GRAPH COLORING FOR MULTI-TASK LEARNING

Anonymous authors

000

001 002 003

004

006

008 009

010

011

012

013

014

016

018

019

021

024

025 026 027

028 029

031

032

033

034

037

040

041

042

043

044

045

046

047

048

051

052

Paper under double-blind review

ABSTRACT

When different objectives conflict with each other in multi-task learning, gradients begin to interfere and slow convergence, thereby potentially reducing the final model's performance. To address this, we introduce SON-GOKU, a scheduler that computes gradient interference, constructs an interference graph, and then applies greedy graph-coloring to partition tasks into groups that align well with each other. At each training step, only one group (color class) of tasks are activated, and the grouping partition is constantly recomputed as task relationships evolve throughout training. By ensuring that each mini-batch contains only tasks that pull the model in the same direction, our method improves the effectiveness of any underlying multi-task learning optimizer without additional tuning. Since tasks within these groups will update in compatible directions, multi-task learning will improve model performance rather than impede it. Empirical results on six different datasets show that this interference-aware graph-coloring approach consistently outperforms baselines and state-of-the-art multi-task optimizers. We provide extensive theory showing why grouping and sequential updates improve multi-task learning, with guarantees on descent, convergence, and the ability to accurately identify what tasks conflict or align.

1 Introduction

Multi-task learning (MTL) trains a single model to solve several tasks simultaneously, sharing knowledge across them to learn more effectively (Caruana, 1997). This allows models to generalize better and converge faster. However, a key issue known as negative transfer arises when tasks don't align very well with each other (Sener & Koltun, 2018; Shi et al., 2023). When two tasks push the shared network in different directions their gradients clash, slowing or even reversing learning. Prior work addresses this issue primarily via (1) gradient manipulation, which reshapes task gradients to reduce conflicts, and (2) loss reweighting, which rescales task objectives to balance their influence. While effective in some specific settings, these strategies typically treat conflict locally at the level of shared-parameter updates and often overlook the evolving global structure of interactions among tasks throughout training.

Some recent works focus on partitioning tasks into subsets (groups) and updating those groups separately. These approaches have been found to improve accuracy and training stability by forming groups with high measured affinity and then updating one group at a time (Fifty et al., 2021; Jeong & Yoon, 2025). Grouping can outperform gradient manipulation and loss reweighting when tasks form clusters with aligned gradients, because each update then reduces direct clashes in the shared layers, lowers gradient variance within the step, and lets compatible tasks reinforce one another while conflicting tasks wait for their turn.

However, grouping methods often face a few key limitations: (1) many rely on dense pairwise affinities that grow noisy and costly as the number of tasks rises (Fifty et al., 2021; Standley et al., 2020; Sherif et al., 2024), and (2) others predetermine or rarely update groups, so they drift as task relations change (Wang et al., 2024; Ruder, 2017), and (3) several use local heuristics that fail to enforce global compatibility or to specify how groups should rotate over time (Zhang & Yang, 2018; Malhotra et al., 2022).

We present SON-GOKU (Scheduling via Optimal INterference-aware Graph-COloring for TasK Grouping in MUltitask Learning). We measure gradient interference, build a graph of tasks from those measurements, greedily color the graph to form non-conflicting compatible task groups, and

update one color group per step during training. This design addresses the earlier issues. We estimate the interference graph from lightweight minibatch statistics and keep it sparse, which avoids noisy dense matrices and scales to many tasks. We recolor the graph at regular intervals so the groups track changing relations during training. Greedy graph coloring ensures we update only compatible tasks in each step, and the color order gives a simple way to cycle through the groups. Our proposed scheduler does not have to work in isolation, it can function on top of existing loss-reweighting and gradient-manipulation MTL approaches.

In our theoretical analysis (Section 5) we show that, under standard conditions, SON-GOKU tends to group tasks whose gradients are, on average, aligned within each group, with high probability. We further show that, over a refresh window, sequentially updating these low-conflict groups yields *at least* as much expected descent as a single mixed update, and *strictly more* when between-group interference is sufficiently negative. We also prove that SON-GOKU preserves descent and reaches the usual non-convex SGD rate under mild assumptions, with only a small factor that depends on the within-group conflict level. In Appendix D we discuss the scheduler's amortized time complexity and the tradeoffs it offers between speed and performance. We discuss ways in which practitioners can reduce its time complexity under certain conditions.

Empirical results from experiments demonstrate that SON-GOKU consistently improves outcomes compared to other MTL approaches, especially when SON-GOKU is coupled with existing approaches. Our contributions are as follows:

- We propose SON-GOKU, an interference-aware scheduler that measures cross-task gradient conflict, builds a conflict graph, colors it to form compatible groups, and activates one group per step. It can be used on top of standard MTL optimizers.
- We provide theoretical analysis that offers guarantees on SON-GOKU's grouping, convergence, scheduling behavior, and more.
- Across six datasets, SON-GOKU improves over strong baselines and pairs well with methods like PCGrad, AdaTask, and GradNorm, delivering consistent gains.
- We perform an ablation study showing that dynamic recoloring and history-averaged conflict estimates are key contributors to performance.

2 RELATED WORK

Many MTL methods (especially earlier ones) adjust task influence by learning or adapting loss weights. Examples include uncertainty-based scaling (Kendall et al., 2018), rate-based schemes such as DWA (Liu et al., 2019), and fast bilevel formulations like FAMO (Liu et al., 2023). FAMO in particular is notable for its $\mathcal{O}(1)$ per-step time complexity. These approaches keep all tasks active each step while modulating relative magnitudes. A completely different approach, which emerged in 2018 with MGDA (Sener & Koltun, 2018), focuses on updating shared-parameter update directions to mitigate interference. Methods like PCGrad (Yu et al., 2020), CAGrad (Liu et al., 2021), and MGDA (Sener & Koltun, 2018) modify the geometry of the shared update to reduce cross-task conflicts while still updating all tasks each step. A smaller body of work forms subsets of tasks to update together, using offline affinity estimation or training-dynamics signals (Fifty et al., 2021; Standley et al., 2020; Wang et al., 2024; Sherif et al., 2024). Most recently, Selective Task Group Updates proposes online grouping with sequential updates, reporting that update order can influence task-specific learning (Jeong & Yoon, 2025). SON-GOKU differs in mechanism from existing approaches (Section 4). It complements loss reweighting and gradient surgery, and we provide explicit guarantees on descent, convergence, and graph/partition recovery. An expanded discussion of more related work is provided in Appendix M.

3 Problem Setup

We formalize multi-task learning (MTL) (Caruana, 1997) as optimizing a shared network while activating only a subset of tasks at each step. Each task contributes a loss whose gradients may align or conflict. We quantify conflict using (the negative of) cosine similarity, embed tasks in a conflict graph, and later use that graph to derive a schedule. This section fixes notation and states the optimization goal that the proposed approach addresses.

3.1 Data and Notation

Let $\mathcal{T} = \{T_1, \dots, T_K\}$ be the set of tasks. The model has shared parameters $\theta \in \mathbb{R}^d$ and task-specific parameters $\phi_k \in \mathbb{R}^{d_k}$ for T_k . Each task draws examples (x, y_k) from a distribution \mathcal{D}_k and defines a per-example loss $\ell_k(\theta, \phi_k; x, y_k)$. Its population loss is

$$L_k(\theta, \phi_k) := \mathbb{E}_{(x, y_k) \sim \mathcal{D}_k} [\ell_k(\theta, \phi_k; x, y_k)]. \tag{1}$$

We minimize the standard weighted MTL objective

$$F(\theta, \phi_1, \dots, \phi_K) = \sum_{k=1}^K w_k L_k(\theta, \phi_k), \tag{2}$$

with nonnegative task weights w_k (default $w_k = 1$). Note that, for simplicity in later sections, we absorb w_k into the per-task gradient estimates. This is permissible since positive scalings do not change cosine signs or the induced conflict graph.

At step t, for any task k that is active we compute stochastic gradients on a mini-batch $\mathcal{B}_k^{(t)} \subset \mathcal{D}_k$:

$$g_k^{(t)} := \nabla_{\theta} L_k(\theta_t, \phi_{k,t}; \mathcal{B}_k^{(t)}), \qquad h_k^{(t)} := \nabla_{\phi_k} L_k(\theta_t, \phi_{k,t}; \mathcal{B}_k^{(t)}). \tag{3}$$

In our proposed method, we form exponential moving averages (EMA) of per-task gradients within a refresh window to stabilize cosine estimates so that they do not become stale (Sec. 4).

3.1.1 Interference Coefficient

We quantify pairwise interaction with the interference coefficient

$$\rho_{ij} = -\frac{\langle \tilde{g}_i, \tilde{g}_j \rangle}{\|\tilde{g}_i\| \|\tilde{g}_j\|},\tag{4}$$

where \tilde{g}_i and \tilde{g}_j are the EMA-smoothed gradients at refresh. Positive ρ_{ij} indicates conflict (negative cosine). $\rho_{ij} \leq 0$ indicates alignment or neutrality.

3.1.2 CONFLICT GRAPH

Fix a tolerance $\tau \in (0,1)$. The conflict graph is

$$G_{\tau} = (\mathcal{T}, E_{\tau}), \qquad E_{\tau} = \{(i, j) : \rho_{ij} > \tau\}. \tag{5}$$

Vertices are tasks. An edge between a pair means to not update that pair together. We will utilize G_{τ} for coloring and scheduling in Section 4

3.2 GOAL

At training step t we choose an active set $S_t \subseteq \mathcal{T}$ and update only those tasks:

$$\theta_{t+1} = \theta_t - \eta_t \sum_{k \in S_t} g_k^{(t)}, \qquad \phi_{k,t+1} = \begin{cases} \phi_{k,t} - \eta_t h_k^{(t)}, & k \in S_t, \\ \phi_{k,t}, & k \notin S_t. \end{cases}$$
(6)

The problem the scheduler addresses is to design the sequence $\{S_t\}_{t=1}^T$ so that: (1) every task is visited regularly; and (2) conflicting tasks seldom appear together. We instantiate this via greedy graph coloring in Section 4 and analyze the guarantees in Section 5.

4 PROPOSED APPROACH

We design an interference-aware scheduler that partitions tasks into low-conflict groups and activates exactly one group per optimization step. The procedure consists of four stages: (1) estimating pairwise interference, (2) building and coloring the conflict graph, (3) generating a periodic schedule, and (4) updating that schedule as training evolves. An overview of the scheduler is provided as Algorithm 1 in Appendix A. A visualization of SON-GOKU is provided in Figure 1 alongside a simple summary in the Figure caption.

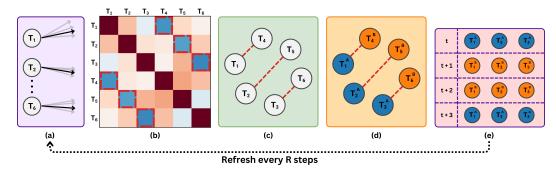


Figure 1: Interference-aware scheduling pipeline: (a) For each task T_i (circles $T_1 \dots T_6$), we smooth recent per-step gradients with an Exponential Moving Average (EMA); (b) From these EMA vectors we compute the pairwise cosine matrix. In the figure, cells outlined with red dashes mark pairs with cosine $<-\tau$. These are flagged as conflicts; (c) We build the conflict graph whose nodes are tasks T_i and whose red dashed edges connect exactly those pairs identified in (b); (d) We apply greedy graph coloring so that no conflict edge lies within a color, producing low-conflict groups. In the example shown, we have two groups: A as blue and B as orange; (e) During training we activate one group per step. After every R steps (here, R=4) we 'refresh' and run the pipeline again from step A, where we update the EMAs with the latest gradients.

4.1 ESTIMATING GRADIENT INTERFERENCE

We absorb task weights into per-task losses, so $g_k^{(t)}$ is the gradient of the weighted loss $w_k L_k$. Cosine calculations and graph construction are not impacted by applying positive scaling.

At step t and for every task T_k appearing in the current mini-batch we compute a task-specific stochastic gradient

$$g_k^{(t)} = \nabla_{\theta} \mathcal{L}_k(\theta_t, \phi_{k,t}; \mathcal{B}_k^{(t)}), \tag{7}$$

using an independent sub-batch $\mathcal{B}_k^{(t)}\subset\mathcal{D}_k$. We then update an exponential moving average

$$\tilde{g}_k^{(t)} = \beta \, \tilde{g}_k^{(t-1)} + (1-\beta) \, g_k^{(t)}, \qquad \beta \in [0,1),$$
 (8)

which stabilizes cosine estimates while requiring only two buffers per task (current and previous). Whenever we refresh the schedule (every R steps) we form the pairwise interference matrix

$$\rho_{ij}^{(t)} = -\frac{\langle \tilde{g}_i^{(t)}, \tilde{g}_j^{(t)} \rangle}{\|\tilde{g}_i^{(t)}\| \|\tilde{g}_i^{(t)}\|}, \qquad i, j \in \{1, \dots, K\}.$$

$$(9)$$

Computing all K(K-1)/2 cosines is $O(K^2d)$ with d representing the shared-parameter dimension. We also write $h_k^{(t)} = \nabla_{\phi_k} \mathcal{L}_k(\theta_t, \phi_{k,t}; \mathcal{B}_k^{(t)})$ for the gradient with respect to the task-specific parameters ϕ_k .

4.2 CONFLICT GRAPH CONSTRUCTION

Given a tolerance $\tau \in (0,1)$, the conflict graph at update round r is

$$G_{\tau}^{(r)} = (V, E_{\tau}^{(r)}), \quad V = \{1, \dots, K\} E_{\tau}^{(r)} = \{(i, j) : \rho_{ij}^{(t_r)} > \tau\}.$$
 (10)

To clarify, tasks are indexed by integers $1 \dots K$ in Equation 10. Edges connect tasks whose averaged gradients have cosine similarity less than $-\tau$. Intuitively, larger τ yields a sparser conflict graph, typically fewer colors (larger per-step groups), and more frequent updates per task. Smaller τ results in a denser graph, more colors (smaller per-step groups), and less frequent updates per task.

4.3 PARTITIONING VIA GREEDY GRAPH COLORING

We apply the Welsh-Powell largest-first greedy heuristic (Welsh & Powell, 1967) to color $G_{\tau}^{(r)}$ and obtain color classes $C_1^{(r)}, \ldots, C_{m_r}^{(r)}$. Classical graph-theory results (West, 2000; Diestel, 2017) guarantee the heuristic uses no more than $\Delta + 1$ colors, where Δ is the maximum vertex degree. In practice Δ is small because many task pairs do not interfere, yielding concise schedules.

4.4.6

4.4 SCHEDULE GENERATION AND EXECUTION

We create a periodic schedule of length m_r :

$$S_t = C_{(t \bmod m_r)+1}^{(r)}, \qquad t_r \le t < t_{r+1} = t_r + R.$$
 (11)

Each training step activates exactly one color class; over one period every task in that class receives a gradient update, while conflicting tasks (edges in $E_{\tau}^{(r)}$) are guaranteed not to co-occur.

4.4.1 MINIMUM UPDATE FREQUENCY

If the greedy coloring yields a singleton class for a rarely updated task, we increase its update frequency by duplicating it only into steps whose active color has no conflict edge to that task.

4.4.2 WARM-UP AND ANNEALING

We start with $\tau=1$ (no edges, full simultaneous training) for the first $T_{\rm warm}$ steps, then logarithmically anneal τ to a target value τ^* . This mitigates noisy gradient signals early in training.

4.5 TIME COMPLEXITY

The proposed scheduler has a time complexity of $\Theta(K^2d)$ per refresh. However, unlike many MTL approaches, our scheduler concentrates its extra work in occasional refreshes. This time complexity therefore becomes $\Theta(K^2d/R)$ amortized per training step where R s the refresh period (the number of training steps between conflict-graph rebuilds). It adds non-trivial overhead which grows quadratically with K (number of tasks) but shrinks as R grows. We provide a full analysis of the time complexity in Appendix D and discuss approaches to reducing time complexity under certain conditions in Appendix D.5.

5 THEORETICAL ANALYSIS

We discuss some of the main guarantees behind SON-GOKU. For a very brief overview: (1) Updating groups of tasks whose gradients are mostly low-conflict (no internal edges) reduces the objective on average and still achieves the usual $1/\sqrt{T}$ convergence rate; (2) Over a refresh window, scheduling several group updates can beat one mixed update that uses all tasks at once; and (3) With a small number of recent gradient measurements per task (via EMA) and a margin separating conflicts, the estimated conflict graph matches the ideal one, giving a short schedule where every task is updated at least once every $\Delta+1$ steps (Δ is the maximum number of conflicts for any task). We provide expanded assumptions, definitions, proofs, reasoning, analysis, etc. in Appendix 5.4–I.

