t)}\right) \quad (12)$$

where $Acc(\hat{Y}_t)$, $Acc^{best}(t)$, and $EO_{viol}(\hat{Y}_t)$, $EO_{viol}^{best}(t)$ are respectively the current and best till current step *accuracy*, and *fairness* values of \mathcal{M} in task t .

5 Experiments

We first evaluate the accuracy-and fairness of our $L2T-FMT^4$ in comparison to other approaches for different MTL problems (Section 5.2) including a more task-specific evaluation (Section 5.3). Next, we analyse the impact of the dynamic loss function selection by $L2T-FMT$ (Section 5.4). The experimental setup including datasets, evaluation measures, and competitors is discussed in Section 5.1.

5.1 Experimental Setup

Datasets We evaluate on one tabular and two visual datasets. The tabular dataset is the recently released *ACS-PUMS* [2], which comprises a superset of the popular Adult dataset from available US Census sources, and consists of 5 different well defined binary classification tasks⁵. We use *gender* as the protected attribute. For training we use the census data from the year “2018”, divided into

⁴ <https://github.com/arjunroyihpa/L2TFMT>

⁵ <https://github.com/zykls/folktables>

train (70%) and validation (30%) sets. For testing we use the data from the following year “2019” (both years of size $\approx 1.65M$). The visual datasets come from CelebA dataset [14] consisting of 202.5K celebrity facial images and 40 different binary attributes. We use the provided⁶ partitioning into train (#162,770 instances), validation (#19,867 instances), and test (#19,962 instances) set. We use two different protected attributes, gender and age, resulting into two versions of the dataset, *CelebA-Gender* and *CelebA-Age*. We don’t consider all 40 attributes as tasks for the MTL since some attributes are extremely skewed towards the protected or non-protected group. For example, the attribute “Mustache” is true only for 3 female instances. We set the filtering threshold to 1.5% or 2.5K instances. The filtering process reduces the number of attributes to 17 for *CelebA-Gender* and 31 for *CelebA-Age*; these are the MTL tasks. Further details on the datasets and experimental setup are provided in the Appendix.

Methods We compared *L2T-FMT* against the following methods:

i) MTA-F: the vanilla fairness-aware MTL method [22] that minimises the weighted sum of accuracy- and fairness-losses (c.f. Eq. 5). For fairness it calculates two separate loss, one for updating θ_t and another for updating θ_{sh} . The weights w_t and λ_t are set via hyperparameter tuning.

ii) G-FMT: a variation of our *L2T-FMT* approach that always chooses *greedily* the best action/loss function among the available choices, by optimising: $\operatorname{argmin}_{\theta} \sum_t w_t \{\max\{\mathcal{L}_t(\theta, U), \mathcal{F}_t(\theta, S)\}\}$.

iii) Vanilla MTL: the vanilla MTL approach that does not consider fairness but aims at minimising the weighted sum of task-specific accuracy losses (c.f. Eq. 2). The task-specific weights w_t are learned via GradNorm [1] as in Eq 3.

iv) STL: trains a separate fair-accurate model on each respective task.

Evaluation Measures Following [22], we report on the relative performance of the MTL model ($Acc(t)_{mtl}$, $EO_{viol}(t)_{mtl}$) to the performance of a STL model trained on each respective task t ($Acc(t)_{stl}$, $EO_{viol}(t)_{stl}$). Specifically, for accuracy we report on the *average relative Acc (ARA)*: $ARA = \frac{1}{T} \sum_t \frac{Acc(t)_{mtl}}{Acc(t)_{stl}}$ and for fairness, on the *average relative EO (AREO)*: $AREO = \frac{1}{T} \sum_t \frac{EO_{viol}(t)_{mtl}}{EO_{viol}(t)_{stl}}$.

5.2 Overall fairness-accuracy evaluation

The overall fairness (AREO) and accuracy (ARA) performance of the different methods on all the datasets is shown in Table 1. As we can see, *L2T-FMT* outperforms all the competitors in *fairness* by producing the lowest AREO scores across all the datasets; the relative reduction in discrimination w.r.t. the best baseline is in the range [12% – 19%]. Interestingly, the second best approach in terms of AREO is our greedy variation, *G-MFT*. In terms of ARA, *L2T-FMT* is best by a small margin comparing to the best baseline with the exception of the ACS-PUMPS dataset for which Vanilla-MTL scores first. In particular, for *CelebA-Gender* and *CelebA-Age*, our *L2T-FMT* beats the best baseline by 2% and 1%, respectively, whereas for the *ACS-PUMPS* dataset *L2T-FMT* scores second with a 3% decrease comparing to the best performing *Vanilla MTL*.

