

AdaMTL: Adaptive Input-dependent Inference for Efficient Multi-Task Learning

Marina Neseem*, Ahmed Agiza*, Sherief Reda
Brown University
Providence, RI

{marina_neseem, ahmed_agiza, sherief_reda}@brown.edu

Abstract

Modern Augmented reality applications require performing multiple tasks on each input frame simultaneously. Multi-task learning (MTL) represents an effective approach where multiple tasks share an encoder to extract representative features from the input frame, followed by task-specific decoders to generate predictions for each task. Generally, the shared encoder in MTL models needs to have a large representational capacity in order to generalize well to various tasks and input data, which has a negative effect on the inference latency. In this paper, we argue that due to the large variations in the complexity of the input frames, some computations might be unnecessary for the output. Therefore, we introduce AdaMTL, an adaptive framework that learns task-aware inference policies for the MTL models in an input-dependent manner. Specifically, we attach a task-aware lightweight policy network to the shared encoder and co-train it alongside the MTL model to recognize unnecessary computations. During runtime, our task-aware policy network decides which parts of the model to activate depending on the input frame and the target computational complexity. Extensive experiments on the PASCAL dataset demonstrate that AdaMTL reduces the computational complexity by 43% while improving the accuracy by 1.32% compared to single-task models. Combined with SOTA MTL methodologies, AdaMTL boosts the accuracy by 7.8% while improving the efficiency by 3.1 \times . When deployed on Vuzix M4000 smart glasses, AdaMTL reduces the inference latency and the energy consumption by up to 21.8% and 37.5%, respectively, compared to the static MTL model. Our code is publicly available ¹.

1. Introduction

Modern computer-vision-based applications require solving multiple tasks simultaneously in order to form a

*The first two authors contributed equally to this work.

¹<https://github.com/scale-lab/AdaMTL.git>

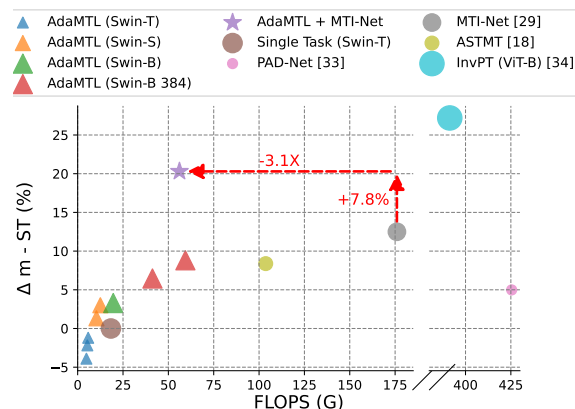


Figure 1. Accuracy-Efficiency trade-off by AdaMTL compared to SOTA MTL techniques. The x-axis shows the FLOPS, while the y-axis represents the average accuracy of the tasks compared to the single-task model.

complete perception of the surrounding visual environment. For example, a simple augmented reality application might need to determine the surface normals, estimate the depths, detect an object of interest, and track it. Similarly, an autonomous vehicle should be able to detect both static and dynamic objects in the scene, determine their proximity and track them [12, 21]. Moreover, all that complex processing needs to be performed on every input frame in real-time, which is extremely challenging, especially since those applications usually run on resource-constrained devices with strict compute, memory, and energy budgets. That is why enabling efficient computation models where these tasks can be performed simultaneously is crucial to make those applications practical.

In recent years, Multi-task learning (MTL) has been used to learn related vision tasks simultaneously [15, 24, 36]. On the one hand, MTL models leverage shared representations and inter-task interactions to improve performance on each task, potentially outperforming single-task models. On the other hand, compared to single-task models, MTL models can reduce both the memory footprint, the energy consumption, and the latency during inference since they avoid re-

computing the features in the shared layers. Different approaches have been proposed to determine which layers should be shared across tasks and which layers should be task-specific in MTL models [20, 26, 33].

One common MTL approach is to have a shared encoder that extracts the critical features from the input scene, followed by some task-specific decoders that predict the output of the corresponding task [28, 29, 33]. Generally, the shared encoder in MTL models needs to have a large representational capacity in order to generalize well to various tasks and input data from different complexities. Vision Transformers have proven to be a powerful tool for extracting strong feature representations from the inputs, improving the performance of the downstream tasks [6, 17, 27]. Hence, few recent works have explored using them as a shared encoder for feature extraction in MTL models [1, 34]. Although these approaches achieve impressive performance outcomes, their computational complexity and memory footprint are usually huge, making them impractical for real-time processing. That’s why, in this work, we focus our efforts on improving the computational complexity of transformer-based Multi-task Learning models.

Real-world visual scenes have large variations in complexity. For example, a scene with a single object and an open background would be easier to process than one with an occluded object and a cluttered background. We argue that treating all input frames equally regarding processing complexity can be wasteful. Therefore, we propose using an adaptive inference policy to reduce the computational complexity of MTL models based on the input complexity. Dynamic Input-dependent inference policies have only been explored for single-task image classification problems [19, 35]. However, MTL models are more complex due to inter-task dependencies and complex architectures.