5.1 DESCENT PRESERVATION WITHIN A LOW-CONFLICT GROUP

If the active set S_t at step t is τ -compatible, then the combined update is a descent direction with a quantitative lower bound:

$$\left\| \sum_{k \in S_{t}} g_{k,t} \right\|^{2} \ge \left(1 - \tau \left(|S_{t}| - 1 \right) \right) \sum_{k \in S_{t}} \|g_{k,t}\|^{2} \tag{12}$$

Thus the step cannot flip to ascent whenever $\tau(|S_t|-1) < 1$. This is proved by expanding the polarization identity and controlling cross terms under the τ -compatibility condition (see Appendix

E). Essentially, this means that SON-GOKU's per-step updates are safe when groups are low conflict. The aggregate direction keeps pointing downhill and the cancellation is quantitatively limited by τ and group size.

5.2 NONCONVEX CONVERGENCE AT THE STANDARD RATE UP TO A SMALL FACTOR

Under standard smoothness and noise conditions (see Appendix I) and with steps $\eta=c/\sqrt{T}$, SON-GOKU achieves the usual nonconvex SGD rate, with a mild $(1+\tau)$ factor that reflects within-group conflict:

$$\min_{t < T} \mathbb{E} \|\nabla F(\theta_t)\|^2 \le \frac{2(F_0 - F^*)}{c\sqrt{T}} (1 + \tau) + \frac{cL\sigma^2}{\sqrt{T}}$$
(13)

When $\tau=0$, the constant matches the classical bound (Bottou et al., 2018; Ghadimi & Lan, 2013); as $\tau\to 1$, it at most doubles, matching the intuition that conflict can cancel up to half of the progress. This demonstrates that scheduling does not degrade asymptotic progress. SON-GOKU preserves the $1/\sqrt{T}$ decay of the gradient norm while controlling the constant through the compatibility threshold τ . In other words, we keep the standard rate of SGD and trade a small constant for reduced interference.

5.3 WHEN SCHEDULED GROUPS OUTPERFORM A SINGLE MIXED UPDATE

We compare two ways to use the same gradients gathered at a refresh: a scheduled sequence of per-group steps (i.e., the scheduler used in SON-GOKU) versus a single aggregated step. Using a telescoping L-smooth bound and evaluating both trajectories at a common linearization—i.e., expanding F at the refresh start θ_{t_r} and applying the same first-order model with the same step size—the scheduled bound is never worse and is strictly better when cross-group interaction terms are sufficiently negative (so mixed updates would cancel progress).

Essentially, when different groups' gradients pull in opposing directions (so adding them together would cancel progress) the scheduler has an advantage. In that case, taking the updates one group at a time is provably better. Our theory guarantees a larger drop in the objective during that refresh than the one-shot step, even though both use the same step size and the same gradients. Under the PL condition, the scheduled path maintains the usual contraction factor and gains a nonnegative extra decrease term over the window.

5.4 EXACT RECOVERY OF THE POPULATION CONFLICT GRAPH AND TASK PARTITION

We show that, after observing gradients for only a modest number of steps, the scheduler can exactly reconstruct the true conflict relations among tasks by averaging recent gradients (EMA), computing pairwise cosines, thresholding at $-\tau$, and coloring the resulting graph. Under a separation margin γ around the threshold (tasks are meaningfully different), bounded noise, and bounded drift within each refresh window, the conflict graph estimated from finite data agrees, with high probability, with the ideal population conflict graph $G^*\tau$ (defined from the pairwise cosines of the true mean gradients $\{\mu_i\}_{i=1}^K$ at the start of the refresh window). Equivalently, when the uniform cosine estimation error is below γ , we have $\hat{G}\tau = G^*\tau$ and the resulting grouping recovers the ground-truth task partition. This explains why the scheduler's group structure is trustworthy and ties the required number of recent gradient measurements per task to interpretable quantities such as noise level, margin, and the number of tasks. For example, an effective sample size of $n_{\rm eff} \gtrsim \frac{\sigma^2}{m_0^2 \gamma^2} \log(K/\delta)$ suffices in our analysis.

5.5 SCHEDULING PROPERTIES WITH FEW GROUPS AND BOUNDED STALENESS

Welsh-Powell greedy coloring uses at most $\Delta+1$ colors on a graph whose maximum degree is Δ (Bonamy et al., 2018). Running the colors in a fixed cycle means each task is updated at least once every $m \leq \Delta+1$ steps. Equivalently, no task waits more than Δ steps between updates (bounded staleness).

This means that the schedule length is controlled by the worst conflict degree Δ rather than by the total number of tasks K. This results in two important benefits: (1) a minimum update-frequency guarantee, since every task receives an update at least once per cycle of length $\leq \Delta + 1$; and (2)

compatibility with standard bounded-delay conditions used in analyses of asynchronous SGD (e.g., Niu et al. 2011; Lian et al. 2015), with delay parameter at most Δ . When $\Delta \ll K$, we achieve both low interference (few conflicts per step) and low staleness (short update gaps).

6 EXPERIMENTAL SETUP

6.1 Datasets

We evaluate across six benchmarks spanning vision, multimodal, and time-series. For each dataset we specify a small set of primary tasks and add positive and negative auxiliaries to stress interference. Architectures are standard backbones (e.g., ResNet-18 for image tasks, CNN/BiLSTM for time-series) with task-specific heads. Full dataset and task definitions, auxiliary construction, and architecture details (including preprocessing and head designs) appear in Appendix J and Table 3.

6.2 Baseline and State-of-the-Art Comparisons

We compare against loss-weighting (Uniform, GradNorm, AdaTask), multi-objective (MGDA, Nash-MTL, FairGrad), projection/surgery (PCGrad, CAGrad), and fast adaptive weighting (FAMO). We provide short method notes in Appendix K and discuss these approaches in Section 2.

6.3 SCHEDULER EXTENSION MODELS

In addition to standalone models, we also evaluate combinations of the scheduler with existing approaches.

1. SON-GOKU + AdaTask. Combines our interference-aware task selection with AdaTask's dynamic loss weighting, applying adaptive weights only to scheduler-selected tasks.

2. SON-GOKU + GradNorm Warm Start. Initializes training with GradNorm for stable gradient magnitudes, then transitions to our scheduler after 3 epochs.

3. SON-GOKU + PCGrad. Applied PCGrad's gradient projection specifically to tasks selected by our scheduler, providing fine-grained conflict resolution within τ -compatible groups.

6.4 ABLATION STUDY

6.4.1 STATIC ONE-SHOT COLORING

We run the greedy graph coloring once at the start of training, freeze the resulting task groups, and never recompute the conflict graph. All other hyperparameters (τ) , history length H, and update interval R) match the full scheduler. As training progresses we expect the fixed coloring to grow stale, mixing tasks whose interference relationships have changed. This ablation isolates the benefit of dynamic recoloring, showing how much performance depends on adapting the schedule to evolving gradient conflicts.

6.4.2 SINGLE-STEP CONFLICT ESTIMATION

 Here, we set the history length to H=1, so every recoloring step relies on only the most recent mini-batch gradients to estimate interference. Without aggregation over many past steps, the conflict graph should become highly noisy, causing unstable task groupings from one update window to the next. This variant tests the importance of historical conflict statistics in the scheduler.

7 RESULTS AND DISCUSSION

Results for all models across every experiment are depicted in Table 1. Across ten metrics on six datasets, our conflict-aware schedulers consistently match or exceed all baseline methods.

378

379 380

391 392 393

394

396

397

403 404 405

406

407

402

412 413

414

415 416 417

418 419

420 421

422 423

424

425 426

427 428 429

430 431

Table 1: Performance of Evaluated Approaches Across Datasets

Model	Accuracy (%) ↑		F&B		HEALTH		NYUv2			
	CIFAR-10	AV-MNIST	MM-IMDb	Acc. (%) ↑	MAE ↓	Acc. (%) ↑	MAE ↓	Angle Error ↓	Seg. MIOU ↑	Depth RMSE ↓
Uniform	55	63	56	45	0.57	52	0.54	21.6	0.059	0.73
GradNorm	61	65	58	47	0.57	53	0.52	21.4	0.054	0.65
MGDA	59	62	56	44	0.57	53	0.53	21.8	0.63	0.75
PCGrad	61	65	58	50	0.55	58	0.48	20.9	0.07	0.69
CAGrad	59	62	57	46	0.58	53	0.52	21.9	0.65	0.73
AdaTask	63	67	59	47	0.59	55	0.52	20.3	0.69	0.65
FAMO	64	70	61	52	0.53	60	0.49	19.9	0.074	0.63
FairGrad	62	66	59	52	0.54	60	0.47	20.7	0.072	0.67
Nash-MTL	63	66	60	52	0.54	60	0.47	20.6	0.073	0.67
Static One-Shot	61	66	58	48	0.56	54	0.51	20.5	0.071	0.65
Single-Step	40	59	20	42	0.60	47	0.55	26.4	0.042	0.81
SON-GOKU + GradNorm	62	69	59	51	0.53	59	0.49	19.6	0.073	0.64
SON-GOKU + AdaTask	67	71	63	52	0.53	59	0.48	20.1	0.68	0.67
SON- $GOKU + PCGrad$	65	70	60	54	0.52	62	0.45	19.7	0.076	0.62
SON-GOKU	65	69	61	51	0.53	58	0.50	19.8	0.073	0.59

7.1 Overall Performance Improvements

Overall, the conflict-aware approaches improve over the uniform baseline by 10%-20% on CIFAR-10 and by 7% on MM-IMDb. This reinforces the idea that grouping tasks according to measured interference is more effective than treating all tasks equally at every update. On NYUv2, we see similar improvements across all the metrics. These results suggest that the scheduler's graph coloring cleanly separates high-conflict tasks, preserving the projection or LR-balancing advantages (stemming from PCGrad's gradient projection and AdaTask's learning-rate adaptation, respectively) while removing residual interference.

7.2 ABLATION STUDY ON SCHEDULER DESIGN

Our ablation study (Section 6.4) further highlights the importance of how SON-GOKU is designed. The results show that: (1) Dynamic recoloring matters. Static One-Shot underperforms the full scheduler on most metrics, indicating that task relations change enough during training that frozen groups become stale. This supports the need for periodic refresh; and (2) History smoothing is very important. Single-Step is markedly worse across datasets, consistent with our claim that per-batch cosines are too noisy to define stable groups. This aligns with our recovery analysis, which requires concentration of EMA gradients within a refresh window.

INTERPRETATION OF ABLATION STUDY RESULTS

Dynamic recoloring lets the schedule track how task interactions change over time, so partitions do not go stale. Averaging gradients over a short history makes the conflict signal less noisy, which results in more stable groups and reliable progress (Section 5.1 and 5.4). Together, these choices satisfy the conditions we use in our theory. They create low-conflict groups that ensure safe per-step descent (Section 5.1–5.2) and provide enough concentration to recover the population conflict graph within each refresh (Section 5.4).

7.3 Additional Analysis

7.3.1 OPTIMIZER-TASK ALIGNMENT

Interestingly, we observe that AdaTask-based approaches tend to be the best on classification tasks (CIFAR-10, AV-MNIST, MM-IMDb) while PCGrad-based approaches tend to be the best on tasks that model regression (NYUv2).

We believe that this stems from unique differences in the features of classification and regressionbased models. For example, cross-entropy gradients near decision boundaries tend to be bursty and high in variance (Shrivastava et al., 2016; Lin et al., 2017; Hoffer et al., 2017). By scaling each task's step size according to its running gradient norm, AdaTask smooths out these spikes.

On the other hand, we believe that PCGrad under the scheduler performs particularly well on regression and dense-prediction tasks as their tasks tend to generate smooth, large-magnitude gradients

Table 2: Wall-clock time (seconds \pm standard deviation) vs. number of tasks K.

Method (R if applicable)	K=3	K=6	K=16	K=40
Uniform	0.2656 ± 0.1201	0.3240 ± 0.0629	0.3798 ± 0.1050	0.4054 ± 0.1190
GradNorm	5.4714 ± 0.7137	5.1201 ± 0.6112	4.9042 ± 0.5869	4.7372 ± 0.9286
AdaTask	2.1816 ± 0.0934	2.1032 ± 0.1012	2.2853 ± 0.0718	2.2278 ± 0.1370
PCGrad	3.6212 ± 0.3517	23.1266 ± 0.8773	176.7566 ± 2.8171	1127.1337 ± 34.2603
MGDA	97.1081 ± 5.4645	121.4371 ± 9.0923	132.4913 ± 3.1752	134.0878 ± 2.2621
FAMO	2.0725 ± 0.2073	1.9980 ± 0.1998	2.1710 ± 0.2171	2.1164 ± 0.2116
FairGrad	3.8020 ± 0.5703	15.2079 ± 2.2812	108.1450 ± 16.2218	675.9065 ± 101.3860
Nash-MTL	5.7030 ± 1.1406	22.8118 ± 4.5624	162.2176 ± 32.4435	1013.8598 ± 202.7720
SON-GOKU ($R=32$)	1.9896 ± 0.3651	3.3202 ± 0.5745	6.0897 ± 0.9425	12.1432 ± 1.2044
SON- $GOKU + AdaTask (R = 32)$	3.7718 ± 0.9654	5.0511 ± 0.6531	7.5903 ± 1.1920	14.5182 ± 2.0660
SON- $GOKU + GradNorm (R = 32)$	7.0202 ± 1.0711	8.1661 ± 0.9355	10.7227 ± 2.2088	16.5760 ± 1.8418
SON- $GOKU + PCGrad (R = 32)$	1.9834 ± 0.3586	3.4971 ± 0.3840	6.1395 ± 0.9425	10.9097 ± 1.5263

whose directions change gradually. PCGrad removes only the small component of the gradient that conflicts across tasks, preserving the main descent direction while reducing interference.

7.3.2 SYNERGY BETWEEN SCHEDULING AND BASELINES

We believe that the superior results found in the combinations of the scheduler and baseline models can be traced to the way scheduling and optimization reinforce one another.

First, greedy graph coloring partitions tasks into τ -compatible groups, segregating tasks with highly divergent gradients. This yields a guaranteed lower bound on descent (Proposition 6), directly improving optimization efficiency.

Within each low-conflict group, the optimizer can do its job under more ideal conditions. PCGrad can remove the remaining minor conflicting components, preserving the majority of the descent direction. AdaTask can adjust each task's learning rate without being impacted by large adversarial gradients.

This $\Delta + 1$ color bound ensures that every task is scheduled at least once per period. This prevents tasks from being essentially starved of updates.

Finally, by computing interference over a window, the scheduler smooths out gradient fluctuations. This prevents the erratic schedule changes that projection-only grouping methods have been shown to face (Yu et al., 2020; Shi et al., 2023; Zhang et al., 2024), thereby better stabilizing convergence.

7.4 SPEED AND TRADEOFFS

The proposed scheduler has a time complexity of $\Theta(K^2d/R)$ amortized per training step (Section 4.5). Table 2 shows near-linear growth over this range of K at K=32, reflecting sparsity in the graphs and batched cosine computation. SON-GOKU's time rises from around 2 seconds (K=3) to 12 seconds (K=40), remaining far below methods that perform heavy per-step conflict handling. For example, PCGrad, FairGrad, and Nash-MTL increase steeply with K. In contrast, FAMO and AdaTask are among the fastest and largely flat with K, as expected from their constant overhead.

These contrasts demonstrate the tradeoffs between speed and fidelity to task interference. Faster methods like FAMO minimize overhead, while methods that model conflicts can improve accuracy. These tradeoffs have to be assessed on a case-by-case basis, based on values that factor into each approach's time complexity and the importance of training speed versus performance on an application.

8 Conclusion

We introduced SON-GOKU, an interference-aware scheduler that estimates cross-task alignment, builds a sparse conflict graph, and greedily colors it to activate one low-conflict group per step. Formally, we provide rigorous theoretical guarantees that justify the design and effectiveness of the scheduler. Empirically, across six benchmarks, SON-GOKU improves over strong baselines and recent approaches. It complements optimizers like PCGrad and AdaTask, indicating that scheduling and gradient shaping are synergistic. By modeling task interactions with a conflict graph and schedule, SON-GOKU offers a simple, scalable, and theory-backed mechanism for robust multitask training.