⁶ <http://mmlab.ie.cuhk.edu.hk/projects/CelebA.html>

To get better insights on the results, in the next section we also report on the task-specific performance using accuracy and EO_{viol} for each task.

Table 1: ARA vs AREO: Higher values better for accuracy/ARA, lower values better for discrimination/AREO . Best values in bold, second best underlined. (%) indicates our relative difference over the performance of the best baseline.

Dataset	#tasks T	Metric	Vanilla MTL	MTA-F	G-FMT	L2T-FMT
ACS-PUMS	5	ARA	1.06	0.97	1.01	1.03 (-3%)
		AREO	2.38	3.52	<u>1.50</u>	1.21 (-19%)
CelebA-Gender	17	ARA	0.89	<u>0.95</u>	0.86	0.97 (2%)
		AREO	2.72	2.29	<u>1.77</u>	1.51 (-15%)
CelebA-Age	31	ARA	<u>0.94</u>	0.85	0.86	0.95 (1%)
		AREO	2.61	1.79	<u>1.72</u>	1.52 (-12%)

5.3 Performance distribution over the tasks

We analyze the distribution of accuracy and fairness scores over the tasks for all methods using boxplots. The results are shown in Fig. 3, 4, and 5 for the *ACS-PUMS*, *CelebA-Gender* and *CelebA-Age*, respectively. For fairness a positively skewed box (median closer to Q1) with low Q1 is better, while for accuracy a negatively skewed box (median closer to Q3) with high Q3 is better.

ACS-PUMS dataset. In Fig 3a, we see that *L2T-FMT* has the lowest median, and the lowest $Q1$ of EO_{viol} , with $Q3$ marginally above the *STLs* but lower than all the competitors. Henceforth, it achieves the best AREO score (c.f., Table 1). In Fig 3b, we see that *L2T-FMT* has the second highest median after *G-FMT*, however it has a higher $Q1$ than *G-FMT* but a lower $Q1$ and lower $Q3$ than *vanilla MTL*. Thus, in Table 1 we find *L2T-FMT* to be second best in ARA score behind *vanilla MTL* on ACS-PUMS. *MTA-F* has the most consistent outcome of EO_{viol} with low spread over the tasks for both fairness and accuracy, but has the highest median and $Q1$ of EO_{viol} , and the lowest median and lowest $Q3$ of accuracy. Thus, in overall it has the worst overall performance as also seen in Table 1. *G-FMT* has the highest spread in both fairness and accuracy and is negatively skewed in accuracy with high accuracy for some tasks. *Vanilla-MTL* comes second in terms of spread and is positively skewed in EO_{viol} . However, its upper whisker is longer with a single point above Q3; this corresponds to task 1 (*Employment Status*) for which the EO_{viol} score is high.

CelebA-Gender dataset. In Fig 4a, we see that all methods have low spread indicating their consistent performance w.r.t. fairness across the tasks. Still, *L2T-FMT* outperforms the competitors with the lowest median, $Q1$, and $Q3$ values. *MTA-F* has the third best median, $Q1$, and $Q3$ of EO_{viol} . As we see in Fig 4a it has a much longer upper whisker indicating larger variance among the larger values, i.e., tasks with higher discrimination with the worst

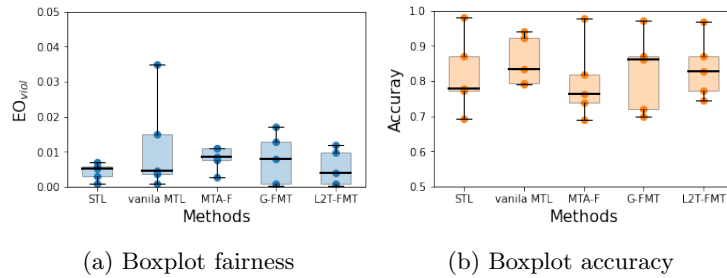


Fig. 3: ACS-PUMS dataset: Performance distribution over the tasks.