To this end, we propose an adaptive MTL framework that recognizes the unnecessary computations in the model depending on the input complexity. We achieve this by learning an adaptive task-aware policy network that ultimately decides on which parts of the MTL model to activate during runtime. Our contributions can be summarized as follows:

- We propose an adaptive Multi-task Learning framework that optimizes Transformer-based MTL models depending on the input complexity.
- We introduce a task-aware policy network, and we show that it is more effective in recognizing the unnecessary computations in MTL models compared to task-agnostic policy networks.
- To prove that our policy network can be easily plugged to improve MTL models’ efficiency, we combine AdaMTL with a SOTA MTL model [29] and show that AdaMTL boosts the accuracy by 7.8% while improving the efficiency by $3.1\times$.

- We deploy AdaMTL on the Vuzix M4000 AR glasses [31], reducing the inference latency and the energy consumption by up to 21.8% and 37.5%, respectively, compared to the static MTL model.

The rest of the paper is organized as follows. We review the related work in Section 2. Then, we introduce our adaptive input-dependent MTL framework in Section 3. Next, we show the qualitative and quantitative analysis of our methodology as well as the ablation study in Section 4. Finally, we conclude in Section 5.

2. Background and Related Work

Deep Multi-task Architectures For dense prediction tasks, deep multi-task architectures usually consist of an encoder that extracts a feature representation from the input frame, followed by task-specific decoders that generate predictions for each of the downstream tasks [20, 29, 33]. Researchers categorize multi-task architectures based on the location of task interactions into encoder- and decoder-focused architectures [28]. Encoder-focused architectures share the information between tasks in the encoder stage [16, 20], while the decoder-focused architectures exchange information between tasks in the decoding stage [29, 33]. Moreover, some approaches share the information across tasks in both encoder and decoder stages [1, 26]. In this paper, our baseline MTL model follows an encoder-focused architecture where a shared encoder is used to extract visual features from the input frame, followed by task-specific decoders. Moreover, our MTL framework can easily adopt other decoder-based task-interaction techniques, as we will show in Subsection 4.3.

Vision Transformer Models Transformers [30] have achieved impressive performance improvements in various domains such as language understanding [3], speech recognition [7], and computer vision [4, 27]. The self-attention modules in transformers have proved to be capable of extracting strong feature representations from the inputs, improving the performance of the downstream tasks. To preserve local visual context, Vision Transformers (ViTs) split the input image into patches which are embedded as tokens [4]. Those tokens pass through successive ViT blocks. Each ViT block has a multi-head attention module followed by a multi-layer perceptron module to extract the global relationship among input tokens. Improvements have been made to ViTs enabling data-efficient training [27] as well as efficient inference [6]. Inspired by ResNets [9], Swin Transformers [17] use a hierarchical ViT block architecture as well as shifted window attention to serve as a general-purpose backbone for computer vision tasks. That is why, ViTs started replacing CNNs as a backbone for different computer vision tasks such as image classification [6, 27], object detection [2], and segmentation [25]. Moreover, several works proposed using them for multi-task learning [1, 34].

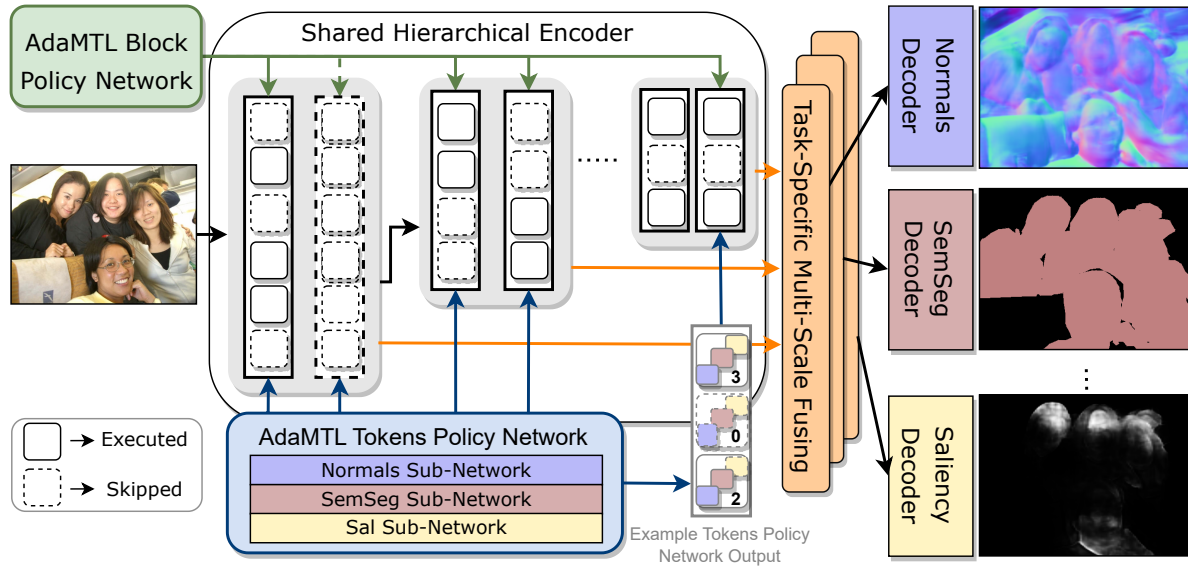


Figure 2. Overview of our proposed AdaMTL framework integrating our AdaMTL Block Policy Network and our AdaMTL Tokens Policy Network. The AdaMTL Block policy network decides on which blocks to activate during runtime. If the decision is to activate a certain block, our AdaMTL Tokens policy network runs to decide which tokens to process through that block. Our policy network achieves a task-aware behavior that improves the quality of the generated policies for multi-task learning models.