REFERENCES

- Rie Kubota Ando and Tong Zhang. A framework for learning predictive structures from multiple tasks and unlabeled data. In *Proceedings of the 24th International Conference on Machine Learning (ICML)*, pp. 3–10, 2005.
- John Arevalo, Thamar Solorio, Manuel Montes y Gómez, and Fabio A. González. Gated multimodal units for information fusion, 2017. URL https://arxiv.org/abs/1702.01992.
 - Andreas Argyriou, Theodoros Evgeniou, and Massimiliano Pontil. Multi-task feature learning. In *Advances in Neural Information Processing Systems*, volume 19, pp. 41–48, 2007.
 - Afiya Ayman, Ayan Mukhopadhyay, and Aron Laszka. Task grouping for automated multi-task machine learning via task affinity prediction, 2023. URL https://arxiv.org/abs/2310.16241.
 - Hao Ban and Kaiyi Ji. Fair resource allocation in multi-task learning. In *Proceedings of the 41st International Conference on Machine Learning*, ICML'24. JMLR.org, 2024.
 - Amir Beck. First-order methods in optimization. SIAM, 2017.
 - Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Springer, New York, NY, 2006. ISBN 978-0387310732.
 - Marthe Bonamy, Tom Kelly, Peter Nelson, and Luke Postle. Bounding χ by a fraction of δ for graphs without large cliques, 2018. URL https://arxiv.org/abs/1803.01051.
 - Thomas Borsani, Andrea Rosani, Giuseppe Nicosia, and Giuseppe Di Fatta. Gradient similarity surgery in multi-task deep learning. *arXiv preprint arXiv:2506.06130*, 2025.
 - Léon Bottou, Frank E Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *SIAM review*, 60(2):223–311, 2018.
 - Said Yacine Boulahia, Abdenour Amamra, Mohamed Ridha Madi, and Said Daikh. Early, intermediate and late fusion strategies for robust deep learning-based multimodal action recognition. *Machine Vision and Applications*, 32(6):121, 2021.
 - Rich Caruana. Multitask learning. *Machine Learning*, 28(1):41–75, July 1997. doi: 10.1023/A: 1007379606734. URL https://doi.org/10.1023/A:1007379606734.
 - Zhao Chen, Vijay Badrinarayanan, Chen-Yu Lee, and Andrew Rabinovich. Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In *International Conference on Machine Learning*, pp. 794–803, 2018. URL https://arxiv.org/abs/1711.02257.
 - Zhiyong Cui, Ruimin Ke, and Yinhai Wang. Deep bidirectional and unidirectional LSTM recurrent neural network for network-wide traffic speed prediction. *CoRR*, abs/1801.02143, 2018. URL http://arxiv.org/abs/1801.02143.
 - Sanjoy Dasgupta and Anupam Gupta. An elementary proof of a theorem of johnson and lindenstrauss. In *Randomization and Approximation Techniques in Computer Science*, RANDOM 2003, pp. 53–62. Springer, 2003. doi: 10.1007/978-3-540-45198-3_4.
 - Victor H De la Pena, Tze Leung Lai, and Qi-Man Shao. Self-normalized processes: Limit theory and Statistical Applications. Springer, 2009.
 - Reinhard Diestel. Graph Theory. Springer, 5th edition, 2017. ISBN 978-3-662-53622-3.
 - Paul Erdős and Alfréd Rényi. On the evolution of random graphs. *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, 5:17–61, 1960.
- Theodoros Evgeniou, Cinzia A. Micchelli, and Massimiliano Pontil. Learning multiple tasks with kernels. *Journal of Machine Learning Research*, 6:615–637, 2005.
 - Caoyun Fan, Wenqing Chen, Jidong Tian, Yitian Li, Hao He, and Yaohui Jin. Maxgnr: A dynamic weight strategy via maximizing gradient-to-noise ratio for multi-task learning. *arXiv* preprint *arXiv*:2302.09352, 2023.

- Christopher Fifty, Ehsan Amid, Zhe Zhao, Tianhe Yu, Rohan Anil, and Chelsea Finn. Efficiently identifying task groupings for multi-task learning, 2021. URL https://arxiv.org/abs/2109.04617.
 - Jorge Fliege and Benar F. Svaiter. Steepest descent methods for multicriteria optimization. *Mathematical Methods of Operations Research*, 51(3):479–494, 2000. doi: 10.1007/s001860000043.
 - Saeed Ghadimi and Guanghui Lan. Stochastic first- and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013. doi: 10.1137/120880811.
 - Mina Ghashami, Edo Liberty, Jeff M. Phillips, and David P. Woodruff. Frequent directions: Simple and deterministic matrix sketching. *SIAM Journal on Computing*, 45(5):1762–1792, 2016. doi: 10.1137/15M1009718.
 - Valerio Guarrasi, Fatih Aksu, Camillo Maria Caruso, Francesco Di Feola, Aurora Rofena, Filippo Ruffini, and Paolo Soda. A systematic review of intermediate fusion in multimodal deep learning for biomedical applications. *Image and Vision Computing*, pp. 105509, 2025.
 - Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015. URL http://arxiv.org/abs/1512.03385.
 - Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *Proceedings of the International Conference on Learning Representations*, 2019.
 - Elad Hoffer, Itay Hubara, and Daniel Soudry. Train longer, generalize better: Closing the generalization gap in large batch training of neural networks. In *Advances in Neural Information Processing Systems*, volume 30, pp. 1731–1741, 2017.
 - Wooseong Jeong and Kuk-Jin Yoon. Selective task group updates for multi-task optimization, 2025.
 - Bo Kågström, Per Ling, and Charles Van Loan. Gemm-based level 3 blas: high-performance model implementations and performance evaluation benchmark. *ACM Transactions on Mathematical Software (TOMS)*, 24(3):268–302, 1998.
 - Zhuoliang Kang, Kristen Grauman, and Fei Sha. Learning with whom to share in multi-task feature learning. In *Proceedings of the 28th International Conference on Machine Learning (ICML)*, pp. 521–528, 2011.
 - Hamed Karimi, Julie Nutini, and Mark Schmidt. Linear convergence of gradient and proximal-gradient methods under the polyak-łojasiewicz condition. In *Joint European conference on machine learning and knowledge discovery in databases*, pp. 795–811. Springer, 2016.
 - Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 7482–7491, 2018. doi: 10.1109/CVPR.2018.00781.
 - Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
 - Vitaly Kurin, Alessandro De Palma, Ilya Kostrikov, Shimon Whiteson, and M. Pawan Kumar. In defense of the unitary scalarization for deep multi-task learning. In *Advances in Neural Information Processing Systems*, volume 35, 2022. URL https://arxiv.org/abs/2201.04122.
 - Harold J Kushner and G George Yin. Stochastic approximation and recursive algorithms and applications. Springer, 2003.
 - Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. doi: 10.1109/5.726791.
 - R Gary Leonard and George Doddington. Tidigits speech corpus. Texas Instruments, Inc., 1993.
 - Xiangru Lian, Yijun Huang, Yuncheng Li, and Ji Liu. Asynchronous parallel stochastic gradient for nonconvex optimization. *Advances in neural information processing systems*, 28, 2015.

- Paul Pu Liang, Yiwei Lyu, Xiang Fan, Zetian Wu, Yun Cheng, Jason Wu, Leslie Chen, Peter Wu, Michelle A Lee, Yuke Zhu, et al. Multibench: Multiscale benchmarks for multimodal representation learning. *Advances in neural information processing systems*, 2021(DB1):1, 2021.
 - Sicong Liang and Yu Zhang. A simple general approach to balance task difficulty in multi-task learning, 2020. URL https://arxiv.org/abs/2002.04792.
 - Edo Liberty. Simple and deterministic matrix sketching. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pp. 581–588, 2013. doi: 10.1145/2487575.2487623.
 - Baijiong Lin, Feiyang Ye, and Yu Zhang. A closer look at loss weighting in multi-task learning, 2022. URL https://openreview.net/forum?id=OdnNBNIdFul.
 - Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pp. 2999–3007, 2017.
 - Bo Liu, Xingchao Liu, Xiaojie Jin, Peter Stone, and Qiang Liu. Conflict-averse gradient descent for multi-task learning. In *Advances in Neural Information Processing Systems*, volume 34, pp. 12345–12355, 2021.
 - Bo Liu, Yihao Feng, Peter Stone, and Qiang Liu. Famo: Fast adaptive multitask optimization. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), Advances in Neural Information Processing Systems, volume 36, pp. 57226–57243. Curran Associates, Inc., 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/b2felee8d936ac08dd26f2ff58986c8f-Paper-Conference.pdf.
 - 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*, pp. 1871–1880, 2019.
 - Y. Liu. Theoretical analysis on how learning rate warmup accelerates gradient descent. *arXiv preprint arXiv:2509.07972*, 2025. URL https://arxiv.org/abs/2509.07972.
 - Ben Lockwood. Pareto efficiency. In *The new Palgrave dictionary of economics*, pp. 1–5. Springer, 2008.
 - László Lovász. Graph minor theory. *Bulletin of the American Mathematical Society*, 43(1):75–86, 2006.
 - Aakarsh Malhotra, Mayank Vatsa, and Richa Singh. Dropped scheduled task: Mitigating negative transfer in multi-task learning using dynamic task dropping. *Transactions on Machine Learning Research*, 2022.
 - Bryan McCann, Nitish Shirish Keskar, Caiming Xiong, and Richard Socher. The natural language decathlon: Multitask learning as question answering. *CoRR*, abs/1806.08730, 2018. URL https://arxiv.org/abs/1806.08730.
 - Florence Merlevède, Magda Peligrad, and Emmanuel Rio. Bernstein inequality and moderate deviations under strong mixing conditions. *High Dimensional Probability VI*, pp. 273–292, 2011.
 - Kaisa Miettinen. *Nonlinear Multiobjective Optimization*. Kluwer Academic Publishers, Boston, MA, 1999. ISBN 978-0792382781.
 - Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space, 2013. URL https://arxiv.org/abs/1301.3781.
 - Salman Mohammadi, Anders Kirk Uhrenholt, and Bjørn Sand Jensen. Odd-one-out representation learning. *arXiv preprint arXiv:2012.07966*, 2020. Shows that distinguishing an "odd" element among "even" ones in auxiliary pretext tasks yields stronger embeddings.
 - MSCI Inc. and S&P Dow Jones Indices. *Global Industry Classification Standard (GICS)*. MSCI Inc., New York, NY, august 2024 edition, 2024. First published January 7, 2020; updated August 2024.

- David Mueller, Mark Dredze, and Nicholas Andrews. The importance of temperature in multi-task optimization. In *OPT 2022: Optimization for Machine Learning (NeurIPS 2022 Workshop)*, 2022. URL https://openreview.net/forum?id=H9UOWMR_Ut.
 - Aviv Navon, Aviv Shamsian, Idan Achituve, Haggai Maron, Kenji Kawaguchi, Gal Chechik, and Ethan Fetaya. Multi-task learning as a bargaining game. *arXiv preprint arXiv:2202.01017*, 2022.
 - Arkadi Nemirovski, Anatoli Juditsky, Guanghui Lan, and Alexander Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–1609, 2009. doi: 10.1137/070704277.
 - Yurii Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Springer, 2004. ISBN 978-1-4419-8853-9.
 - Feng Niu, Benjamin Recht, Christopher Ré, and Stephen J. Wright. Hogwild!: A lock-free approach to parallelizing stochastic gradient descent. In *Advances in Neural Information Processing Systems*, volume 24, pp. 693–701, 2011.
 - Vilfredo Pareto. Manual of political economy: a critical and variorum edition. OUP Oxford, 2014.
 - Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pp. 400–407, 1951.
 - Sebastian Ruder. An overview of multi-task learning in deep neural networks. *arXiv preprint* arXiv:1706.05098, 2017. URL https://arxiv.org/abs/1706.05098.
 - Ozan Sener and Vladlen Koltun. Multi-task learning as multi-objective optimization. In *Advances in Neural Information Processing Systems*, volume 31, pp. 525–536, 2018.
 - Ammar Sherif, Abubakar Abid, Mustafa Elattar, and Mohamed ElHelw. Stg-mtl: scalable task grouping for multi-task learning using data maps. *Machine Learning: Science and Technology*, 5(2):025068, June 2024. ISSN 2632-2153. doi: 10.1088/2632-2153/ad4e04. URL http://dx.doi.org/10.1088/2632-2153/ad4e04.
 - Guangyuan Shi, Qimai Li, Wenlong Zhang, Jiaxin Chen, and Xiao-Ming Wu. Recon: Reducing conflicting gradients from the root for multi-task learning. In *ICLR 2023 Workshop on Multi-Task Learning*, 2023.
 - Abhinav Shrivastava, Abhinav Gupta, and Ross Girshick. Training region-based object detectors with online hard example mining. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 761–769, 2016.
 - Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *European Conference on Computer Vision (ECCV)*, pp. 746–760. Springer, 2012.
 - Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
 - Trevor Standley, Amir R. Zamir, Dawn Chen, Leonidas Guibas, Jitendra Malik, and Silvio Savarese. Which tasks should be learned together in multi-task learning? In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, pp. 9120–9132, 2020.
 - J Michael Steele. *The Cauchy-Schwarz master class: an introduction to the art of mathematical inequalities.* Cambridge University Press, 2004.
 - Swabha Swayamdipta, Roy Schwartz, Nicholas Lourie, Yizhong Wang, Hannaneh Hajishirzi, Noah A. Smith, and Yejin Choi. Dataset cartography: Mapping and diagnosing datasets with training dynamics. In Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 9275–9293, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020. emnlp-main.746. URL https://aclanthology.org/2020.emnlp-main.746/.

- Andrea Tacchetti, Stephen Voinea, and Georgios Evangelopoulos. Trading robust representations for sample complexity through self-supervised visual experience. In *Advances in Neural Information Processing Systems*, volume 31, pp. 1686–1696, 2018. Section 4.2 demonstrates a "Transfer learning: even/odd MNIST" auxiliary task that boosts few-shot performance.
 - Lovre Torbarina, Tin Ferkovic, Lukasz Roguski, Velimir Mihelcic, Bruno Sarlija, and Zeljko Kraljevic. Challenges and opportunities of using transformer-based multi-task learning in nlp through ml lifecycle: A survey, 2023. URL https://arxiv.org/abs/2308.08234.
 - 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. doi: 10.1109/TPAMI.2021. 3054719.
 - Valentin Vielzeuf, Alexis Lechervy, Stéphane Pateux, and Frédéric Jurie. Centralnet: a multilayer approach for multimodal fusion. *CoRR*, abs/1808.07275, 2018. URL http://arxiv.org/abs/1808.07275.
 - Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge university press, 2019.
 - Chenguang Wang, Xuanhao Pan, and Tianshu Yu. Towards principled task grouping for multi-task learning. *arXiv preprint arXiv:2402.15328*, 2024.
 - D. J. A. Welsh and M. B. Powell. An upper bound for the chromatic number of a graph and its application to timetabling problems. *The Computer Journal*, 10(1):85–86, 01 1967. ISSN 0010-4620. doi: 10.1093/comjnl/10.1.85. URL https://doi.org/10.1093/comjnl/10.1.85.
 - Douglas B. West. *Introduction to Graph Theory*. Prentice Hall, 2nd edition, 2000. ISBN 978-0130144003.
 - Enneng Yang, Junwei Pan, Ximei Wang, Haibin Yu, Li Shen, Xihua Chen, Lei Xiao, Jie Jiang, and Guibing Guo. Adatask: A task-aware adaptive learning rate approach to multi-task learning. In *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence (AAAI)*, 2023. URL https://arxiv.org/abs/2211.15055.
 - Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. Gradient surgery for multi-task learning. In *Advances in Neural Information Processing Systems*, volume 33, pp. 18524–18536, 2020.
 - Wenxin Yu, Xueling Shen, Jiajie Hu, and Dong Yin. Revisiting the loss weight adjustment in object detection. *arXiv preprint arXiv:2103.09488*, 2021.
 - Amir R. Zamir, Alexander Sax, Teresa Yeo, Oguzhan Kar, Nikhil Cheerla, Rohan Suri, Zhangjie Cao, Jitendra Malik, and Leonidas Guibas. Robust learning through cross-task consistency. *CoRR*, abs/2006.04096, 2020. URL https://arxiv.org/abs/2006.04096.
 - Yu Zhang and Qiang Yang. An overview of multi-task learning. *National Science Review*, 5(1): 30–43, 2018.
 - Zhi Zhang, Jiayi Shen, Congfeng Cao, Gaole Dai, Shiji Zhou, Qizhe Zhang, Shanghang Zhang, and Ekaterina Shutova. Proactive gradient conflict mitigation in multi-task learning: A sparse training perspective. arXiv preprint arXiv:2411.18615, 2024. URL https://arxiv.org/abs/2411.18615.
 - Han Zhao, Yifan Guo, Aleksandar Risteski, et al. Robust multi-task learning with excess risks. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, 2024.
 - Mo Zhou, Rong Ge, and Chi Jin. A local convergence theory for mildly over-parameterized two-layer neural network. In *Conference on Learning Theory*, pp. 4577–4632. PMLR, 2021.