discrimination of 0.175 which corresponds to task 11 (*Narrow Eyes*). However, in accuracy all the methods vary substantially (high spread) as we see in Fig 4b. *L2T-FMT* holds the highest median and $Q3$, but its $Q1$ is marginally lower than *MTA-F* which indicates that in a few tasks *L2T-FMT* gets outperformed by *MTA-F* (second best median, $Q3$) on accuracy. This explains the ARA scores in Table 1, where *L2T-FMT* scores first followed by *MTA-F*. *G-FMT* has the second best median, $Q3$, and $Q1$ of EO_{viol} score, but has the worst median, $Q3$, and $Q1$ of accuracy. Thus, in Table 1 we see that *G-FMT* bags the second best AREO score but has the worst ARA score.

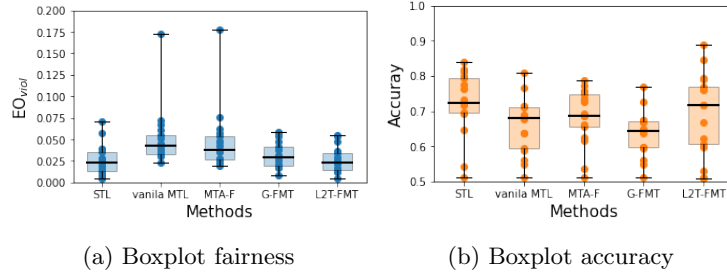


Fig. 4: CelebA-Gender: Performance distribution over the tasks.

CelebA-Age dataset. There is large spread accuracy (Fig 5b) across the different methods. *L2T-FMT* has the lowest median, $Q3$, and $Q1$ of EO_{viol} score, and highest median, $Q3$, and $Q1$ of accuracy, respectively. This reflects in Table 1 where *L2T-FMT* achieves the best AREO, and ARA scores. *MTA-F* delivers consistent fair performance over the tasks with the second best median, $Q3$, and $Q1$ of EO_{viol} score almost same as *G-FMT*. This is the reason why the AREO score of *G-FMT* and *MTA-F* are nearly same, with *G-FMT* marginally ahead, making *MTA-F* the third best on AREO score. On accuracy *G-FMT* has the third best median, $Q3$, and $Q1$, having marginal improvements over *MTA-F*. *Vanilla MTL* with no fairness treatment has the worst median, $Q3$, and $Q1$ of

of EO_{viol} score over the tasks positioning it at the last place on AREO score, while having median, $Q3$, and $Q1$ of accuracy over the tasks very similar to that of $L2T-FMT$ making it a very close second on ARA score.

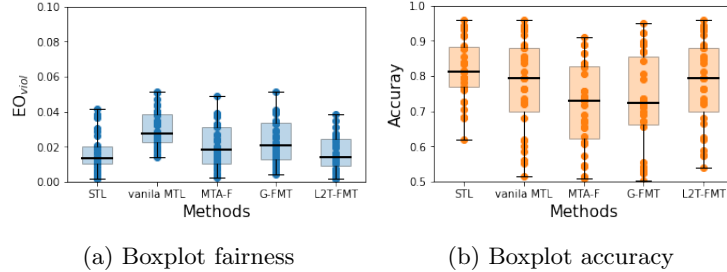


Fig. 5: CelebA-Age: Performance distribution over the tasks.

Summary. For all datasets, $L2T-FMT$ stands out among the competitors with a very low median of EO_{viol} , and a very high median of accuracy. The fairness performance of $MTA-F$ depends on the number of tasks; for larger MTL problems (like *CelebA-Gender* and *CelebA-Race*) the performance varies across the tasks including tasks with high discrimination (high upper whisker). $G-FMT$ consistently delivers descent fairness performance positioning it always in the second place, however its accuracy gets affected when the number of tasks is high (*CelebA-Gender* and *CelebA-Age*). The *Vanilla MTL* without any fairness treatment on EO_{viol} score is often bad with very high upper whiskers.

5.4 Dynamic loss selection

We focus on the (dynamic) loss selection of the teacher network by looking at which function among the two available options: accuracy loss (\mathcal{L}) and fairness loss (\mathcal{F}) is used for each task over the training process. The results for the different datasets are shown in Figure 6.