Sparsely-Activated Vision Models As computer vision models become increasingly complex, researchers have started sparsifying the models to prevent redundant computations. For example, DSelect-k [8] and M³ViT [14] proposed Mixture-of-Experts (MoE) architectures that use trainable sparse gates to activate a subset of experts depending on the given input. Other methods employ early-exiting strategies to adaptively allocate computations depending on input complexity [11, 37]. Other dynamic models skip redundant layers [32], while others use a bottleneck layer to direct the computations in an input-dependent manner [22]. More recent approaches use lightweight policy networks to generate execution strategies based on the input complexity [19, 35]. Despite the effectiveness of these approaches for image classification problems, adaptive policies have not yet been explored for MTL scenarios. Sparsifying MTL models is more challenging; the multi-objective nature of those models complicates the overall objective of optimizing the performance of the different tasks in addition to the policy networks. Therefore, there is a need to explore and develop techniques that can adaptively sparsify MTL models based on input complexity while being aware of the multi-objective nature of those models.

3. Methodology

In this section, we propose AdaMTL – an adaptive end-to-end Multi-task Learning framework. We start by presenting our transformer-based MTL architecture in Subsection 3.1. In Subsection 3.2, we introduce our task-aware policy network that dynamically recognizes the unnecessary com-

putations in the MTL model depending on the input complexity. Finally, we explain our proposed multi-staged task-aware training recipe in Subsection 3.3.

3.1. Overview

Figure 2 illustrates an overview of our proposed adaptive end-to-end Multi-task Learning framework - AdaMTL. AdaMTL consists of two main components: a static MTL model and a lightweight policy network. Our MTL model has three parts: a Shared Hierarchical Encoder, a learnable task-specific multi-scale fusing layer, and a pool of task-specific decoders. We adopt an off-the-shelf hierarchical Vision Transformer *Swin* [17] as our shared encoder to extract visual features from the input frames. Vision Transformers (ViTs) usually split the input image into patches which are embedded as tokens. Those tokens pass through successive ViT blocks. Each ViT block has a multi-head attention module followed by a multi-layer perceptron module to extract the global relationship among input tokens. Our learnable multi-scale fusing layers use a residual blocks-based architecture [9]. They are added to combine the features at different scales (i.e., receptive fields) in an informative way for every downstream task. Finally, we use simple task-specific decoders consisting of two convolutional layers. Each decoder takes the output from the corresponding multi-scale fusing layer to generate the corresponding task predictions.

On top of our static MTL model, AdaMTL adds a lightweight policy network that runs alongside the original model to decide which parts of the model to activate, adapt-

ing to the complexity of the input frame. Our policy network works as a multi-grained decision maker; it consists of an *AdaMTL Block policy network* and an *AdaMTL Tokens policy network*. The *AdaMTL Block policy network* decides on which blocks to activate during runtime. If the decision is to activate a certain block, our *AdaMTL Tokens policy network* runs to decide which tokens to process through that block. Our intuition behind AdaMTL is that not all patches in the input frame are equally informative; for example, a patch of an input frame from the background is less informative than a patch with a person for a task such as human-parts detection. Moreover, the receptive field at which different patches should be processed differs depending on the scale of the objects in the input frame. For example, an image with multiple smaller objects might benefit from being processed by the first few layers of the models where the receptive field is small. However, the later layers with larger receptive fields are essential for an accurate output on a zoomed-in frame with one object. In other words, our *AdaMTL Policy Network* decides which patches are needed to accurately perform the downstream tasks and the receptive field at which those patches should be processed.

3.2. AdaMTL Policy Network

Our policy network works as a multi-grained decision maker; the *Block Policy Network* decides on which blocks to activate, while the *Tokens Policy Network* decides on which tokens (i.e., patches) to process through the activated blocks. Our *Block Policy Network* generates a learnable binary mask for each block in the shared encoder. We use this binary mask to recognize the necessary blocks needed by the MTL model in order to perform well on the downstream tasks. Similarly, for each block, we attach a *Tokens Policy Network* that generates a learnable policy to determine the tokens that need to be processed through the activated block. Intuitively, the policy network should learn to recognize the most informative patches in every input frame as well as the receptive field at which it needs to be processed. Each policy network has two simple fully-connected layers, followed by a Gumbel Softmax activation to generate binary masks [13]. We devise two different settings for the policy network: a *task-agnostic* policy network and a *task-aware* policy network.

Task-agnostic Policy: In the *task-agnostic* setting, the policy network is unaware of the number of downstream tasks. We achieve this by co-training the whole policy network alongside the MTL model. We mainly experiment with this setting to show the necessity of task awareness while creating an effective policy network for multi-task scenarios.

Task-aware Policy: In the *task-aware* setting, we want our policy network to capture task-specific computational needs. We achieve this by dividing the policy network into sub-networks, where each sub-network is responsible

for recognizing the necessary blocks/tokens for the corresponding task. As shown in Figure 2, the *AdaMTL Tokens Policy Network* consists of a sub-network for each task (i.e., Normals Controller, Semantic Segmentation Controller, Saliency Controller, etc.). Each sub-network makes a decision that is plausible to its respective task. Finally, to get a unified policy, we make a decision to activate a token if at least one of the tasks needs to process it. The intuitive way to combine the masks would be to apply the logical ORing operation on all the generated task-specific masks. However, to make it more learnable, we combine the masks using addition and clamping. As shown in the example output from the Tokens Policy Network in Figure 2, the policy network only activates a token if at least one task-specific policy sub-network decides to activate it.