A FULL ALGORITHM BLOCK FOR PROPOSED APPROACH

Algorithm 1 SON-GOKU Scheduler

757 758

759

791 792

793

794

796

797798

799 800

801

802

804

808

809

```
760
            Require: Initial shared params \theta_0, heads \{\phi_k\}_{k=1}^K, EMA buffers \tilde{g}_k^{(0)} = 0, total steps T, learning-rate schedule \{\eta_t\}, refresh length R, warm-up T_{\text{warm}}, target threshold \tau^*, minimum coverage f_{\min},
761
762
                  EMA parameter \beta
              1: Gradients follow the weighted-loss convention (Sec. 4).
764
              2: r \leftarrow 0, t_r \leftarrow 0
                                                                                                     > current refresh round and start index
             3: \tau \leftarrow 1; m_0 \leftarrow 1; C_1^{(0)} \leftarrow \{1, \dots, K\}
4: for t = 0, \dots, T - 1 do
765
                                                                                                                            766
767
                        Warm-up/Anneal: \tau \leftarrow \text{ANNEAL}(t)

    □ approach in Sec. 4.4

                       Scheduling: S_t \leftarrow C_{(t \bmod m_r)+1}^{(r)}
Forward/Backward:
768
              6:
769
              7:
770
                        for all k \in S_t do
              8:
771
                             compute per-task gradients g_k^{(t)} and h_k^{(t)} (defs: Sec. 4.1)
              9:
772
             10:
773
                        Parameter update (shared): \theta_{t+1} \leftarrow \theta_t - \eta_t \sum_{k \in S_t} g_k^{(t)}
             11:
774
            12:
                        Parameter update (task-specific):
775
            13:
                        for all k \in S_t do
                             \phi_{k,t+1} \leftarrow \phi_{k,t} - \eta_t h_{\iota}^{(t)}
776
            14:
                        end for
777
            15:
                        EMA:
            16:
778
                       \begin{array}{l} \mbox{ for all } k \in S_t \mbox{ do } \\ \mbox{ update } \tilde{g}_k^{(t+1)} \mbox{ (Eq. 8)} \end{array}
            17:
779
            18:
                        end for
            19:
781
            20:
                        if (t+1) \mod R = 0 then

⊳ refresh

782
                             EMA refresh: update all \tilde{g}_i using small mini-batches (Sec. 4.1)
            21:
783
                             Interference matrix: compute \rho_{ij}^{(t+1)} via Eq. 9
            22:
784
                              Conflict graph: build G_{\tau}^{(r+1)} via Eq. 10
785
            23:
                             Greedy coloring: Welsh–Powell \rightarrow \{C_1^{(r+1)}, \dots, C_{m_{r+1}}^{(r+1)}\}
786
            24:
787
                             Minimum coverage: enforce f_i \ge f_{\min} using compatible-slot duplication (Sec. 4.4.1)
            25:
788
            26:
                             r \leftarrow r + 1; t_r \leftarrow t + 1
                        end if
789
            27:
            28: end for
790
```

Algorithm block 1 provides an overview of the SON-GOKU scheduler. At a high level, the procedure consists of four stages: (1) estimating pairwise interference, (2) building and coloring the conflict graph, (3) generating a periodic schedule, and (4) updating that schedule as training evolves.

B EXACT RECOVERY OF POPULATION CONFLICT GRAPH & TASK PARTITION

B.1 SETTING, DEFINITIONS, AND POPULATION OBJECTS

Let $K \geq 2$ be the number of tasks and $d \geq 1$ the parameter dimension. At designated refresh iterations, the scheduler:

- (i) computes a per-task exponential moving average (EMA) of stochastic gradients over a probe window of R iterations,
- (ii) forms a cosine-similarity matrix from the K EMA vectors,
- (iii) builds a conflict graph by thresholding negative cosines at a fixed level $-\tau$ with $\tau \in (0,1)$,
- (iv) computes a proper coloring of the conflict graph, and
- (v) schedules one color class per iteration until the next refresh

Definition B.1. At the beginning of a refresh window (i.e., at a fixed iterate θ), let

$$\mu_i \in \mathbb{R}^d \qquad (i = 1, \dots, K) \tag{14}$$

denote the population task gradients (or the window-stationary means). Define the population cosine matrix $C^* \in [-1,1]^{K \times K}$ by

$$C_{ij}^{\star} = \frac{\langle \mu_i, \mu_j \rangle}{\|\mu_i\| \|\mu_j\|}, \qquad i \neq j, \quad C_{ii}^{\star} = 1.$$
 (15)

Definition B.2. Fix $\tau \in (0,1)$. The population conflict graph $G^* = (V, E^*)$ on vertex set $V = \{1, \ldots, K\}$ has an edge $\{i, j\}$ iff $C_{ij}^* < -\tau$. The true grouping \mathcal{P}^* is one of:

(A) Component Model: the vertex partition given by the connected components of G^* .

 (B) Multipartite model: a partition $V = \bigsqcup_{r=1}^m P_r$ (with $m \ge !$) such that G^* is the complete m-partite graph induced by $\{P_r\}_{r=1}^m$ (no edges within any P_r , all cross-part edges present)

When we later speak of group recovery, we mean equality of the empirical partition (defined from data) with \mathcal{P}^* , up to label permutation in case (B).

B.2 ASSUMPTIONS

We adopt the following assumptions, which are standard in analyses of stochastic-gradient methods and verifiable in practice (see, e.g., Robbins & Monro 1951; Kushner & Yin 2003; Nemirovski et al. 2009; Bottou et al. 2018; Wainwright 2019; for concentration of geometrically weighted and mixing sequences, see Merlevède et al. 2011; De la Pena et al. 2009).

Assumption 1 (Separation margin around the threshold). There exists $\gamma \in (0, 1 - \tau)$ such that for all $i \neq j$:

$$\begin{cases} C_{ij}^{\star} \leq -(\tau + \gamma), & \text{if } i \text{ and } j \text{ lie in different groups of } \mathcal{P}^{\star}, \\ C_{ij}^{\star} \geq -(\tau - \gamma), & \text{if } i \text{ and } j \text{ lie in the same group of } \mathcal{P}^{\star}. \end{cases}$$
(16)

Assumption 2 (Probe noise model and EMA). *In the refresh window of length R, the per-iteration stochastic task gradients admit the decomposition*

$$g_{i,t} = \mu_i + \xi_{i,t}, \qquad t = 1, \dots, R,$$
 (17)

where $\{\xi_{i,t}\}_{t=1}^R$ are mean-zero, sub-Gaussian with parameter σ^2 , and satisfy a ϕ -mixing or martingale-difference condition ensuring concentration with geometric weights. The EMA for task i is

$$\tilde{g}_i = \sum_{t=1}^R w_t g_{i,t}, \qquad w_t = \frac{(1-\beta)\beta^{R-t}}{1-\beta^R}, \qquad \beta \in [0,1).$$
 (18)

Define the effective sample size $n_{\rm eff}$ by

$$n_{\text{eff}}^{-1} := \sum_{t=1}^{R} w_t^2 = \frac{(1-\beta)^2 (1-\beta^{2R})}{(1-\beta^R)^2 (1-\beta^2)}.$$
 (19)

In particular, as $R \to \infty$ (with fixed $\beta \in [0,1)$), we have $n_{\text{eff}} \to \frac{1+\beta}{1-\beta}$

Assumption 3 (Slow drift within a refresh). Over the refresh window, the changes in μ_i are small enough to be absorbed in the concentration bounds below (equivalently, one can regard μ_i as constant within the window by working at the start-of-window iterate and moving any drift into the noise process).

Assumption 4 (Minimum norm and task inclusion). There exists $m_0 > 0$ such that $\|\mu_i\| \ge m_0$ for all tasks included in the graph. In our implementation, we make it so that tasks with $\|\tilde{g}_i\| < \nu$ (for a small $\nu \ll m_0$) are temporarily excluded from graph construction until stabilized.

Assumption 5 (Threshold selection). The threshold τ is fixed across refreshes or selected using data independent of the probe window used to form $\{\tilde{g}_i\}$ (e.g., via a separate pilot set). The analysis below treats τ as deterministic with respect to the probe sample.

B.3 Deterministic group recovery from the conflict graph

We begin with basic graph-theoretic facts that we will use once we have established that the empirical conflict graph coincides with its population counterpart.

Proposition 1 (Chromatic number of a complete multipartite graph). If G^* is complete m-partite with parts $\{P_r\}_{r=1}^m$, then $\chi(G^*)=m$.

Proof. Picking one vertex from each part yields a clique of size m, hence $\chi(G^*) \geq m$. Coloring each part with a distinct color is proper, hence $\chi(G^*) \leq m$. Therefore $\chi(G^*) = m$.

Theorem 1 (Identifiability via optimal coloring under model (B)). Assume model (B), i.e., G^* is complete m-partite with parts $\{P_r\}_{r=1}^m$. Let $c: V \to \{1, \ldots, m\}$ be a proper coloring of G^* that uses exactly $\chi(G^*)$ colors. Then each color class equals some part P_r (up to relabeling).

Proof. In a complete multipartite graph, any two vertices from different parts are adjacent. Thus, no color class can contain vertices from two different parts, so each color class is contained in some P_r . By Proposition 1, $\chi(G^\star)=m$, so any optimal coloring uses exactly m colors. Since there are m nonempty parts, none can be split across two colors. Hence, the color classes coincide with $\{P_r\}_{r=1}^m$ up to permutation. \square

Proposition 2 (Identifiability via components under model (A)). Under model (A), the grouping \mathcal{P}^* equals the connected components of G^* . Consequently, any procedure that returns the connected components of the empirical graph recovers \mathcal{P}^* whenever the empirical graph equals G^* .

B.4 Uniform control of empirical cosines from EMA gradients

We now quantify the deviation of the empirical cosine matrix \hat{C} formed from $\{\tilde{g}_i\}$ relative to C^* .

Lemma 1 (EMA vector concentration in directions of interest). Assume Assumption 2 and Assumption 3. There exists a constant c > 0 depending only on the mixing parameters such that for any fixed unit vector $u \in \mathbb{S}^{d-1}$ and any $\varepsilon > 0$.

$$\Pr(\left|\langle \tilde{g}_i - \mu_i, u \rangle\right| > \varepsilon) \le 2 \exp(-c n_{\text{eff}} \varepsilon^2 / \sigma^2).$$
 (20)

In particular, for any finite set of unit vectors $\{u_j\}_{j=1}^M$, a union bound yields

$$\Pr\left(\max_{1 \le j \le M} \left| \langle \tilde{g}_i - \mu_i, u_j \rangle \right| > \varepsilon \right) \le 2M \exp\left(-c n_{\text{eff}} \varepsilon^2 / \sigma^2 \right). \tag{21}$$

Proof. The scalar process $\{\langle \xi_{i,t}, u \rangle\}_{t=1}^R$ is sub-Gaussian with variance proxy σ^2 and satisfies the same mixing condition. Exponential-weighted averages of such sequences obey Hoeffding-Azuma/Berstein-type tail bounds with variance proxy $\sigma^2 \sum_t w_t^2 = \sigma^2/n_{\rm eff}$. The stated inequality follows.

Lemma 2 (Cosine stability under perturbations). Assume Assumption 4 and let $\epsilon > 0$. If for a pair (i, j) we have

$$\left| \langle \tilde{g}_{i} - \mu_{i}, \frac{\mu_{j}}{\|\mu_{j}\|} \rangle \right| \leq \epsilon, \qquad \left| \langle \tilde{g}_{j} - \mu_{j}, \frac{\mu_{i}}{\|\mu_{i}\|} \rangle \right| \leq \epsilon, \qquad \left| \langle \tilde{g}_{i} - \mu_{i}, \frac{\mu_{i}}{\|\mu_{i}\|} \rangle \right| \leq \epsilon, \qquad \left| \langle \tilde{g}_{j} - \mu_{j}, \frac{\mu_{j}}{\|\mu_{j}\|} \rangle \right| \leq \epsilon,$$

$$(22)$$

then

$$\left|\widehat{C}_{ij} - C_{ij}^{\star}\right| \leq \frac{6\,\epsilon}{m_0} + \frac{4\,\epsilon^2}{m_0^2}.\tag{23}$$

Proof. Write $\tilde{g}_i = \mu_i + \delta_i$, $\tilde{g}_j = \mu_j + \delta_j$. Decompose the numerator and denominator in the cosine:

$$\langle \tilde{g}_i, \tilde{g}_j \rangle - \langle \mu_i, \mu_j \rangle = \langle \delta_i, \mu_j \rangle + \langle \mu_i, \delta_j \rangle + \langle \delta_i, \delta_j \rangle, \tag{24}$$

and

$$\|\tilde{g}_i\| = \|\mu_i\| \sqrt{1 + 2\langle \delta_i, \mu_i \rangle / \|\mu_i\|^2 + \|\delta_i\|^2 / \|\mu_i\|^2}$$
(25)

Using Assumption 4,

$$|\langle \delta_i, \mu_i / || \mu_i || \rangle| \le \epsilon \tag{26}$$

and

imply

and

$$|\langle \delta_i, \mu_i / || \mu_i || \rangle| \le \epsilon \tag{27}$$

 $|\langle \delta_i, \mu_i \rangle| \leq \epsilon \|\mu_i\|$ (28)

 $|\langle \delta_i, \mu_i \rangle| \le \epsilon \|\mu_i\|$ (29)

A second-order expansion of the cosine in (δ_i, δ_j) with the above controls yields the bound. The constants 6 and 4 arise from collecting the linear and quadratic contributions in ϵ/m_0 .

Combining Lemma 1 and Lemma 2 with a union bound over all unordered pairs (i, j) shows that the empirical cosines are uniformly close to their population counterparts.

Proposition 3 (Uniform cosine accuracy with high probability). Assume Assumption 2, Assumption 3, and Assumption 4. For any $\epsilon > 0$ there exist absolute constants c, C > 0 such that if

$$n_{\text{eff}} \ge C \frac{\sigma^2}{m_0^2 \epsilon^2} \tag{30}$$

then, with probability $1 - \delta$,

$$\max_{i < j} \left| \widehat{C}_{ij} - C_{ij}^{\star} \right| \le \epsilon \tag{31}$$

Proof. For each unordered pair (i, j), apply Lemma 1 with the four unit vectors $\mu_j / \|\mu_j\|$, $\mu_i / \|\mu_i\|$, and use Lemma 2 to convert these directional deviations into a cosine deviation bound. A union bound over the $O(K^2)$ pairs yields the claimed logarithmic factor. The constants absorb the quadratic term in ϵ by requiring $\epsilon \leq m_0$.

B.5 EXACT EDGE RECOVERY AND GROUP RECOVERY

We first show that a uniform cosine error smaller than the margin γ implies exact equality of empirical and population conflict graphs.

Theorem 2 (Exact conflict-graph recovery under the margin). Assume Assumptions 1–5. If

$$\max_{i < j} \left| \hat{C}_{ij} - C_{ij}^{\star} \right| \le \epsilon \quad \text{with} \quad \epsilon < \gamma, \tag{32}$$

then the empirical conflict graph equals the population graph:

$$\widehat{G} = G^{\star}. \tag{33}$$

Equivalently, for every $i \neq j$,

$$C_{ij}^{\star} \le -(\tau + \gamma) \Rightarrow \widehat{C}_{ij} < -\tau \quad and \quad C_{ij}^{\star} \ge -(\tau - \gamma) \Rightarrow \widehat{C}_{ij} > -\tau.$$
 (34)

Proof. For any pair
$$(i,j)$$
, if $C_{ij}^{\star} \leq -(\tau+\gamma)$, then $\widehat{C}_{ij} \leq -(\tau+\gamma) + \epsilon < -\tau$, hence $\{i,j\} \in \widehat{E}$. If $C_{ij}^{\star} \geq -(\tau-\gamma)$, then $\widehat{C}_{ij} \geq -(\tau-\gamma) - \epsilon > -\tau$, hence $\{i,j\} \notin \widehat{E}$.

Combining Proposition 3 and Theorem 2 yields a high-probability statement.

Corollary 1 (High-probability exact recovery of G^*). Under Assumptions 1–5, there exists a universal constant C > 0 such that if

$$n_{\text{eff}} \ge C \frac{\sigma^2}{m_0^2 \gamma^2} \log \left(\frac{K^2}{\delta}\right),$$
 (35)

then $\Pr(\widehat{G} = G^*) \ge 1 - \delta$.

Theorem 3 (Group recovery under the component model). *Under model* (A) and the conditions of Corollary 1, with probability at least $1 - \delta$, the connected components of \widehat{G} equal \mathcal{P}^* .

Proof. Immediate from $\widehat{G} = G^*$ and the definition of \mathcal{P}^* .

Theorem 4 (Group recovery under the multipartite model). Under model (B) and the conditions of Corollary 1, with probability at least $1 - \delta$, $\chi(\widehat{G}) = m$ and any optimal coloring of \widehat{G} yields color classes equal to $\{P_r\}_{r=1}^m$ up to label permutation.

Proof. If $\widehat{G} = G^*$, then \widehat{G} is complete m-partite. Proposition 1 gives $\chi(\widehat{G}) = m$. Theorem 1 implies identifiability up to permutation by any optimal coloring.

B.6 QUANTITATIVE PROBE-BUDGET REQUIREMENT

Combining the bounds above yields the following sample-complexity statement.

Corollary 2. Under assumptions 1–5, there exist absolute constants c, C > 0 such that the following holds. If the EMA parameters (R, β) are chosen to ensure

$$n_{\text{eff}} \ge C \frac{\sigma^2}{m_0^2 \gamma^2} \qquad \left(equivalently, \quad \sum_{t=1}^R w_t^2 \le c \frac{m_0^2 \gamma^2}{\sigma^2} \frac{1}{\log(K/\delta)} \right)$$
 (36)

then $\Pr(\widehat{G} = G^{\star}) \geq 1 - \delta$, and consequently Theorems 3–4 apply. In particular, for fixed β and large R, $n_{\text{eff}} \to \frac{1+\beta}{1-\beta}$ (i.e., it saturates). Thus, to meet the required budget as K grows, one increases n_{eff} by choosing β closer to 1 (e.g., $1 - \beta \approx 1/\log(K^2/\delta)$), or by switching to a unnormalized averaging approach.