Regarding *ACS-PUMS* (Fig. 6a), the decision of selecting \mathcal{L} or \mathcal{F} is almost equally distributed over the tasks. Using *Vanilla MTL* (trained only for accuracy) as a reference method, we see in Table 1 that to achieve the best accuracy (best ARA) in this dataset one needs to always tune with \mathcal{L} . However, this comes with depreciation in fairness (high AREO). Thus, the necessary balance as selected by $L2T-FMT$ is required.

For *CelebA-Gender* (Fig. 6b), the accuracy loss \mathcal{L} is selected more often in some tasks (2, 10, 17) and the fairness loss \mathcal{F} in other tasks (6, 9, 11, 15), whereas there are also tasks with balanced selection (e.g., 1). *Vanilla MTL* that only tunes for \mathcal{L} does not produce the best accuracy (c.f., Table 1); in contrast, $L2T-FMT$ with dynamic loss selection achieves the best accuracy and fairness, justifying the selection (im-) balance.

For *CelebA-Age*, (Fig. 6c) we notice that for the majority of the tasks (21 out of 31) the loss selection is unevenly distributed with 13 tasks in favour of \mathcal{L} , and 8 in favour of \mathcal{F} . Interestingly, in 9 tasks (1, 5, 6, 10, 15, 18, 23, 25, 29) \mathcal{L} is selected continuously over epochs at a stretch (≥ 7), signifying that in these tasks tuning for accuracy is far more important. This is also reflected in the very close ARA performance of *L2T-FMT* and *Vanilla MTL* (c.f., Table 1). Similarly, in 8 tasks (7, 14, 16, 17, 24, 26, 30, 31) \mathcal{F} is chosen more frequently, which ultimately leads to the superiority of *L2T-FMT* in the AREO score.

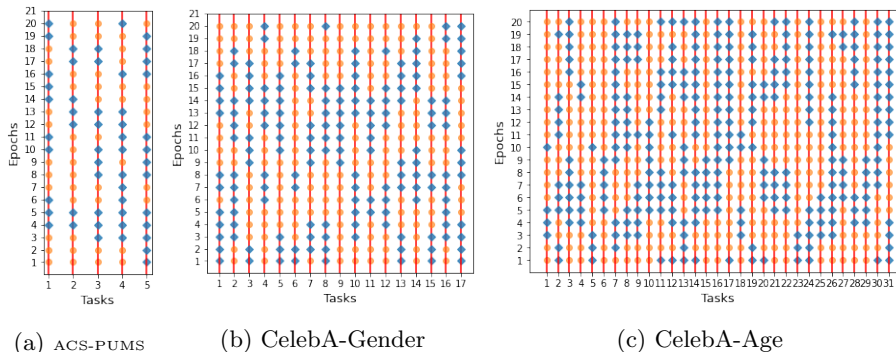


Fig. 6: Loss function selection for each task over the training epochs. Accuracy-loss (\mathcal{L}) selection in orange, fairness-loss (\mathcal{F}) selection in blue.

6 Conclusion

We proposed *L2T-FMT*, an approach for fairness-aware multi-task learning that dynamically selects for each task the best loss function to be used at each training step among the available: accuracy loss and fairness loss. Our approach is a student-teacher network framework, where the student learns to solve the fair-MTL problem while the teacher decides the action/loss function that the student should use for its update. The teacher is implemented as a DQN, whereas the student is implemented as a deep MTL network. In contrast to a rigid fairness-accuracy trade-off formulation [22], *L2T-FMT* allows for a flexible model update based on which objective (accuracy or fairness) is harder to learn for each task. Moreover, it reduces the number of parameters to be learned to half. Our experiments on three real datasets show that *L2T-FMT* improves on both fairness and accuracy over state-of-the-art approaches. Moreover, we also show the effectiveness of learning to select the best loss in producing such favourable outcomes.

As part of our future work, instead of jointly training on all the tasks, we plan to identify the tasks that would benefit from training together. Such task groupings might differ based on whether fairness or accuracy is considered.

Acknowledgements. The work of A.Roy is supported by Volkswagen Foundation under the call “Artificial Intelligence and the Society of the Future” project “Bias and Discrimination in Big Data and Algorithmic Processing - BIAS”.