3.3. AdaMTL Training Recipe

For our adaptive MTL framework to work effectively, we need to learn the MTL model weights as well as the execution policy (i.e., policy network’s binary masks) that achieves the target efficiency without compromising the accuracy of the various tasks in our MTL model. Our end-to-end training recipe consists of 3 stages:

Stage 1: Static MTL Model Training

First, we train a static MTL model. We adopt a shared encoder along with task-specific decoders to perform multi-task learning as explained in Subsection 3.1. For our shared encoder, we use the publicly-available pre-trained Swin Transformer backbones [17]. Then, we attach the multi-scale fusing layer as well as task-specific decoders, and we fine-tune the end-to-end MTL model to get our static baseline. In this stage, our loss function is the weighted sum of the losses of the various downstream tasks as follows.

$$L_{stage1} = \sum_i^m \omega_{task.i} \times L_{task.i} \quad (1)$$

where $\omega_{task.i}$ and $L_{task.i}$ are the weight and the loss of the various tasks in the MTL model, respectively, and m is the number of downstream tasks in the MTL model. We adopt the task weights used by Vandenhende *et. al.* [29].

Stage-2 Policy Network Initialization

We aim to co-train both the policy network as well as the MTL model. Randomly initializing the policy network while co-training can lead to degrading the model accuracy since the policy network would make random decisions in the earlier epochs. That’s why we choose initialization weights for the policy network that activates all the blocks and the tokens. To achieve this, we freeze the static MTL model and pre-train our AdaMTL policy network with the following loss function:

$$L_{stage2} = \sum_k^{blocks} (1 - M_{b/k}) + \sum_k^{blocks} (1 - M_{t/k}) \quad (2)$$

Where $M_{b/k}$ and $M_{t/k}$ are the output masks generated by the Block Policy Network and the Tokens Policy Network attached to the Encoder block k , respectively. This results in an adaptive model that behaves exactly like the static model, where all blocks and tokens are activated. This acts as a good initialization point to start co-training the policy network and the MTL model.

Stage 3: Policy Network/MTL Model Co-training

In this stage, we co-train the policy network along with the MTL model. Our goal is to learn the binary masks that our policy network should generate in order to meet the target computational budget (i.e., the target percentage of the MTL model components to be activated) while maintaining the accuracy of the downstream tasks. Thus, we use a multi-objective loss as shown in Equation 3. Our loss function incorporates the various task losses L_{tasks} as well as the efficiency loss L_{eff} multiplied by a factor α . α represents the efficiency weight which controls the trade-off between accuracy and efficiency. In our experiments, we set α to unity; however, different values can be used to control the trade-off depending on the application requirements.

$$L_{stage3} = L_{tasks} + \alpha L_{eff} \quad (3)$$

The efficiency loss incorporates the efficiency of the decisions made by the blocks as well as the tokens controller, as shown in Equation 4. In order to minimize the number of activated blocks, we set L_{blocks} as mean squared error (MSE) between the actual percentage of the activated blocks and the target percentage of activated blocks as shown in Equation 5. Similarly, we set L_{tokens} as the MSE between the actual percentage of the activated tokens and the target percentage of activated tokens. However, it is common for hierarchical ViTs to have more tokens in the earlier layers compared to the later layers (i.e., the earlier layers have smaller receptive fields, thus more patches, while later layers have larger receptive fields, thus fewer patches). This means that the number of tokens does not linearly reflect the computational complexity since layers with more tokens have smaller embedding dimensions per token, while layers with fewer tokens have larger embedding dimensions. That's why we multiply the percentage of tokens in each layer by a weight ω_d equivalent to the embedding dimension in this layer as shown in Equation 6.

$$L_{eff} = L_{blocks} + L_{tokens} \quad (4)$$

$$L_{blocks} = MSE\left(\frac{B_{activ}}{B_{total}}, \frac{B_{target}}{B_{total}}\right) \quad (5)$$

$$L_{tokens} = MSE\left(\frac{\omega_d \times T_{activ}}{T_{total}}, \frac{\omega_d \times T_{target}}{T_{total}}\right) \quad (6)$$

To make our policy network task-aware, we co-train each task-specific policy sub-network independently with the end-to-end model. Sequentially co-training the policy sub-networks along with the end-to-end model suffers from

Algorithm 1: AdaMTL - Alternating Task Training (ATT)

Input: *model*: Model w/ initialized policy network
Tasks: Task names in MTL model
m: Number of tasks in MTL model
epochs_{att}: Number of epochs for ATT
L_i: Loss for the epoch *i*

```

1 for i in 0, ..., epochsatt do
2   current_task = Tasks[i%m];
3   Li = Lcurrent_task +  $\alpha$ Leff;
4   for task in Tasks do
5     if task = current_task then
6       model.unfreeze_decoder(task);
7       model.enable_policy_network(task);
8     end
9     else
10      model.freeze_decoder(task);
11      model.disable_policy_network(task);
12    end
13  end
14  train_one_epoch(model, Li);
15 end
```

catastrophic forgetting (i.e., the model gets biased towards behaving well on the last task, and the performance deteriorates on the earlier tasks). That's why we propose an *Alternating Task Training* (ATT) where we shift the focus between the tasks every one epoch. Algorithm 1 shows the steps of our ATT technique. For each epoch, we choose a task to focus on, and we set the loss accordingly, as shown in lines 2 and 3, respectively. Then, we only activate the decoder and the policy sub-network corresponding to the chosen task, as shown in lines 4-13. We train the MTL model for one epoch using that setting. Then, we move on to co-training the decoders and the policy sub-networks of the other tasks. Finally, we perform end-to-end fine-tuning to improve the overall model accuracy. The training loss remains the same as in Equation 3, and we unfreeze all the model's components. This results in an adaptive MTL model that generates task-aware inference policies depending on the complexity of the input.