B.7 SUMMARY OF THE RECOVERY ARGUMENT

We summarize the logical flow leading to consistency of the scheduler.

- (i) Assumptions: Assumptions 1–5 define the conditions in which in which across-group population cosines lie below $-(\tau + \gamma)$, within-group cosines lie above $-(\tau \gamma)$, EMA gradients concentrate with effective sample size $n_{\rm eff}$, and all included tasks have nonnegligible gradient norm.
- (ii) Uniform cosine accuracy: Lemmas 1–2 together with Proposition 3 yield a high-probability uniform cosine approximation:

$$\max_{i \le j} \left| \hat{C}_{ij} - C_{ij}^{\star} \right| \le \epsilon, \tag{37}$$

with probability at least $1 - \delta$, where ϵ decreases as n_{eff} increases.

(iii) Exact recovery of edges: If the approximation tolerance satisfies $\epsilon < \gamma$, Theorem 2 converts the uniform bound into exact edge recovery of the conflict graph:

$$\widehat{G} = G^{\star}. \tag{38}$$

(iv) Recovery of the grouping: Given $\widehat{G} = G^{\star}$, Theorem 3 implies group recovery under the component model (groups are the connected components). Under the multipartite model, Proposition 1 and Theorem 1 yield $\chi(\widehat{G}) = m$ and Theorem 4 shows that any optimal coloring returns the true parts (up to label permutation).

Quantitative consequence. Assume Assumptions 1–5 and fix $\delta \in (0,1)$. Let $m_0 = \min_i \|\mu_i\|$ and let σ^2 be the variance proxy from Assumption 2. If the EMA probe budget satisfies

$$n_{\text{eff}} \ge C \frac{\sigma^2}{m_0^2 \gamma^2} \log \left(\frac{K^2}{\delta} \right)$$
 (39)

for a universal constant C > 0, then with probability at least $1 - \delta$ the empirical conflict graph equals the population graph: $\hat{G} = G^*$. Consequently:

- (i) under the component model (A), the connected components of \widehat{G} coincide with \mathcal{P}^{\star} .
- (ii) under the multipartite model (B), $\chi(\widehat{G})=m$ and any optimal coloring of \widehat{G} recovers \mathcal{P}^{\star} up to permutation of labels.

C DESCENT BOUNDS FOR SCHEDULED VERSUS AGGREGATED UPDATES

We compare two update procedures over a single refresh: a scheduled sequence of per-group steps (i.e., the approach we propose in our paper) and a single aggregated step that combines all groups at once. Both use the same step size η and the same gradient information measured at the start of the refresh, and our analysis operates at the level of L-smooth (descent) upper bounds. We identify when the scheduled bound is strictly tighter and summarize implications under PL / strong convexity.

Throughout, $F: \mathbb{R}^d \to \mathbb{R}$ is differentiable and L-smooth, i.e.

$$F(y) \le F(x) + \langle \nabla F(x), y - x \rangle + \frac{L}{2} |y - x|^2, \quad \forall x, y.$$
 (40)

We write $\nabla F(x) = \sum_{r=1}^{m} G_r(x)$, where each $G_r(x)$ is the group gradient for color r (any fixed linear aggregator of task gradients assigned to color r for the current refresh). We use a refresh step size $\eta \in (0, 1/L]$.

C.1 SINGLE REFRESH BASELINES AND NOTATION

C.1.1 SINGLE AGGREGATED STEP

Definition C.1 (Aggregated step). Starting from the same point x, with step size $\eta \in (0, 1/L]$ and group gradients $G_r^0 := G_r(x)$ (with $\nabla F(x) = \sum_{r=1}^m G_r^0$), define

$$x^{\text{agg}} := x - \eta \sum_{r=1}^{m} G_r^0.$$
 (41)

One-shot L-smoothness bound. Applying L-smoothness with $y = x^{agg}$ yields

$$F(x^{\text{agg}}) \le F(x) - \eta \left\langle \nabla F(x), \sum_{r=1}^{m} G_r^0 \right\rangle + \frac{L\eta^2}{2} \left\| \sum_{r=1}^{m} G_r^0 \right\|^2.$$
 (42)

C.1.2 SCHEDULED GROUP SEQUENCE OVER ONE REFRESH

Definition C.2 (Scheduled refresh). *Starting from the same point x, define*

$$x_0 := x, \qquad x_r := x_{r-1} - \eta G_r(x_{r-1}) \quad (r = 1, \dots, m), \qquad x^{\text{sch}} := x_m.$$
 (43)

Order and notation. The within refresh order (1, ..., m) may be fixed or randomly permuted each refresh. We write $H(\cdot)$ for the Hessian of F and take $\eta \in (0, 1/L]$.

Our goal is to compare upper bounds derived from L-smoothness for $F(x^{\text{sch}})$ and $F(x^{\text{agg}})$.

C.2 Telescoping bound for scheduled updates

Lemma 3 (Smoothness Expansion for Two Scheduled Groups). Let m=2 and $G_r^0:=G_r(x)$. For any $\eta\in(0,1/L]$,

$$F(x^{sch}) \leq F(x) - \eta \langle \nabla F(x), G_1^0 \rangle + \frac{L\eta^2}{2} \|G_1^0\|^2 - \eta \langle \nabla F(x), G_2(x_1) \rangle + \frac{L\eta^2}{2} \|G_2(x_1)\|^2 + \eta^2 \int_0^1 \langle H(x - t\eta G_1^0) G_1^0, G_2(x_1) \rangle dt.$$

$$(44)$$

Proof sketch. Apply the L-smoothness inequality at the first step to bound $F(x_1)$. For the second step, use L-smoothness at x_1 and expand

$$\nabla F(x_1) = \nabla F(x) - \int_0^1 H(x - t\eta G_1^0) \, \eta G_1^0 \, dt \tag{45}$$

by the fundamental theorem of calculus along the segment $x \to x_1$.

C.2.1 Start-of-refresh reduction under per-group Lipschitzness

We adopt the following assumption whenever we compare bounds solely in terms of start-of-refresh measurements. It will be used throughout Sections C.3–C.6

Assumption 6 (Per-group lipschitzness). Each group map $G_r(\cdot)$ is L_r -lipschitz:

$$||G_r(u) - G_r(v)|| \le L_r ||u - v||$$
 for all u, v . (46)

Under this assumption, for m=2 we have $G_2(x_1)=G_2^0+\delta_2$ with $\|\delta_2\|\leq L_2\eta\|G_1^0\|$, hence

$$||G_2(x_1)|| \le ||G_2^0|| + L_2 \eta ||G_1^0|| \tag{47}$$

For general m

$$||G_r(x_{r-1})|| \le ||G_r^0|| + L_r \eta \sum_{p \le r} ||G_p^0|| \qquad (r = 2, \dots, m)$$
 (48)

When these substitutions are made in scheduled bounds, the induced drift contributions are collected into a nonnegative penalty $R_m(x; \eta)$

C.3 UPPER BOUNDS FOR SCHEDULED AND AGGREGATED UPDATES (GENERAL m)

Applying L-smoothness m times yields the scheduled upper bound

$$UB_{sch}(x;\eta) := F(x) - \eta \sum_{r=1}^{m} \left\langle \nabla F(x), G_r(x_{r-1}) \right\rangle + \frac{L\eta^2}{2} \sum_{r=1}^{m} \|G_r(x_{r-1})\|^2 + \eta^2 \sum_{1 \le n \le q \le m} \int_0^1 \left\langle H(x - t\eta G_p(x_{p-1})) G_p(x_{p-1}), G_q(x_{q-1}) \right\rangle dt.$$

$$(49)$$

The aggregated upper bound is the one-shot bound from Equation 42, restated as

$$UB_{agg}(x;\eta) := F(x) - \eta \left\langle \nabla F(x), \sum_{r=1}^{m} G_r^0 \right\rangle + \frac{L\eta^2}{2} \left\| \sum_{r=1}^{m} G_r^0 \right\|^2$$
 (50)

The integrals in Equation 49 are over ordered pairs p < q along the specific sequence $x_0 \to x_1 \to \cdots \to x_m$; the bound therefore depends on the within-refresh order. Randomizing the order yields an expected version.

In Sections C.4–C.6 we express the scheduled bound in terms of $\{G_r^0\}$ under the per-group lipschitzness assumption. The associated drift terms are aggregated into $R_m(x;\eta)$.

C.4 SCHEDULED AND AGGREGATED GAP AT A COMMON LINEARIZATION

Define the shorthand

$$I_{pq}(x;\eta) := \int_0^1 \langle H(x - t\eta G_p^0) G_p^0, G_q^0 \rangle dt$$
 (51)

By expanding UB_{sch} around $\{G_r^0\}$ and collecting the lipschitz drift penalties into $R_m(x;\eta) \geq 0$, we obtain:

Theorem 5 (Upper-bound gap under per-group lipschitzness). Assuming per-group lipschitzness, for any partition $\{G_r\}$ and $\eta \in (0, 1/L]$,

$$UB_{sch}(x;\eta) - UB_{agg}(x;\eta) \leq \eta^2 \sum_{1 \leq p \leq q \leq m} \left(-L \langle G_p^0, G_q^0 \rangle + I_{pq}(x;\eta) \right) + R_m(x;\eta). \quad (52)$$

Using $||H(\cdot)||_{op} \leq L$ and Cauchy-Schwarz (Steele, 2004)

$$I_{pq}(x;\eta) \le L \|G_p^0\| \|G_q^0\|$$
 (53)

which gives the envelope

$$UB_{sch}(x;\eta) - UB_{agg}(x;\eta) \leq L\eta^{2} \sum_{p < q} (\|G_{p}^{0}\| \|G_{q}^{0}\| - \langle G_{p}^{0}, G_{q}^{0} \rangle) + R_{m}(x;\eta) \geq 0$$
 (54)

Interpretation This shows that without additional structure, the scheduled smoothness bound can be looser than the aggregated bound. The gap is governed by Hessian-weighted cross terms I_{pq}

Proposition 4 (Drift penalty bound under per-group lipschitzness). Assume each group map G_r is L_r -lipschitz. Then for $r \geq 2$,

$$||G_r(x_{r-1})|| \le ||G_r^0|| + L_r \eta \sum_{p < r} ||G_p^0|| := ||G_r^0|| + L_r \eta S_{r-1},$$
 (55)

and the scheduled start substitution error satisfies

$$R_{m}(x;\eta) \leq \eta^{2} \left(\sum_{p=1}^{m} \|G_{p}^{0}\| \right) \sum_{r=2}^{m} L_{r} S_{r-1} + \frac{L\eta^{2}}{2} \sum_{r=2}^{m} \left(2 \|G_{r}^{0}\| L_{r} \eta S_{r-1} + (L_{r} \eta S_{r-1})^{2} \right),$$

$$(56)$$

so $R_m(x;\eta) = O(\eta^2)$ with constants controlled by $\{L_r\}$ and $\{\|G_r^0\|\}$.

C.5 SUFFICIENT CONDITIONS FOR A TIGHTER SCHEDULED BOUND

The terms $I_{pq}(x;\eta)$ encode Hessian-weighted interactions between groups and determine when scheduling is advantageous at the bound level.

Assumption 7 (Hessian-weighted negative cross-terms). There exist nonnegative margins $\{\Gamma_{pq}\}_{p < q}$ such that

$$I_{pq}(x;\eta) = \int_{0}^{1} \langle H(x - t\eta G_{p}^{0}) G_{p}^{0}, G_{q}^{0} \rangle dt \leq -\Gamma_{pq} \|G_{p}^{0}\| \|G_{q}^{0}\| \text{ for all } p < q$$
 (57)

Theorem 6 (Strict upper-bound improvement under per-group lipschitzness and negative Hessian-weighted cross-terms). Assuming per-group lipschitzness and 57, for any $\eta \in (0, 1/L]$,

$$UB_{sch}(x;\eta) - UB_{agg}(x;\eta) \leq \eta^{2} \sum_{p < q} \left(-L \langle G_{p}^{0}, G_{q}^{0} \rangle - \Gamma_{pq} \|G_{p}^{0}\| \|G_{q}^{0}\| \right) + R_{m}(x;\eta)$$
 (58)

In particular, if

$$\sum_{p < q} \left(\Gamma_{pq} \| G_p^0 \| \| G_q^0 \| + L \langle G_p^0, G_q^0 \rangle \right) > \frac{R_m(x; \eta)}{\eta^2}$$
 (59)

then $UB_{sch}(x;\eta) < UB_{agg}(x;\eta)$

C.6 PL OR STRONG CONVEXITY: STANDARD RATE AND UPPER-BOUND GAINS FOR SCHEDULING

Assume F satisfies the Polyak–Łojasiewicz (PL) inequality with parameter $\mu > 0$:

$$\frac{1}{2}|\nabla F(x)|^2 \ge \mu(F(x) - F^*), \qquad \forall x \tag{60}$$

For any $\eta \in (0, 1/L]$, the single aggregated update satisfies the standard GD bound

$$F(x^{\text{agg}}) \le F(x) - \eta \left(1 - \frac{L\eta}{2}\right) |\nabla F(x)|^2 \le \left(1 - 2\mu\eta \left(1 - \frac{L\eta}{2}\right)\right) \left(F(x) - F^{\star}\right)$$
 (61)

Define the upper-bound gain (under per-group lipschitzness, so both bounds are expressed at start-of-refresh):

$$\Delta UB(x;\eta) := UB_{agg}(x;\eta) - UB_{sch}(x;\eta) \ge 0$$
(62)

whenever 59 holds. Since $F(x^{\mathrm{sch}}) \leq \mathrm{UB}_{\mathrm{sch}}(x;\eta)$ and $\mathrm{UB}_{\mathrm{agg}}(x;\eta)$ upper-bounds the one-shot decrease term in 61, we obtain the bound-level contraction

$$F(x^{\operatorname{sch}}) - F^{\star} \leq \left(1 - 2\mu\eta\left(1 - \frac{L\eta}{2}\right)\right) \left(F(x) - F^{\star}\right) - \Delta_{\operatorname{UB}}(x;\eta). \tag{63}$$

Consequently, under per-group lipschitzness and 59, the scheduled refresh satisfies the standard gradient-descent contraction and, in addition, achieves an extra nonnegative decrement $\Delta_{\mathrm{UB}}(x;\eta)$ in the upper bound.

C.7 WHY THE ASSUMPTIONS ARE MILD

The assumptions we use are mild. They are standard and naturally align with our training pipeline.

C.7.1 L-SMOOTHNESS

This is the same regularity used throughout the main paper and in our baselines. Each task loss we optimize is L_i -smooth, so the overall objective is L-smooth. We only use this to apply the standard smoothness (descent) inequality (Nesterov, 2004; Beck, 2017).

C.7.2 Per-group Lipschitzness of G_r

Each G_r is a fixed linear combination of the task gradients assigned to group r. If each task gradient is L_i -lipschitz, then G_r is lipschitz with constant $Lr \leq \sum i \in rL_i$. In other words, this property falls out of task-level smoothness. The same smoothness estimates we already use for step-size selection upper-bound the L_r .

C.7.3 NEGATIVE HESSIAN-WEIGHTED CROSS-TERMS

The condition we use asks that, over the short moves we actually take $(\eta \le 1/L)$, groups that are separated by the scheduler continue to exhibit negative interaction under the local Hessian (i.e., the Hessian-weighted cross-terms remain negative). This aligns with how the scheduler is built. It separates tasks that exhibit sustained negative interactions and it periodically refreshes assignments so the local geometry does not drift far. Thus the assumption matches the mechanism we deploy.

C.7.4 PL AND STRONG CONVEXITY

We invoke PL only to convert a per-refresh decrease into a standard contraction factor. We do not require global strong convexity. A local PL inequality around the iterates is enough, which is commonly observed after warm-up and annealing we already use (Karimi et al., 2016; Zhou et al., 2021; Liu, 2025).

C.8 CONCLUDING REMARKS

This appendix formalizes a bound-level comparison between scheduled and aggregated updates. Without additional structure the scheduled bound need not be tighter, but under per-group lipschitzness and negative Hessian-weighted cross-terms it becomes strictly tighter, and under PL the scheduled refresh inherits the standard GD contraction with an additional nonnegative decrement. In practice, these conditions arise naturally once the task-group assignments stabilize, so the scheduler will typically achieve tighter descent bounds without changing step sizes or gradient information.

D COMPUTATIONAL COMPLEXITY OF ONE REFRESH (AND AMORTIZED OVER TRAINING)

We analyze the computational and memory complexity of the proposed interference-aware scheduler per refresh and its amortized cost over training. The former accounts for the cost of a single refresh operation while the latter represents the average cost distributed across all training steps. We distinguish the work required by the underlying multi-task training objective (e.g., backpropagation

to obtain gradients) from the scheduler overhead (EMA maintenance, cosine computation, conflict graph construction, and color).