References

1. Chen, Z., Badrinarayanan, V., Lee, C.Y., Rabinovich, A.: Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In: ICML. pp. 794–803. PMLR (2018)
2. Ding, F., Hardt, M., Miller, J., Schmidt, L.: Retiring adult: New datasets for fair machine learning. *NeurIPS* **34** (2021)
3. Dong, N., Kampffmeyer, M., Voiculescu, I.: Self-supervised multi-task representation learning for sequential medical images. In: ECML PKDD. pp. 779–794. Springer (2021)
4. Donini, M., Oneto, L., Ben-David, S., Shawe-Taylor, J., Pontil, M.: Empirical risk minimization under fairness constraints. In: *NeurIPS*. pp. 2796–2806 (2018)
5. Dwork, C., Hardt, M., Pitassi, T., Reingold, O., Zemel, R.: Fairness through awareness. In: ITCS. pp. 214–226 (2012)
6. Fan, Y., Tian, F., Qin, T., Li, X.Y., Liu, T.Y.: Learning to teach. In: ICLR (2018)
7. Feldman, M., Friedler, S.A., Moeller, J., Scheidegger, C., Venkatasubramanian, S.: Certifying and removing disparate impact. In: KDD. pp. 259–268 (2015)
8. Gretton, A., Borgwardt, K., Rasch, M., Schölkopf, B., Smola, A.: A kernel method for the two-sample-problem. *NeurIPS* **19**, 513–520 (2006)
9. Guo, P., Deng, C., Xu, L., Huang, X., Zhang, Y.: Deep multi-task augmented feature learning via hierarchical graph neural network. In: ECML PKDD. pp. 538–553. Springer (2021)
10. Hardt, M., Price, E., Srebro, N.: Equality of opportunity in supervised learning. *NeurIPS* **29**, 3315–3323 (2016)
11. Hessel, M., Soyer, H., Espenholt, L., Czarnecki, W., Schmitt, S., van Hasselt, H.: Multi-task deep reinforcement learning with popart. In: AAI. vol. 33, pp. 3796–3803 (2019)
12. Kacupaj, E., Premnadh, S., Singh, K., Lehmann, J., Maleshkova, M.: Vogue: Answer verbalization through multi-task learning. In: ECML PKDD. pp. 563–579. Springer (2021)
13. Kendall, A., Gal, Y., Cipolla, R.: Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7482–7491 (2018)
14. Liu, Z., Luo, P., Wang, X., Tang, X.: Deep learning face attributes in the wild. In: ICCV (December 2015)
15. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., Galstyan, A.: A survey on bias and fairness in machine learning. *CSUR* **54**(6), 1–35 (2021)
16. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G., et al.: Human-level control through deep reinforcement learning. *Nature* **518**(7540), 529–533 (2015)
17. Oneto, L., Donini, M., Pontil, M.: General fair empirical risk minimization. In: IJCNN. pp. 1–8. IEEE (2020)
18. Oneto, L., Doninini, M., Elders, A., Pontil, M.: Taking advantage of multitask learning for fair classification. In: AIES. pp. 227–237 (2019)

19. Padala, M., Gujar, S.: Fmnc: achieving fairness through neural networks. In: IJCAI (2020)
20. Rezaei, A., Fathony, R., Memarrast, O., Ziebart, B.: Fairness for robust log loss classification. In: AAAI. vol. 34, pp. 5511–5518 (2020)
21. Vandenhende, S., Georgoulis, S., Van Gansbeke, W., Proesmans, M., Dai, D., Van Gool, L.: Multi-task learning for dense prediction tasks: A survey. TPAMI (2021)
22. Wang, Y., Wang, X., Beutel, A., Prost, F., Chen, J., Chi, E.H.: Understanding and improving fairness-accuracy trade-offs in multi-task learning. In: KDD. pp. 1748–1757 (2021)
23. Wu, L., Tian, F., Xia, Y., Fan, Y., Qin, T., Lai, J.H., Liu, T.Y.: Learning to teach with dynamic loss functions. In: NeurIPS (2018)
24. Zhao, C., Chen, F.: Rank-based multi-task learning for fair regression. In: ICDM. pp. 916–925. IEEE (2019)