4. Experiments

4.1. Setup

Dataset: We evaluate our method on the PASCAL dataset [5]. Following other papers in MTL literature [29, 33, 34], we use the PASCAL-Context split that has annotations for various dense prediction tasks such as semantic segmentation, human part detection, surface normals estimation, and saliency distillation. It has 4,998 images in the training split and 5,105 in the validation split.

Implementation and Training details: We implemented

Table 1. Quantitative analysis of AdaMTL on PASCAL dataset. The table shows the accuracy-efficiency trade-off by our adaptive MTL model compared to the single-task model as well as static MTL models. H and L represent high and low computational complexity targets for AdaMTL, respectively. Δm (ST) and Δ FLOPS (ST) show the change in the average accuracy of the tasks and percentage of FLOPS compared to the single-task model, respectively. \downarrow means the lower the better, while \uparrow means the higher the better. **Bolded** values represent the Pareto-frontier of the accuracy-efficiency trade-off.

Policy	Backbone	Image Size	Sal maxF \uparrow	Human Parts mIoU \uparrow	Sem Seg mIoU \uparrow	Normals mERR \downarrow	Δm (ST) (%) \uparrow	GFLOPS \downarrow	Δ FLOPS (ST) (%) \downarrow	Params (M) \downarrow
Single-Task	Swin-T	224	71.93	48.63	60.35	18.45	0.00	18.33	1 \times	111.42
Static MTL		224	74.15	47.62	59.08	19.20	-1.20	5.79	0.32 \times	34.77
AdaMTL (H)	Swin-T	224	73.72	47.64	57.13	19.16	-2.18	5.34	0.29 \times	34.87
AdaMTL (L)		224	73.01	46.85	55.9	19.54	-3.86	4.82	0.26 \times	34.87
Static MTL		224	75.00	50.66	61.84	18.71	+2.35	12.5	0.68 \times	67.02
AdaMTL (H)	Swin-S	224	75.23	51.22	61.88	18.79	+2.65	12.51	0.66 \times	67.12
AdaMTL (L)		224	74.75	50.07	59.91	18.61	+1.31	10.35	0.57 \times	67.12
Static MTL		384	73.93	56.68	67.47	17.69	+8.81	59.39	3.24 \times	108.66
AdaMTL (H)	Swin-B	384	76.55	56.28	65.04	17.74	+8.43	51.29	2.80 \times	108.88
AdaMTL (L)		384	76.23	55.04	62.63	17.92	+6.46	41.229	2.25 \times	108.88

AdaMTL using PyTorch. As mentioned in Subsection 3.1, we adopt the publicly available pre-trained Swin Transformer backbone [17] as our shared encoder. To get our adaptive MTL model, we employ a three-stage training recipe as explained in Subsection 3.3. In Stage 1, we fine-tune the Swin Transformer backbone along with our task-specific decoders for 1000 epochs. Then in stage 2, we freeze the MTL model and initialize the policy network by training it to activate the whole MTL model. We run this stage for another 80 epochs. Finally, in stage 3, we use our proposed *Alternating Task Training technique* to co-train the policy network along with the MTL model for another 150 epochs, followed by fine-tuning the end-to-end AdaMTL model for another 150 epochs. Our method does not only increase the model efficiency during inference, but it also reduces the carbon emission since we avoid retraining complex ViTs from scratch by reusing off-the-shelf pre-trained backbones; AdaMTL needs only 1 V100 GPU for around 24-48 hours (i.e., depending on the used backbone) in order to run our end-to-end training recipe.

4.2. Quantitative Analysis

Accuracy-Efficiency Trade-off: We apply AdaMTL on three SOTA ViTs from the Swin Transformer family [17]. We include *Swin-Tiny*, *Swin-Small*, and *Swin-Base*, representing three different scales of ViTs in terms of computational complexity. We include two different computational complexity targets: H and L . H represents a higher computational budget where the target percentage of activated tokens and blocks are 60% and 90%, respectively. L represents a lower computational budget where both the target percentage of activated tokens and blocks are set to 50%. In both settings, the accuracy-efficiency trade-off weight (i.e., α in Equation 3) is set to unity. To evaluate the accuracy-efficiency trade-off by AdaMTL, we compare it to two baselines: (1) Single-Task models and (2) Static MTL models

Table 2. Comparison with SOTA MTL models.