D.1 NOTATION

- $K \in \mathbb{N}$ number of tasks
- $d \in \mathbb{N}$ dimension of the gradient EMA vector per task
- $R \in \mathbb{N}$ refresh period (number of training steps between graph rebuilds)
- $\beta \in [0,1)$ exponential moving average (EMA) parameter
- $T \in \mathbb{N}$ total number of training steps
- ullet G > 0 time to compute one backward pass to obtain a task gradient at a refresh
- $\tau \in (0,1)$ conflict threshold; an undirected edge $\{i,j\}$ is present iff $\widehat{C}_{ij} < -\tau$
- $T_{\text{refresh}} > 0$ time cost of a single scheduler refresh
- $S_{\text{refresh}} > 0$ peak additional memory used during a refresh
- $N_{\text{refresh}} \in \mathbb{N}$ number of refreshes over T steps with period R (satisfies $N_{\text{refresh}} \in \{\lfloor T/R \rfloor, \lceil T/R \rceil\}$ and $N_{\text{refresh}} \leq T/R + 1$)

D.2 PER-REFRESH COMPLEXITY (TIME AND SPACE)

At a refresh, the scheduler performs a finite sequence of deterministic operations on the current collection of task-wise exponential moving averages (EMAs) of gradients. Let

$$M \in \mathbb{R}^{K \times d} \tag{64}$$

denote the matrix whose i-th row m_i^{\top} is the EMA for task i. A refresh first updates these rows through a scalar EMA rule

$$m_i \leftarrow \beta m_i + (1 - \beta)g_i \tag{65}$$

using the most recent probe (or reused) gradient g_i . It then constructs the cosine-similarity matrix

$$\widehat{C} = \widetilde{M}\widetilde{M}^{\top} \tag{66}$$

where \widetilde{M} is the row-normalized version of M. It thresholds \widehat{C} at $-\tau$ to obtain the conflict adjacency. Finally, it applies a graph-coloring routine to the resulting simple graph (Welsh & Powell, 1967).

EMA maintenance uses a constant number of vector operations per task: one multiply-add on each of the d coordinates of m_i . Aggregating over all K tasks gives a time proportional to Kd. The storage required to hold all EMAs is the $K \times d$ array M, so the working set devoted to EMAs is $\Theta(Kd)$ numbers.

The construction of \widehat{C} proceeds by normalizing each row of M and then multiplying \widetilde{M} by its transpose. Row normalization touches each entry exactly once and therefore costs $\Theta(Kd)$ time. The Gram product \widetilde{MM}^{\top} consists of K^2 dot products of length d, which is $\Theta(K^2d)$ time (Kågström et al., 1998). The cosine matrix itself occupies K^2 entries. If it is retained after thresholding, it uses $\Theta(K^2)$ space. If dropped right after graph construction, that $\Theta(K^2)$ storage is only temporary.

Thresholding linearly scans the off-diagonal of \widehat{C} , adding an undirected edge when $\widehat{C}_{ij} < -\tau$; this costs $\Theta(K^2)$ time. The result is either a dense $K \times K$ boolean array requiring $\Theta(K^2)$ space, or a sparse adjacency whose size depends on the number of conflicts (e.g., $\Theta(kK)$) when retaining the k most negative entries per row).

Putting these pieces together yields the following statement.

Proposition 5 (Per-refresh scheduler overhead). *Under the standard RAM model with dense matrix multiplication costed as* $\Theta(K^2d)$, *the time required by a single scheduler refresh is*

$$T_{refresh} = \Theta(Kd) + \Theta(K^2d) + \Theta(K^2) + O(K^2) = \Theta(K^2d),$$
 (67)

and the additional space required by the scheduler during the refresh is

 $S_{\textit{refresh}} = \Theta(Kd) + \Theta(K^2),$

where the $\Theta(K^2)$ term is transient if \widehat{C} is not retained after coloring.

Proof. The EMA update costs $\Theta(Kd)$ by a direct count of coordinate-wise multiply-adds. Row normalization also costs $\Theta(Kd)$. The Gram matrix requires K^2 inner products of length d, which is $\Theta(K^2d)$. This term dominates $\Theta(Kd)$. Thresholding scans $O(K^2)$ entries and is therefore $\Theta(K^2)$. The greedy coloring performs a sort of K keys and then assigns at most one color per edge incident on the current vertex, which is $O(K^2)$ in the worst case. This is dominated by $\Theta(K^2d)$ whenever $d \geq 1$. Summing these contributions and absorbing lower-order terms yields $T_{\text{refresh}} = \Theta(K^2d)$. The EMA matrix occupies $\Theta(Kd)$ memory, and storing \widehat{C} uses $\Theta(K^2)$. But if \widehat{C} is discarded immediately after thresholding, only $\Theta(Kd)$ remains.

D.3 AMORTIZED COST OVER TRAINING

Let $R \in \mathbb{N}$ denote the refresh period as the scheduler executes a refresh once every R training steps. Consider a training run of length T steps. The number of refreshes executed is $\lfloor T/R \rfloor$ or $\lceil T/R \rceil$ depending on whether one refresh occurs at step 0. In either case it is bounded by T/R+1. Multiplying the per-refresh time T_{refresh} by the number of refreshes and dividing by T shows that the amortized scheduler time per training step satisfies

$$\frac{1}{T} N_{\text{refresh}} T_{\text{refresh}} \le \frac{1}{T} \left(\frac{T}{R} + 1 \right) T_{\text{refresh}} = \frac{1}{R} T_{\text{refresh}} + \frac{1}{T} T_{\text{refresh}}$$
 (69)

Letting $T\to\infty$ (or simply taking T large compared to one refresh) eliminates the $T^{-1}T_{\text{refresh}}$ boundary term, yielding the asymptotic amortized bound

$$\frac{1}{R}T_{\text{refresh}} = \frac{1}{R}\Theta(K^2d) \tag{70}$$

(68)

If probe gradients are computed only at refreshes, their contribution $K\mathsf{G}$ per refresh adds $\frac{1}{R}\Theta(K\mathsf{G})$ to the amortized time per step. If, instead, the training loop already computes task-wise gradients each step and these are reused to update the EMAs, then the probe term is absent and the amortized scheduler overhead remains $\frac{1}{R}\Theta(K^2d)$.

The amortized space usage is simpler. The EMA matrix M must be retained throughout training and therefore contributes $\Theta(Kd)$ at all times. The cosine matrix \widehat{C} and the adjacency are constructed only during the refresh. They're released after coloring, so the $\Theta(K^2)$ space does not persist. Consequently, the persistent memory overhead attributable to the scheduler is $\Theta(Kd)$, while the peak overhead during a refresh is $\Theta(Kd) + \Theta(K^2)$.

D.4 CONDITIONS FOR NEGLIGIBLE OVERHEAD

Let the amortized per-step costs be

$$C_{\text{sched}} = \frac{a}{R} K^2 d$$
 and $C_{\text{probe}} = \frac{b}{R} K G,$ (71)

where a,b>0 are platform-dependent constants and G denotes the per-task backpropagation cost of the optional probe at a refresh. For fixed R,

$$\frac{C_{\text{sched}}}{C_{\text{probe}}} = \frac{a}{b} \frac{K^2 d}{K \,\mathsf{G}} = \frac{a}{b} \frac{K d}{\mathsf{G}}.\tag{72}$$

Hence C_{sched} is negligible relative to C_{probe} whenever

$$\frac{C_{\rm sched}}{C_{\rm probe}} \le \varepsilon \quad \text{for some } 0 < \varepsilon \ll 1 \quad \Longleftrightarrow \quad Kd \le \frac{b}{a} \varepsilon \, \mathsf{G}. \tag{73}$$

D.5 REDUCING TIME COMPLEXITY

In this section, we detail approaches that can be taken under certain circumstances to optimize time complexity.

D.5.1 RANDOM PROJECTIONS

We replace the EMA matrix $M \in \mathbb{R}^{K \times d}$ by a lower-dimensional sketch $\widetilde{M} = MR$ with $R \in \mathbb{R}^{d \times r}$ and $r \ll d$ (Dasgupta & Gupta, 2003). The sketching multiply costs O(Kdr) and the cosine Gram becomes $O(K^2r)$ instead of $\Theta(K^2d)$. Storage for the sketched EMAs is O(Kr). By the Johnson-Lindenstrauss (JL) random projection guarantee, if we map the K task-EMA vectors from \mathbb{R}^d to \mathbb{R}^r using a suitable random matrix with $r = \Theta(\epsilon^{-2} \log K)$, then after row normalization all pairwise inner products (hence cosines) are preserved within $\pm \epsilon$ with high probability. We assume a uniform row-norm floor $\min_i \|m_i\| \geq m_0 > 0$ (which can be enforced in practice by skipping tasks with $\|m_i\| < \nu \ll m_0$) so cosine errors remain controlled. Choosing $\epsilon < \gamma$, where γ is the cosine margin from the recovery analysis, ensures that every pair remains on the same side of the threshold $-\tau$. Therefore the set $\{(i,j): \widehat{C}_{ij} < -\tau\}$ and the resulting coloring are unchanged with high probability.

In short, dimensionality drops from d to r, the refresh cost drops from $\Theta(K^2d)$ to $O(Kdr + K^2r)$, and decisions are preserved as long as the chosen r makes the embedding error smaller than the margin.

D.5.2 DETERMINISTIC COVARIANCE SKETCHING VIA FREQUENT DIRECTIONS

We maintain a deterministic sketch $B \in \mathbb{R}^{\ell \times d}$ of the row space of M using Frequent Directions and either project rows onto $\operatorname{span}(B)$ or form an approximate Gram from the sketch (Liberty, 2013; Ghashami et al., 2016). Maintaining the sketch costs $O(Kd\ell)$, the cosine Gram in the sketch space costs $O(K^2\ell)$, and storage for the sketch is $O(\ell d)$. Frequent Directions gives a spectral-norm bound

$$|MM^{\top} - \widehat{MM^{\top}}|_2 \le \epsilon |M|_F^2 \tag{74}$$

when $\ell = \Theta(\epsilon^{-2})$, which yields a uniform bound on inner-product and squared-norm errors. Assuming a row-norm floor $\min_i \|m_i\| \ge m_0 > 0$ and applying a standard cosine perturbation bound after row normalization, one obtains

$$\left|\cos(m_i, m_j) - \widehat{\cos}(m_i, m_j)\right| \le \frac{2\epsilon \|M\|_F^2}{m_0^2} + O\left(\frac{\epsilon^2 \|M\|_F^4}{m_0^4}\right)$$
 (75)

Taking ϵ small enough so that the right-hand side is $<\gamma$ ensures that all threshold decisions and the resulting coloring are preserved deterministically. Thus the effective dimension drops from d to ℓ in the worst case, and the refresh cost becomes $O(Kd\ell + K^2\ell)$.

D.5.3 EDGE SAMPLING FOR CONFLICT GRAPHS WITH ADAPTIVE REFINEMENT

We reduce the number of cosine evaluations by computing \widehat{C}_{ij} for only $O(K \log K)$ randomly chosen task pairs to build a provisional conflict graph and then refining by evaluating additional pairs that are near the threshold or needed to certify connectivity and chromatic structure. We still compute all K row norms once in O(Kd) time for normalization, and the first pass costs $O(Kd \log K)$ for the sampled dot products. The total cost adds only the refinement work, which remains small when only few pairs are ambiguous. Under a planted separation model with margin γ and reasonably dense cross-group conflicts, one can show with high probability that the sampled graph already captures the correct inter-group connectivity, so the coloring or component structure is recovered after the first pass and only boundary pairs need refinement. This reduces the pairwise work from K^2 to near $K \log K$ while preserving the final decisions under stated assumptions (Erdős & Rényi, 1960).

D.5.4 INCREMENTAL GRAM UPDATES

We avoid rebuilding the full cosine matrix when only a small subset of tasks has meaningfully changed since the last refresh. If s rows of M cross a chosen change threshold, we first renormalize these rows and then recompute both the corresponding s rows and s columns of the Gram by taking

dot products against all K rows, which costs O(sKd), with an additional O(sd) to update norms, instead of $\Theta(K^2d)$, and we leave all unchanged entries as they are. This update is exact for the affected entries, so conflict edges and coloring decisions are preserved by construction, and the reduction is deterministic whenever $s \ll K$. To prevent slow drift in the unchanged entries, we can periodically force a full rebuild and reset the change counters.

E DESCENT PRESERVATION UNDER au-Compatibility

E.1 Proof of Proposition 6

Proposition 6. Let $S \subseteq \{1, ..., K\}$ be a τ -compatible task set. That is, every pair of gradients satisfies

$$\langle g_i, g_j \rangle \ge -\tau \|g_i\| \|g_j\|, \quad \forall i \ne j \in S, \quad 0 \le \tau < 1$$
 (76)

Then

$$\left\| \sum_{k \in S} g_k \right\|^2 \ge \left(1 - \tau(|S| - 1) \right) \sum_{k \in S} \|g_k\|^2. \tag{77}$$

Proof. We begin with the polarization identity for any finite set of vectors:

$$\left\| \sum_{k \in S} g_k \right\|^2 = \sum_{k \in S} \|g_k\|^2 + 2 \sum_{\substack{i,j \in S \\ i < j}} \langle g_i, g_j \rangle.$$
 (78)

E.1.1 LOWER-BOUNDING THE CROSS TERMS

Because S is τ -compatible, inequality (76) gives

$$\langle g_i, g_j \rangle \geq -\tau \|g_i\| \|g_j\|. \tag{79}$$

Insert this bound into (78) to obtain

$$\left\| \sum_{k} g_{k} \right\|^{2} \geq \sum_{k} \|g_{k}\|^{2} - 2\tau \sum_{i < j} \|g_{i}\| \|g_{j}\|.$$
 (80)

E.1.2 Symmetrizing the mixed sum

Observe that

$$\sum_{i < j} \|g_i\| \|g_j\| = \frac{1}{2} \sum_{\substack{i,j \\ i \neq j}} \|g_i\| \|g_j\|.$$
(81)

Substituting (81) into (80) yields

$$\left\| \sum_{k} g_{k} \right\|^{2} \geq \sum_{k} \|g_{k}\|^{2} - \tau \sum_{\substack{i,j\\i \neq j}} \|g_{i}\| \|g_{j}\|.$$
 (82)

E.1.3 BOUNDING THE MIXED SUM VIA CAUCHY-SCHWARZ

Apply the Cauchy-Schwarz inequality in $\mathbb{R}^{|S|}$ to the vectors $a=(\|g_1\|,\ldots,\|g_{|S|}\|)$ and $\mathbf{1}=(1,\ldots,1)$:

$$\sum_{k} ||g_{k}|| = \langle a, \mathbf{1} \rangle \le ||a|| \, ||\mathbf{1}|| = \left(\sum_{k} ||g_{k}||^{2}\right)^{1/2} \sqrt{|S|}.$$
 (83)

Using $(\sum_{k} a_{k})^{2} \leq |S| \sum_{k} a_{k}^{2}$ and (84),

$$\sum_{i \neq j} \|g_i\| \|g_j\| = \left(\sum_k \|g_k\|\right)^2 - \sum_k \|g_k\|^2, \tag{84}$$

we obtain the standard estimate

$$\sum_{i \neq j} \|g_i\| \|g_j\| \le (|S| - 1) \sum_k \|g_k\|^2. \tag{85}$$

Hence,

$$\tau \sum_{i \neq j} \|g_i\| \|g_j\| \le \tau \left(|S| - 1 \right) \sum_k \|g_k\|^2. \tag{86}$$

E.1.4 COMBINING BOUNDS

1514 Insert (86) into (82):

$$\left\| \sum_{k} g_{k} \right\|^{2} \geq \sum_{k} \|g_{k}\|^{2} - \tau (|S| - 1) \sum_{k} \|g_{k}\|^{2} = (1 - \tau (|S| - 1)) \sum_{k} \|g_{k}\|^{2}, \tag{87}$$

which is (77).

E.2 INTERPRETATION AND PRACTICAL IMPLICATIONS

Equation (77) guarantees that whenever we restrict an SGD step to a τ -compatible group (i.e., a set of tasks whose gradients are not too conflicting) the resulting joint update preserves at least a $\left(1-\tau(|S|-1)\right)$ fraction of the summed squared step lengths.

Below, we provide a strictly stronger version that is assumption free.

Proposition 7 (Data-Dependent Lower Bound via the Aggregate Conflict Ratio). *Define the aggregate conflict ratio*

$$\tau_{\text{eff}}(S) := \frac{\sum_{i \neq j} (-\langle g_i, g_j \rangle)_+}{\sum_{k} \|g_k\|^2}, \qquad (x)_+ := \max\{x, 0\}.$$
 (88)

Then, without additional assumptions,

$$\left\| \sum_{k \in S} g_k \right\|^2 \ge \left(1 - \tau_{\text{eff}}(S) \right) \sum_{k \in S} \|g_k\|^2,$$
 (89)

and under τ -compatibility we always have $\tau_{\text{eff}}(S) \leq \tau(|S|-1)$, so (89) is never weaker than (77).

Our takeaways from this are as follows:

- (i) Descent direction safety. The aggregated step is guaranteed to be a descent direction whenever $\tau_{\rm eff}(S) < 1$ (data-dependent) and, in particular, whenever $\tau(|S|-1) < 1$ (worst-case).
- (ii) Convergence-rate constant. In analyses for smooth SGD, one may replace $||g_t||^2$ by the right-hand side of either (89) (which is tighter) or (77) (worst-case), leading respectively to constants involving $\tau_{\text{eff}}(S_t)$ or $\tau(|S_t|-1)$.

F CONVERGENCE RATE WITH au -DEPENDENT CONSTANT

Theorem 7 (Baseline $O(1/\sqrt{T})$ convergence of the full gradient). Let $F(\theta) = \sum_{k=1}^K \mathcal{L}_k(\theta, \phi_k)$ be L-smooth in the shared parameters θ . Assume the stochastic gradient g_t obtained at step t satisfies $\mathbb{E}[g_t \mid \theta_t] = \nabla F(\theta_t)$ and $\mathbb{E}[\|g_t - \nabla F(\theta_t)\|^2 \mid \theta_t] \leq \sigma^2$. Let the step size be $\eta = \frac{c}{\sqrt{T}}$ with $0 < c \leq \frac{1}{L}$, and suppose the scheduler selects a τ -compatible task set S_t at each step (this will be used below for a refinement). Then

$$\min_{1 \le t \le T} \mathbb{E} \left[\|\nabla F(\theta_t)\|^2 \right] \le \frac{2(F_0 - F^*)}{c\sqrt{T}} + \frac{cL\sigma^2}{\sqrt{T}}.$$
 (90)

Proof. Because F is L-smooth, for any $\eta \leq \frac{1}{L}$ the standard non-convex SGD inequality (Ghadimi & Lan 2013, Lemma 3.2) gives

$$\mathbb{E}[F(\theta_{t+1})] \leq \mathbb{E}[F(\theta_t)] - \frac{\eta}{2} \mathbb{E}[\|\nabla F(\theta_t)\|^2] + \frac{\eta^2 L \sigma^2}{2}. \tag{91}$$

Summing equation 91 over t = 0, ..., T-1 and using $\mathbb{E}[F(\theta_T)] \geq F^*$ yields

$$\frac{\eta}{2} \sum_{t=0}^{T-1} \mathbb{E}[\|\nabla F(\theta_t)\|^2] \le F_0 - F^* + \frac{\eta^2 L \sigma^2 T}{2}. \tag{92}$$

Dividing by T, using $\min_t x_t \leq \frac{1}{T} \sum_t x_t$, and substituting $\eta = \frac{c}{\sqrt{T}}$ gives equation 90.

Data-dependent τ -refinement for the scheduled gradient energy. For a finite set S, define the aggregate conflict ratio

$$\tau_{\text{eff}}(S) := \frac{\sum_{i \neq j \in S} \left(-\langle g_i, g_j \rangle \right)_+}{\sum_{k \in S} \|g_k\|^2} \in [0, \infty), \qquad (x)_+ = \max\{x, 0\}.$$
 (93)

Then for every step t,

$$\left\| \sum_{k \in S_t} g_{k,t} \right\|^2 \ge \left(1 - \tau_{\text{eff}}(S_t) \right) \sum_{k \in S_t} \|g_{k,t}\|^2.$$
 (94)

Consequently,

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \Big[\sum_{k \in S_t} \|g_{k,t}\|^2 \Big] \le \underbrace{\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \Big[\frac{1}{1 - \tau_{\text{eff}}(S_t)} \Big]}_{=: \Gamma_T} \cdot \underbrace{\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \big[\|g_t\|^2 \big]}_{}. \tag{95}$$

Using $\mathbb{E}\|g_t\|^2 = \mathbb{E}\|\nabla F(\theta_t)\|^2 + \mathbb{E}\|g_t - \nabla F(\theta_t)\|^2 \le \mathbb{E}\|\nabla F(\theta_t)\|^2 + \sigma^2$ and the average version of equation 91,

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|\nabla F(\theta_t)\|^2] \le \frac{2(F_0 - F^*)}{\eta T} + L\eta \sigma^2, \tag{96}$$

we obtain the τ -dependent, data-driven control

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \Big[\sum_{k \in S_t} \|g_{k,t}\|^2 \Big] \le \Gamma_T \left(\frac{2(F_0 - F^*)}{\eta T} + L \eta \sigma^2 + \sigma^2 \right). \tag{97}$$

If, in addition, each S_t is pairwise τ -compatible with $|S_t| = s_t$ and $\tau\left(s_t - 1\right) \leq \rho < 1$ uniformly in t, then $\tau_{\text{eff}}(S_t) \leq \tau(s_t - 1) \leq \rho$ and hence $\Gamma_T \leq \frac{1}{1-\rho}$. With $\eta = \frac{c}{\sqrt{T}}$, equation 97 becomes

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \Big[\sum_{k \in S_t} \|g_{k,t}\|^2 \Big] \le \frac{1}{1-\rho} \left(\frac{2(F_0 - F^*)}{c\sqrt{T}} + \frac{cL\sigma^2}{\sqrt{T}} + \sigma^2 \right). \tag{98}$$

F.1 DISCUSSION AND INTUITION

Equation equation 90 is the classical $O(1/\sqrt{T})$ rate for non-convex SGD with unbiased and bounded variance gradients and constant-over-time step size $\eta = c/\sqrt{T}$. Under these conditions, the convergence rate in terms of the full gradient norm $\|\nabla F(\theta_t)\|^2$ does not depend on τ . However, the scheduler's τ structure does control the per step energy of the scheduled gradient through equation 95–equation 98. Less cross-task conflict (smaller Γ_T) results in a tighter bound on $\frac{1}{T}\sum_{k\in S_t}\|g_{k,t}\|^2$, which is the quantity governed by the descent preservation inequalities used throughout the analysis.

G BOUNDED STALENESS VIA GREEDY GRAPH COLORING

Proposition 8 (Staleness Bound). Let $G = (\mathcal{T}, E)$ be the task-conflict graph whose vertices are tasks and whose edges connect pairs with interference coefficient exceeding the threshold τ . Denote by Δ its maximum degree. Greedy graph coloring produces a proper coloring C_1, \ldots, C_m with

$$m \le \Delta + 1. \tag{99}$$

If the scheduler activates the color classes in the cyclic order $C_1 \to C_2 \to \ldots \to C_m \to C_1 \to \ldots$, then every task is updated at least once every

$$s_{\text{max}} = m - 1 \le \Delta \tag{100}$$

iterations. In particular, the schedule enforces a bounded inter-update delay of at most Δ iterations per task, consistent with the bounded-delay assumption of Recht et al. (Niu et al., 2011).

Proof. We proceed in two parts.

Part A: Color count bound. A greedy algorithm scans vertices in some order and assigns to each vertex the smallest available color not used by its already colored neighbors. When the *i*-th vertex v is reached, at most $\deg(v) \leq \Delta$ of its neighbors are already colored, so at most Δ colors are unavailable. Therefore one of the first $\Delta + 1$ colors is always free, implying $m \leq \Delta + 1$ (Lovász, 2006).

Part B: Staleness of cyclic execution. Fix any task $T \in \mathcal{T}$ and let it belong to color C_j for some $1 \leq j \leq m$. Under cyclic scheduling, C_j is executed at steps $t = j, \ j + m, \ j + 2m, \ldots$ The number of intervening steps between two consecutive executions of C_j is exactly m-1. Hence task T never waits more than $s_{\max} = m-1$ iterations for an update. Combining with Equation 8 yields $s_{\max} \leq \Delta$.

G.1 INTERPRETATION

The bound (Equation 100) guarantees that the shared parameters used by any task are refreshed at least once every Δ iterations in the worst case (e.g., when the conflict graph is a clique of size $\Delta+1$). This aligns with the bounded-delay assumption common in analyses of asynchronous SGD and lock-free training, so convergence proofs built under that assumption apply to our cyclic schedule with delay parameter at most Δ when iterations are used as the unit of delay (Niu et al., 2011; Lian et al., 2015). In practice Δ is often much smaller than the total number of tasks, so the scheduler achieves low interference and low parameter staleness simultaneously.

H Greedy Graph-Coloring Uses at Most $\Delta+1$ Colors

H.1 Proof of Proposition 9

Proposition 9 (Coloring Period Bound). Let G = (V, E) be a finite, simple, undirected graph with maximum degree $\Delta := \max_{v \in V} \deg(v)$. The greedy (first-fit) coloring algorithm (e.g., Welsh-Powell order)¹ produces a proper vertex coloring with no more than

$$\chi_{greedy}(G) \le \Delta + 1 \tag{101}$$

distinct colors. Consequently, when the scheduler activates the color classes in a cyclic order, the cycle length is bounded by $\Delta+1$. This is a quantity depending only on the structure of the conflict graph.

Proof. Let the vertices be processed in the chosen order $v_1, v_2, \ldots, v_{|V|}$ (e.g., Welsh-Powell). Assume inductively that after coloring the first k-1 vertices the algorithm has used at most $\Delta+1$ colors. Consider vertex v_k . Since $\deg(v_k) \leq \Delta$, at most Δ neighbors of v_k can appear before v_k in the ordering. Hence, at the moment of coloring v_k , at most Δ colors are forbidden (one for each previously colored neighbor). Among the palette $\{1,2,\ldots,\Delta+1\}$ there is therefore at least one color still available. Assigning the smallest such color to v_k maintains a proper coloring and never introduces a new color beyond $\Delta+1$.

Proceeding vertex-by-vertex, no step ever requires more than $\Delta+1$ colors, establishing equation 101.

H.2 IMPLICATIONS FOR THE SCHEDULER

A coloring with at most $\Delta+1$ classes means the scheduler's cycle period (the number of batches needed before every task reappears) is bounded by a graph invariant independent of the number of tasks. Even if thousands of tasks exist, as long as each one conflicts with at most Δ others, the memory footprint (one shared backbone plus $\Delta+1$ sets of head activations) and the maximum waiting time between successive updates for any task (bounded by Δ , see Proposition 8) remain predictable and small. This guarantee is essential for scaling the scheduler to large, heterogeneous tasks.

¹Order the vertices in non-increasing degree and assign to each the smallest positive integer (color) not used by its previously colored neighbors.

I BASELINE NON-CONVEX SGD CONVERGENCE RATE

I.1 PROOF OF THEOREM 8

Theorem 8 (Classical $O(1/\sqrt{T})$ bound). Let $F : \mathbb{R}^d \to \mathbb{R}$ be an L-smooth, possibly non-convex objective and suppose the stochastic gradient g_t computed at iteration t satisfies

$$\mathbb{E}[g_t \mid \theta_t] = \nabla F(\theta_t), \qquad \mathbb{E}[\|g_t - \nabla F(\theta_t)\|^2 \mid \theta_t] \le \sigma^2. \tag{102}$$

Run SGD with the constant step size $\eta = \frac{c}{\sqrt{T}}$, $0 < c \le \frac{1}{L}$, for T iterations starting from θ_0 . Then

$$\min_{0 \le t < T} \mathbb{E}[\|\nabla F(\theta_t)\|^2] \le \frac{2(F_0 - F^*)}{c\sqrt{T}} + \frac{cL\sigma^2}{\sqrt{T}},\tag{103}$$

where $F^* = \inf_{\theta} F(\theta)$.

Proof. The proof is a streamlined restatement of ((Ghadimi & Lan, 2013; Nemirovski et al., 2009)). By *L*-smoothness,

$$F(\theta_{t+1}) \leq F(\theta_t) + \langle \nabla F(\theta_t), \theta_{t+1} - \theta_t \rangle + \frac{L}{2} \|\theta_{t+1} - \theta_t\|^2.$$
 (104)

With $\theta_{t+1} = \theta_t - \eta g_t$ and taking conditional expectation,

$$\mathbb{E}[F(\theta_{t+1})] \le \mathbb{E}[F(\theta_t)] - \eta \,\mathbb{E}[\|\nabla F(\theta_t)\|^2] + \frac{\eta^2 L}{2} \,\mathbb{E}[\|g_t\|^2]. \tag{105}$$

Decompose the squared stochastic gradient:

$$\mathbb{E}[\|g_t\|^2] = \mathbb{E}[\|\nabla F(\theta_t)\|^2] + \mathbb{E}[\|g_t - \nabla F(\theta_t)\|^2] \le \mathbb{E}[\|\nabla F(\theta_t)\|^2] + \sigma^2$$
(106)

Thus, and using $\eta \leq 1/L$ so that $\eta - \frac{L\eta^2}{2} \geq \frac{\eta}{2}$,

$$\mathbb{E}[F(\theta_{t+1})] \le \mathbb{E}[F(\theta_t)] - \frac{\eta}{2} \mathbb{E}[\|\nabla F(\theta_t)\|^2] + \frac{\eta^2 L \sigma^2}{2}.$$
 (107)

Summing from t = 0 to T - 1 and telescoping gives

$$\frac{\eta}{2} \sum_{t=0}^{T-1} \mathbb{E}[\|\nabla F(\theta_t)\|^2] \le F_0 - F^* + \frac{\eta^2 L \sigma^2 T}{2}.$$
 (108)

Dividing by ηT and inserting $\eta = c/\sqrt{T}$ yields equation 103.

I.2 CONNECTION TO THE SCHEDULER

At $\tau=0$, pairs with negative inner product are incompatible, so the conflict graph on tasks can be colored into m classes $\{C_1,\ldots,C_m\}$, and a simple policy activates one color class per step. Under a deterministic (cyclic) activation order, the update $g_t=\sum_{k\in S_t}g_{k,t}$ generally satisfies

$$\mathbb{E}[g_t \mid \theta_t] = \sum_{k \in S_t} \nabla \mathcal{L}_k(\theta_t, \phi_{k,t}) \neq \sum_{k=1}^K \nabla \mathcal{L}_k(\theta_t, \phi_{k,t}), \tag{109}$$

so it is biased for the full gradient.

I.2.1 CONSISTENCY WITH THE UNBIASED SGD ASSUMPTION

The analysis in Theorem 8 assumes an unbiased stochastic gradient, $\mathbb{E}[g_t \mid \theta_t] = \nabla F(\theta_t)$. This assumption is met under either of the following implementations.

(i) Randomized class sampling with scaling. Draw $J_t \sim \text{Unif}\{1,\ldots,m\}$ independently each step and set

$$\tilde{g}_t = m \sum_{k \in C_{J_t}} g_{k,t}. \tag{110}$$

Table 3: Information on the datasets utilized in experimentation. (*Some samples were removed during preprocessing)

Dataset	Main Tasks	(+) Aux. Tasks	(-) Aux Tasks	Modalities	Samples
	Semantic Segmentation				
	Depth Estimation				
NYUv2	Surface Normal Prediction	_	Color Temp. Estimation	Image	250*
		Quadrant Localization	Corruption-Type Prediction		
CIFAR-10	Image Classification	Texture Classification	Rotation Angle Prediction	Image	2,500*
AV-MNIST	Digit Classification	Digit Parity		Audio, Image	56.0k
MM-IMDb	Genre Classification	Release Decade	Title-Iniial Classification	Image, Text	25.9k
		Five-Day Rolling Volatility	Day of the Week Prediction		
STOCKS-F&B	$4\times$ Stock Return Prediction	Sector-Average Next-Day Return	Lag-0 Reconstruction of Today's Open-Price	Timeseries $\times 18$	75.5k
		Five-Day Rolling Volatility	Day of the Week Prediction		
STOCKS-HEALTH	7× Stock Return Prediction	Sector-Average Next-Day Return	Lag-0 Reconstruction of Today's Open-Price	Timeseries ×63	75.5k

Then $\mathbb{E}[\tilde{g}_t \mid \theta_t] = \sum_{k=1}^K \nabla \mathcal{L}_k(\theta_t, \phi_{k,t}) = \nabla F(\theta_t)$, so Theorem 8 applies (with the variance bound adjusted for the scaled estimator). Equivalently, one may keep $g_t = \sum_{k \in C_{J_t}} g_{k,t}$ and use an effective step size $m\eta$.

(ii) Deterministic cyclic schedule. If the classes are visited in a fixed periodic order, then generally $\mathbb{E}[g_t \mid \theta_t] \neq \nabla F(\theta_t)$ at the per-step level. Nonetheless, standard analyses of nonconvex smooth cyclic block updates yield an $O(1/\sqrt{T})$ decay of the average gradient norm under usual step-size conditions, with constants depending on the number of blocks.

Either implementation delivers an $O(1/\sqrt{T})$ convergence guarantee.

J EXPERIMENTAL SETUP FOR DATASETS

We evaluate the proposed scheduler alongside numerous baselines and state-of-the-art models across multiple datasets to reliably assess its general performance relative to other approaches. In total, it is evaluated across 6 datasets.

Across all datasets, we incorporate positive and/or negative auxiliary tasks into training. Positive auxiliary tasks share structure or predictive signals with the main tasks (e.g., common features or correlated outputs) and so can improve the learned representations by providing relevant supervision. In contrast, negative auxiliary tasks are uncorrelated or directly conflicting with the main objectives, inducing gradient interference that can slow or degrade primary performance. Including both creates controlled variation in task alignment, letting us test whether SON-GOKU (1) groups compatible tasks, (2) separates conflicting tasks, and (3) maintains main-task performance under interference created by auxiliary tasks.