Method	Backbone	Δm (ST) \uparrow (%)	Δ FLOPS (ST) \downarrow (%)
PAD-Net [33]	HRNet-18	+4.98	23.21 \times
AdaMTL (Ours)	Swin-B	+6.46	2.25 \times
ASTMT [18]	R26-DLv3	+8.38	5.66 \times
AdaMTL (Ours)	Swin-B	+8.81	3.24 \times
MTI-Net [29]	ResNet-50	+13.49	9.6 \times
InvPT [34]	ViT-B	+27.20	21.35 \times

(i.e., our base MTL model before attaching our task-aware policy network). Table 1 shows the accuracy, the computational complexity (i.e., FLOPS) as well as the model size (i.e., Params) of AdaMTL compared to the single-task model and the static MTL models. Δm (ST) and Δ FLOPS (ST) show the change in the average accuracy of the tasks and percentage of FLOPS compared to the single-task model, respectively. Results show how our method enhances the accuracy-efficiency trade-off for MTL models. For example, by applying AdaMTL to *Swin-S* backbone, we can get an MTL model with 43% less FLOPS and 1.31% more accuracy compared to the single-task model. Similarly, by applying AdaMTL to *Swin-T* backbone, we can get an MTL model with 71% less FLOPS and only 2.18% drop in accuracy compared to the single-task model. Therefore, given any target computational complexity, AdaMTL can meet it while potentially improving the accuracy.

Comparison to SOTA MTL models: We also compare the accuracy-efficiency trade-off of AdaMTL to four SOTA MTL models that vary in computational complexity and accuracy. Figure 1 shows the accuracy and the computational complexity of AdaMTL compared to those of *PAD-Net* [33], *ASTMT* [18], *MTI-Net* [29], and *InvPT* [34]. We can notice that AdaMTL dominates both *PAD-Net* and *ASTMT*. Moreover, AdaMTL achieves a more efficient trade-off compared to *MTI-Net* and *InvPT*. As shown in Table 2, AdaMTL has 3 \times less FLOPS than *MTI-Net* and 7 \times less



Figure 3. Qualitative insights on the computational budget allocated by AdaMTL to process input frames of different complexity. We can notice that AdaMTL assigns fewer computations for simpler scenes, as in (a) and (b), while it justly assigns more computations to process more complex and cluttered scenes as in (c) and (d).

Table 3. Combining AdaMTL with SOTA MTL components.

Method	Backbone	Δm %	Δ FLOPS ↓	Params (M)
Single-Task	Swin-T	+0.00	18.33 G	111.42
AdaMTL	Swin-B	+8.81	1.1 \times	108.66
MTI-Net [29]	ResNet-50	+13.49	9.6 \times	91.00
AdaMTL + MTI-Net	Swin-B	+20.29	3.1 \times	101.14

FLOPS than *InvPT*. It is important to note that our effort in AdaMTL is directed specifically toward enabling efficient Multi-Task learning. That is why we argue that while the performance of *MTI-Net* and *InvPT* is impressive, their demanding computational complexity might not be suitable for real-time processing on resource-constrained devices.

4.3. Combining with SOTA MTL components

Our adaptive MTL framework can easily adopt other SOTA MTL components to further enhance the accuracy-efficiency trade-off. In this section, we show a case study where we integrate AdaMTL with the SOTA MTL concepts in *MTI-Net* [29] in order to improve both its efficiency and accuracy. *MTI-Net* has two main modules: (1) A multi-scale multi-modal distillation unit to model task interactions at different scales and (2) A feature propagation module that propagates distilled task information from lower to higher scales. In this experiment, we attach those two modules between the shared hierarchical encoder and the task-specific decoders in our AdaMTL framework. Following the exact same training recipe introduced in Subsection 3.3, the results in Table 3 show that AdaMTL can be integrated with other MTL modules from *MTI-Net* to boost *MTI-Net*'s accuracy by 7.8% while improving its efficiency by 3.1 \times .

4.4. Qualitative Analysis

Figure 3 shows some insights about the allocated amount of computations by AdaMTL for input frames of different visual complexity. Figures 3a, 3b, 3c, and 3d shows examples where AdaMTL allocated 55%, 65%, 75%, and 85% of the static model FLOPS respectively. We can see that AdaMTL allocates fewer computations to simple frames with fewer objects as compared to complex scenes




Figure 4. Sample of the generated masks by AdaMTL tokens policy network. The white areas in (b)-(d) represent activated tokens.

with multiple objects and cluttered backgrounds. Figure 3e shows a histogram of AdaMTL's computational complexity per example for all the images in PASCAL validation set. The red line represents the average number of FLOPS needed to process all the images. We can notice that AdaMTL adapts to the large variation in the computational complexity requirement by various images in the dataset.

To gain more insights into the decisions made by our policy network, we visualize a sample of the generated tokens masks in Figure 4. Column 4a represents the input frames, while columns 4b, 4c, and 4d represent the generated tokens masks by our policy network at different layers, such that the white areas represent activated tokens. We can notice that the policy network tends to activate more tokens in the earlier layers to understand the global features of the input frame. Then, it narrows down its scope in the later layers, focusing on the most informative patches (i.e., patches with the main objects) in the input frame.

4.5. Deployment on Vuzix M4000 AR glasses

We compile our AdaMTL model using PyTorch for Android [23], and we deploy it on the Vuzix M4000 AR glasses [31]. The Vuzix glasses have an 8 Core 2.52Ghz Qualcomm XR1 board with 6GB RAM. It operates using Android 11.0. Using the *Battery Historian* tool [10] to profile the energy consumption on the device, we record the average latency

Table 4. Performance Analysis on Vuzix Augmented Reality Glasses [31]. We analyze the percentage of the inference latency and the energy consumption by our AdaMTL model compared to its corresponding static model.