J.1 NYUv2

The NYU Depth Dataset v2 (NYUv2) (Silberman et al., 2012) consists of RGB-D indoor scenes with 1,449 densely labeled pairs of RGB and depth images. To demonstrate auxiliary task value in data-scarce conditions, we employ a subset of 250 training samples randomly selected from the original training set.

We formulate a multi-main-task setup with three primary objectives: (1) semantic segmentation (14 classes), (2) depth estimation where the model predicts per-pixel depth values from RGB images, and (3) surface normal prediction where 3-channel surface normals are estimated from RGB input. The negative auxiliary task is color temperature estimation, a synthetically generated task that predicts global color temperature properties designed to interfere with the main tasks by emphasizing global color distribution rather than local semantic and geometric features.

All tasks utilize RGB images as the sole input modality, with depth maps and surface normals serving as prediction targets rather than input features. A ResNet-18 (He et al., 2015) backbone trained from scratch processes the RGB input, with task-specific decoder heads for segmentation (with 32×10^{-5} upsampling), depth regression, surface normal regression, and color temperature estimation.

J.2 CIFAR-10

1782

1783 1784

1785

1786

1787

1788

1789

1790

1791

1792

1793

1794

1795 1796

1797

1798 1799

1801

1802

1803

1805

1807

1808

1809

1810

1811

1812

1813 1814

1815 1816

1817

1818

1819

1820

1821

1822

1824

1825 1826

1827

1828

1829

1830 1831

1832

1833

1834

The CIFAR-10 (Krizhevsky et al., 2009) dataset contains $60,000 32 \times 32$ color images across 10 generic classes. To evaluate our interference-aware scheduler in a data-scarce environment where auxiliary tasks provide maximum benefit, we employ a subset of 2,500 training samples (250 per class) from the original 50,000 training images.

For the multi-task learning setup, we set image classification as the main task and construct three auxiliary tasks synthetically from the RGB images. The positive auxiliary tasks include: (1) quadrant localization, where the model predicts which quadrant contains the primary object, and (2) texture classification using Gabor filter responses clustered into 8 texture categories via k-means clustering. The negative auxiliary tasks consist of: (3) corruption-type prediction, where images are artificially corrupted using 15 different corruption types from the ImageNet-C corruption suite (Hendrycks & Dietterich, 2019), and (4) rotation angle prediction, where images are rotated by 0° , 90° , 180° , or 270° and the model predicts the rotation angle.

All tasks share a ResNet-18 (He et al., 2015) backbone trained from scratch without pretraining, with task-specific heads for each auxiliary task.

J.3 AV-MNIST

The AV-MNIST benchmark (Vielzeuf et al., 2018) pairs MNIST images (Lecun et al., 1998) with a log-mel spectrogram of the corresponding spoken digit from TIDIGITS (Leonard & Doddington, 1993). It is a synthetic benchmark that has significant noise applied to audio and feature reduction applied to images, making it far more difficult than the original MNIST.

We use all paired samples in our experiments. Our primary task is 10-way digit classification. Following (Vielzeuf et al., 2018), we encode images with a small 4-layer convolutional network and spectrograms with a 2-layer CNN, both built and trained from scratch. These embeddings are projected and fused for processing by a simple MLP in intermediate fusion (Boulahia et al., 2021; Guarrasi et al., 2025), as are the models trained on MM-IMDb and STOCKS. We include only one positive auxiliary class, Digital Parity. This task aims to identify the digits as either even or odd, which has been shown to be a positive auxiliary task for improving representations on MNIST-like datasets (Tacchetti et al., 2018; Mohammadi et al., 2020).

J.4 MM-IMDB

The MM-IMDb dataset (Arevalo et al., 2017) contains 25,959 movies with genre annotations over 23 categories. We extract poster images and plot summaries for every movie in the dataset.

The images and summaries are encoded by a frozen VGG16 (Simonyan & Zisserman, 2014) and Google word2vec (Mikolov et al., 2013) model, respectively. Our main task is movie genre prediction. We add one positive auxiliary task, Release Decade, and one negative auxiliary task, the classification of the title's first word as either a vowel or consonant.

J.5 STOCKS

The STOCKS datasets we use, introduced in (Liang et al., 2021), contain stock market timeseries data across two categories. Specifically: (1) STOCKS-F&B, which has 14 input and 4 output stocks in the GICS Restaurants or Packaged Food & Meats category (MSCI Inc. & S&P Dow Jones Indices, 2024), and (2) STOCKS-HEALTH, which contains 56 input and 7 output stocks in the Health Care category.

Every input stock consists of 500 trading days, with the goal of predicting returns over the next day. We discretize the continuous return variable R into three non-overlapping categories: (1) Low, where $0 \le R < 0.1$, (2) Medium, where $0.1 \le R < 0.5$, and (3) High, where $R \ge 0.5$. Mean Absolute Error (MAE) is calculated by mapping the three classes to numbers ($Low \rightarrow 0$, $Medium \rightarrow 1$, $High \rightarrow 2$) and then deriving MAE as usual. Each input series is encoded by the 1835 same CNN-BiLSTM network. This consists of 3 CNNs and 1 BiLSTM (Cui et al., 2018).

We augment the main prediction task with two positive auxiliaries and two negative auxiliaries. The first positive task, Five-Day Rolling Volatility, is calculated as the standard deviation of daily logarithmic returns over a sliding five-trading-day window. This feature captures short-term fluctuations in a stock's price. In Sector-Average Next-Day Return, for each date we compute the mean of the actual next-day returns of all stocks within the same GICS sector, providing a simple measure of sector-level momentum and drift

The negative tasks focus on useless information that is meant to distract the model. Namely, day of the week prediction (in the range of Monday to Friday) and Lag-0 Open-Price Reconstruction, which requires the model to reproduce the same day's opening price verbatim. The first is information that contains little to no signals that would contribute to overall performance, and the second is a trivial identity mapping that contributes no real predictive challenge.

K Models Used for Comparison

K.1 BASELINE MODELS

- 1. *Uniform*. This baseline assigns equal weights to all tasks throughout training, representing the simplest approach where all task losses are weighted equally.
- Gradnorm (Chen et al., 2018). Balances task learning rates by normalizing gradient magnitudes relative to target loss ratios. This maintains consistent training dynamics across tasks.
- 3. MGDA (Sener & Koltun, 2018). Formulates multi-task learning as a multi-objective optimization problem, finding Pareto-optimal solutions (Lockwood, 2008; Pareto, 2014) through gradient descent in the convex hull of gradients (Fliege & Svaiter, 2000; Miettinen, 1999).

K.2 State-of-the-art models

- 1. *PCGrad (Yu et al., 2020)*. Projects conflicting gradients onto orthogonal subspaces when negative cosine similarity is detected, eliminating destructive interference between task gradients.
- 2. CAGrad (Liu et al., 2021). Extends PCGrad by adaptively adjusting gradient magnitudes based on conflict severity. This proves more nuanced modifications to gradients than binary projection.
- 3. Adatask (Yang et al., 2023). Dynamically reweighs task losses using relative loss changes, adapting to varying task learning rates during training.
- 4. FAMO (Liu et al., 2023). Fast Adaptive Multitask Optimization dynamically adjusts task weights to equalize each task's rate of loss improvement. It uses an online, per-step rule (no pairwise gradient ops), adding negligible overhead while remaining robust to loss-scale differences.
- 5. Fair Resource Allocation in MTL (FairGrad) (Ban & Ji, 2024). Views the shared update as a limited resource and chooses it to maximize an α -fair utility of per-task improvements. The parameter α controls the trade-off between average performance and fairness.
- 6. Nash-MTL (Navon et al., 2022). Frames multitask training as a bargaining game and computes a scale-invariant weighted combination of task gradients given by the Nash bargaining solution. Weights are obtained by solving a small inner problem (e.g., via CCP) using the gradient Gram matrix. Updates are balanced across tasks.

L EXPANDED WALL-CLOCK TIME STUDY

We provide more results from our wall-clock time study. The expanded table includes results from testing refresh rates $R \in \{4, 32, 256\}$ for scheduler-based methods.

L.1 EXPERIMENTAL SETUP FOR WALL-CLOCK TIME STUDY

We benchmark wall-clock time with a controlled synthetic workload to remove the effects of data loading and I/O. For each configuration (number of tasks K and scheduler refresh rates R), we

Table 4: We present wall-clock time (seconds \pm standard deviation) across all K and scheduler refresh rates $R \in \{4, 32, 256\}$. We split results into sub-tables by R for readability. Non-scheduler methods do not depend on R, so they are shown in the R=4 sub-table and omitted in the R=32, 256 subtables to avoid redundancy.

(a) R=4 (all methods)

Method	K=3	K=6	K=16	K=40
Uniform	0.2656 ± 0.1201	0.3240 ± 0.0629	0.3798 ± 0.1050	0.4054 ± 0.1190
GradNorm	5.4714 ± 0.7137	5.1201 ± 0.6112	4.9042 ± 0.5869	4.7372 ± 0.9286
MGDA	97.1081 ± 5.4645	121.4371 ± 9.0923	132.4913 ± 3.1752	134.0878 ± 2.2621
PCGrad	3.6212 ± 0.3517	23.1266 ± 0.8773	176.7566 ± 2.8171	1127.1337 ± 34.2603
CAGrad	102.8651 ± 18.3422	136.1034 ± 2.4218	134.3585 ± 4.0791	132.7034 ± 1.2412
AdaTask	2.1816 ± 0.0934	2.1032 ± 0.1012	2.2853 ± 0.0718	2.2278 ± 0.1370
FAMO	2.0725 ± 0.2073	1.9980 ± 0.1998	2.1710 ± 0.2171	2.1164 ± 0.2116
FairGrad	3.8020 ± 0.5703	15.2079 ± 2.2812	108.1450 ± 16.2218	675.9065 ± 101.3860
Nash-MTL	5.7030 ± 1.1406	22.8118 ± 4.5624	162.2176 ± 32.4435	1013.8598 ± 202.7720
SON-GOKU	2.0904 ± 0.3506	3.6770 ± 0.4974	6.3225 ± 0.7895	14.3280 ± 1.4073
SON- $GOKU + AdaTask$	4.1011 ± 0.4174	5.2126 ± 0.6066	7.6798 ± 0.7107	14.7528 ± 1.8671
SON- $GOKU + GradNorm$	7.3223 ± 0.4994	8.5898 ± 0.8203	12.1065 ± 2.5850	16.8329 ± 1.9803
SON- $GOKU + PCGrad$	2.3489 ± 0.3258	3.5925 ± 0.4100	6.1549 ± 0.8461	12.5729 ± 1.2657

(b) R=32 (scheduler-based approaches)

Method	K=3	K=6	K=16	K=40
SON-GOKU	1.9896 ± 0.3651	3.3202 ± 0.5745	6.0897 ± 0.9425	12.1432 ± 1.2044
SON- $GOKU + AdaTask$	3.7718 ± 0.9654	5.0511 ± 0.6531	7.5903 ± 1.1920	14.5182 ± 2.0660
SON- $GOKU + GradNorm$	7.0202 ± 1.0711	8.1661 ± 0.9355	10.7227 ± 2.2088	16.5760 ± 1.8418
SON- $GOKU + PCGrad$	1.9834 ± 0.3586	3.4971 ± 0.3840	6.1395 ± 0.9425	10.9097 ± 1.5263

(c) R=256 (scheduler-based approaches)

Method	K=3	K=6	K=16	K=40
SON-GOKU	1.7593 ± 0.2280	3.0024 ± 0.3942	4.8411 ± 0.7302	11.4162 ± 1.6076
SON- $GOKU + AdaTask$	3.7224 ± 0.2696	4.4548 ± 0.5837	7.5276 ± 0.6230	13.0608 ± 3.2925
SON- $GOKU + GradNorm$	6.0221 ± 1.0418	7.8659 ± 0.7917	9.5029 ± 1.2168	15.6860 ± 2.3680
SON- $GOKU + PCGrad$	1.6776 ± 0.4104	3.0189 ± 0.7854	5.9893 ± 1.3797	7.1915 ± 0.2021

pre-generate a fixed sequence of per-task gradient vectors and loss values directly on the target device, and then feed the same exact tensors, in the same exact order, to every method. We set the gradient dimensionality to 1024. Timing uses a high-resolution clock with a device synchronize before starting and after finishing to capture only on-device compute. We also accumulate the norm of the combined gradient into a scalar accumulator (also known as a scalar sink) so the backend must realize the computation, avoiding lazy evaluation. Each MTL approach is run for 900 steps and repeated 10 times.

M EXTENDED RELATED WORK

Multi-task learning (MTL) methods have evolved from simple loss-weighting approaches to larger and more sophisticated optimization techniques that manage task conflict and cooperation (Yang et al., 2023). Early adaptive-weighting approaches sought to balance losses automatically (Vandenhende et al., 2022; Fan et al., 2023), while more recent work modifies gradients directly (Yu et al., 2020). Task scheduling and grouping methods, though far less popular than adaptive weighting techniques (Torbarina et al., 2023), have contributed to the field by controlling the timing of updates.

M.1 TUNED LOSS WEIGHTING

From early MTL work it became clear that simply summing task losses often favors one objective at the expense of others (Kurin et al., 2022; Zhao et al., 2024; Mueller et al., 2022), especially when losses have different scales or noise levels. To address this, practitioners manually tuned per-task weight coefficients (λ -values) to rebalance learning (Argyriou et al., 2007; Ando & Zhang, 2005;

Evgeniou et al., 2005; Kang et al., 2011; Liang & Zhang, 2020; Lin et al., 2022; Yu et al., 2021), but this process was laborious and dataset-specific. Thus, researchers began to develop automated methods.

M.2 Adaptive Loss Weighting

(Kendall et al., 2018) introduced uncertainty weighting, learning each task's homoscedastic (constant-variance) (Bishop, 2006) noise to scale losses automatically and improve depth and semantics on NYUv2 (Silberman et al., 2012).

GradNorm automatically balances multiple loss functions by tuning each task's gradient magnitude so that all tasks train at comparable speeds (Chen et al., 2018). It does this by introducing a single asymmetry hyperparameter α that governs how much each task's loss is scaled. This eliminates the need for expensive grid searches over manual weights. GradNorm was also a major leap empirically as it surpassed exhaustive search baselines on both regression and classification tasks. Dynamic Weight Averaging (DWA) extended this idea by adjusting weights based on loss rate of change, reducing oscillations between tasks (Liu et al., 2019).

More recently AdaTask applies task-specific learning rates that adapt to each head's gradient norm, yielding significant gains on multi-label classification benchmarks (Yang et al., 2023).

M.3 GRADIENT-LEVEL CONFLICT MITIGATION

Rather than rescaling losses, gradient surgery methods alter update directions. PCGrad projects gradients that conflict (negative cosine) onto each other's normal plane, significantly boosting efficiency on supervised vision and RL problems (Yu et al., 2020). CAGrad frames task balance as a min-max optimization, finding updates that maximize the worst-case task improvement (Liu et al., 2021). The Multiple Gradient Descent Algorithm (MGDA) computes a Pareto-optimal convex combination of task gradients, ensuring no task is harmed (Sener & Koltun, 2018). More recent variants such as SAM-GS incorporate momentum into conflict detection, smoothing gradient estimates while preserving the benefits of surgery (Borsani et al., 2025).

M.4 EMPIRICAL TASK GROUPING

Task grouping aims to decide which tasks should train together so that helpful transfer is amplified and harmful interference is limited. It typically groups tasks into subsets that update jointly, rather than updating all tasks at once. This is different from approaches that keep all tasks active or reweight the joint gradient (adaptive loss weighting, gradient surgery).

Early approaches under this category used round-robin and random sampling-based approaches that ignored any task relationships (McCann et al., 2018; Zamir et al., 2020). Standley et al. (2020) exhaustively searches over small subsets to identify beneficial groupings, demonstrating the potential of selective updates but failing to scale beyond eight tasks due to computational complexity.

Task Affinity Groupings (TAG) (Fifty et al., 2021) performs one joint training run to measure intertask 'affinity'. It quantifies how an update for task i (its gradient) would change task j's loss, and it uses these cross-effects to select partitions of tasks that should share updates. The key idea is to treat grouping as an outcome of measured gradient interactions.

Ayman et al. (Ayman et al., 2023) train a predictor that maps single-task statistics and dataset features to an estimate of whether two or more tasks should be grouped. They then use that predictor to guide a randomized search over groups, which dramatically reduces the number of multi-task trainings (or 'MTL trials') needed to find a good partition.

Using a completely different approach, Towards Principled Task Grouping (PTG) (Wang et al., 2024) formulates grouping as a mathematical program with a theoretically motivated objective capturing beneficial transfer while respecting resource constraints (e.g., compute budgets). It builds a principled optimization over candidate groups that is meant to generalize across application domains.

Scalable Task Grouping via Training Dynamics (STG-MLT) (Sherif et al., 2024) avoids expensive affinity estimation by extracting Data Maps (Swayamdipta et al., 2020) (simple summaries of training dynamics per task) and then clustering tasks using those features. The clusters are intended to push