Method	Backbone	Inference Latency (%)	Energy Consumption (%)
AdaMTL	Swin-T	-20.6%	-35.3%
AdaMTL	Swin-S	-21.8%	-37.5%

Table 5. Comparison between the quality of our task-aware policy, the task-agnostic, and the random execution policy.

Policy	Backbone	Δm (ST) (%) \uparrow	Δ FLOPS (ST) \downarrow
Random	Swin-T	-34.53	0.24 \times
Random+		-8.64	0.24 \times
Task-Agnostic		-5.68	0.24 \times
Task-Aware		-3.86	0.26\times
Random	Swin-B	-34.86	0.82 \times
Random+		-8.57	0.82 \times
Task-Agnostic		+0.77	0.84 \times
Task-Aware		+1.12	0.83\times

and energy consumption across random samples from the PASCAL validation dataset. Table 4 shows that AdaMTL reduces the inference latency and the energy consumption by up to 21.8% and 37.5%, respectively, compared to its corresponding static MTL model.

4.6. Ablation Study

Quality of the learnt inference policies: To analyze the quality of the learned inference policies by our task-aware policy network. Table 5 compares the accuracy-efficiency trade-off achieved by our task-aware policy network (i.e., referred to as Task-Aware) to three other baselines: (1) Random where we activate random blocks and tokens from the static MTL model, (2) Random+ where we again activate random blocks and tokens from the static MTL model but after performing our adaptive training pipeline, and (3) Task-Agnostic policy network explained in Subsection 3.2. The results in Table 5 show that the policies learned by our task-aware policy network outperform the other baselines. We can also notice that performing adaptive training gives the MTL model robustness towards sparsification (i.e., Using a random policy on the static model reduced the accuracy by 34%, while it only reduced the accuracy by 8% when applied to the adaptively trained MTL model).

Analysis of the adaptation along each component of our policy network To understand the contribution of both components of our policy network (i.e., blocks policy network and tokens policy network) to the accuracy-efficiency trade-off of AdaMTL, we compare the results from four different settings: (1) the static model where neither component of the policy network is activated, (2) AdaMTL while activating the block policy network only, (3) AdaMTL while activating the tokens policy network only, and (4)

Table 6. The contribution of different adaptive dimensions to AdaMTL. Δm is measured relative to the static MTL model.

Adaptive Blocks	Adaptive Tokens	Δm %	FLOPS (G)
\times	\times	-0.00	5.8
\checkmark	\times	-1.66	5.23
\times	\checkmark	-1.46	5.08
\checkmark	\checkmark	-0.85	5.37




Figure 5. Loss function ablation: The figure illustrates the average percentage of activated tokens over the 12 blocks in Swin-T using (a) non-weighted tokens loss and (b) weighted tokens loss.

AdaMTL where both policy networks are activated. Table 6 shows that enabling both components enhances the accuracy-efficiency trade-off of AdaMTL.

Analysis of the behavior of our loss function To analyze the importance of the weight factor ω_d in Equation 6 of our loss function, we visualize the average percentage of activated tokens across the blocks of Swin-T encoder with and without ω_d in Figure 5. We can notice that adding the ω_d factor to the loss function leads to a more distributed FLOPS reduction across different blocks, which is essential for the effectiveness of our adaptive MTL framework.

5. Conclusion

In this paper, we propose AdaMTL - an adaptive framework that learns task-aware inference policies for the MTL models in an input-dependent manner. We achieve this by co-training a lightweight policy network along with our MTL model. During runtime, our policy network recognizes the unnecessary computations and dynamically chooses an execution strategy depending on the input complexity and the target computational budget. Our experiments on PASCAL dataset demonstrate that AdaMTL reduces the computational complexity by 43% while improving the accuracy by 1.32% compared to the single task models. Combined with SOTA MTL components, AdaMTL boosts the accuracy by 7.8% while improving the SOTA MTL model efficiency by 3.1 \times . Finally, we deployed AdaMTL on Vuzix M4000 AR glasses showing up to 21.8% and 37.5% reduction in inference latency and energy consumption, respectively, compared to the static MTL model.

Acknowledgements: This work is partially supported by NSF grant 1814920 and DoD ARO grant W911NF-19-1-0484.

References

- [1] Deblina Bhattacharjee, Tong Zhang, Sabine Süsstrunk, and Mathieu Salzmann. Mult: an end-to-end multitask learning transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12031–12041, 2022. 2
- [2] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, pages 213–229. Springer, 2020. 2
- [3] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 2
- [4] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2
- [5] Mark Everingham, Luc Van Gool, Christopher Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88:303–338, 06 2010. 5
- [6] Benjamin Graham, Alaeldin El-Nouby, Hugo Touvron, Pierre Stock, Armand Joulin, Hervé Jégou, and Matthijs Douze. Levit: a vision transformer in convnet’s clothing for faster inference. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 12259–12269, 2021. 2
- [7] Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, et al. Conformer: Convolution-augmented transformer for speech recognition. *arXiv preprint arXiv:2005.08100*, 2020. 2
- [8] Hussein Hazimeh, Zhe Zhao, Aakanksha Chowdhery, Maheswaran Sathiamoorthy, Yihua Chen, Rahul Mazumder, Lichan Hong, and Ed Chi. Dselect-k: Differentiable selection in the mixture of experts with applications to multi-task learning. *Advances in Neural Information Processing Systems*, 34:29335–29347, 2021. 3
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 2, 3
- [10] Battery historian, August 2016. <https://github.com/google/battery-historian>. 7
- [11] Gao Huang, Danlu Chen, Tianhong Li, Felix Wu, Laurens Van Der Maaten, and Kilian Q Weinberger. Multi-scale dense networks for resource efficient image classification. *arXiv preprint arXiv:1703.09844*, 2017. 3
- [12] Keishi Ishihara, Anssi Kanervisto, Jun Miura, and Ville Hautamaki. Multi-task learning with attention for end-to-end autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2902–2911, 2021. 1
- [13] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016. 4
- [14] Hanxue Liang, Zhiwen Fan, Rishov Sarkar, Ziyu Jiang, Tianlong Chen, Kai Zou, Yu Cheng, Cong Hao, and Zhangyang Wang. M³vit: Mixture-of-experts vision transformer for efficient multi-task learning with model-accelerator co-design. *arXiv preprint arXiv:2210.14793*, 2022. 3
- [15] Shikun Liu, Edward Johns, and Andrew J Davison. End-to-end multi-task learning with attention. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1871–1880, 2019. 1
- [16] Shikun Liu, Edward Johns, and Andrew J Davison. End-to-end multi-task learning with attention. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1871–1880, 2019. 2
- [17] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 2, 3, 4, 6
- [18] Kevis-Kokitsi Maninis, Ilija Radosavovic, and Iasonas Kokkinos. Attentive single-tasking of multiple tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1851–1860, 2019. 6
- [19] Lingchen Meng, Hengduo Li, Bor-Chun Chen, Shiyi Lan, Zuxuan Wu, Yu-Gang Jiang, and Ser-Nam Lim. Adavit: Adaptive vision transformers for efficient image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12309–12318, 2022. 2, 3
- [20] Ishan Misra, Abhinav Shrivastava, Abhinav Gupta, and Martial Hebert. Cross-stitch networks for multi-task learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3994–4003, 2016. 2
- [21] Eslam Mohamed and Ahmad El Sallab. Spatio-temporal multi-task learning transformer for joint moving object detection and segmentation. In *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, pages 1470–1475. IEEE, 2021. 1
- [22] Marina Neseem and Sherief Reda. Adacon: Adaptive context-aware object detection for resource-constrained embedded devices. In *2021 IEEE/ACM International Conference On Computer Aided Design (ICCAD)*, pages 1–9. IEEE, 2021. 3
- [23] Pytorch mobile: End-to-end workflow from training to deployment for ios and android mobile devices, March 2023. <https://pytorch.org/mobile/android>. 7
- [24] Trevor Standley, Amir Zamir, Dawn Chen, Leonidas Guibas, Jitendra Malik, and Silvio Savarese. Which tasks should be learned together in multi-task learning? In *International Conference on Machine Learning*, pages 9120–9132. PMLR, 2020. 1
- [25] Robin Strudel, Ricardo Garcia, Ivan Laptev, and Cordelia Schmid. Segmenter: Transformer for semantic segmenta-

- tion. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7262–7272, 2021. 2
- [26] Ximeng Sun, Rameswar Panda, Rogerio Feris, and Kate Saenko. Adashare: Learning what to share for efficient deep multi-task learning. *Advances in Neural Information Processing Systems*, 33:8728–8740, 2020. 2
- [27] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International conference on machine learning*, pages 10347–10357. PMLR, 2021. 2
- [28] 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. 2
- [29] Simon Vandenhende, Stamatios Georgoulis, and Luc Van Gool. Mti-net: Multi-scale task interaction networks for multi-task learning. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16*, pages 527–543. Springer, 2020. 2, 4, 5, 6, 7
- [30] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 2
- [31] Vuzix m4000 smart glasses, March 2023. <https://www.vuzix.com/products/m4000-smart-glasses>. 2, 7, 8
- [32] Xin Wang, Fisher Yu, Zi-Yi Dou, Trevor Darrell, and Joseph E Gonzalez. Skipnet: Learning dynamic routing in convolutional networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 409–424, 2018. 3
- [33] Dan Xu, Wanli Ouyang, Xiaogang Wang, and Nicu Sebe. Pad-net: Multi-tasks guided prediction-and-distillation network for simultaneous depth estimation and scene parsing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 675–684, 2018. 2, 5, 6
- [34] Hanrong Ye and Dan Xu. Inverted pyramid multi-task transformer for dense scene understanding. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXVII*, pages 514–530. Springer, 2022. 2, 5, 6
- [35] Zhongzhi Yu, Yonggan Fu, Sicheng Li, Chaojian Li, and Yingyan Lin. Mia-former: efficient and robust vision transformers via multi-grained input-adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 8962–8970, 2022. 2, 3
- [36] Amir R Zamir, Alexander Sax, William Shen, Leonidas J Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling task transfer learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3712–3722, 2018. 1
- [37] Wangchunshu Zhou, Canwen Xu, Tao Ge, Julian McAuley, Ke Xu, and Furu Wei. Bert loses patience: Fast and robust inference with early exit. *Advances in Neural Information Processing Systems*, 33:18330–18341, 2020. 3