

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 SENTINEL: STAGEWISE INTEGRITY VERIFICATION FOR PIPELINE PARALLEL DECENTRALIZED TRAINING

Anonymous authors

Paper under double-blind review

ABSTRACT

Decentralized training introduces critical security risks when executed across untrusted, geographically distributed nodes. While existing Byzantine-tolerant literature addresses data parallel (DP) training through robust aggregation methods, pipeline parallelism (PP) presents fundamentally distinct challenges. In PP, model layers are distributed across workers where the activations and their gradients flow between stages rather than being aggregated, making traditional DP approaches inapplicable. We propose SENTINEL, a verification mechanism for PP training *without computation duplication*. SENTINEL employs lightweight momentum-based monitoring using exponential moving averages (EMAs) to detect corrupted inter-stage communication. Unlike existing Byzantine-tolerant approaches for DP that aggregate parameter gradients *across replicas*, our approach verifies sequential activation/gradient transmission *between layers*. We provide theoretical convergence guarantees for this new setting that recovers classical convergence rates when relaxed to standard training. Experiments demonstrate successful training of billion-parameter LLMs across untrusted distributed environments with hundreds of workers while maintaining model convergence and performance.

1 INTRODUCTION

Large Language Models (LLMs) have fundamentally reshaped artificial intelligence, demonstrating exceptional performance across diverse tasks (OpenAI, 2023; Yang et al., 2024; Jiang et al., 2024; Dubey et al., 2024; MetaAI, 2025; Bi et al., 2024; DeepSeek-AI et al., 2025). Training state-of-the-art LLMs, however, requires substantial computational resources (reportedly tens of thousands of co-located GPUs for models like GPT-4 (Walker II, 2023), Llama-4 (MetaAI, 2025), Qwen2.5 (Yang et al., 2024), etc.) with corresponding energy and financial costs. This has motivated research into decentralized training approaches to broaden participation in LLM development (Ryabinin et al., 2021; Yuan et al., 2022; Ryabinin et al., 2023). Decentralized training, which extends distributed training to trustless settings, allows independent collaborators to pool their computational resources, potentially over large distances, to develop models without relying on massive centralized infrastructure.

Decentralized training of LLMs over networks of interconnected devices is made possible through two primary parallelization approaches: data and pipeline parallelism. Data parallelism (DP) (Li et al., 2020; Zhao et al., 2023) distributes different batches of training data across workers nodes, however requires each node to fit the entire model which is not practical for billion-parameter models in decentralized settings. Pipeline parallelism (PP) (Krizhevsky et al., 2017; Huang et al., 2019) partitions the model across stage-wise across worker nodes, with each responsible for distinct model stages (groups of layers), but requires high-bandwidth connections and suffers from node dropout. Combining these complementary approaches reduces the size limitation of DP and the vulnerability of PP, and have enabled frameworks such as SWARM (Ryabinin et al., 2023) to train billion-parameter LLMs through internet-scale communication among distributed nodes. By leveraging these parallelization techniques, such frameworks aim to achieve high node utilization while minimizing bandwidth requirements, hoping to make large-scale model training more widely accessible.

While optimizing communication bandwidth and fault tolerance have been the primary focus in decentralized training research (Ryabinin et al., 2021; Douillard et al., 2023; 2025; Ajanthan et al., 2025), the success of incentive-driven decentralized training critically hinges on the integrity and trustworthiness of participating nodes (Lu et al., 2024).

054
055
056
057
058
059
060
061
062
063
064
Malicious actors in DP configurations can corrupt global
updates through parameter gradient poisoning, while in PP
(layer-wise model parallelism), adversaries can sabotage
intermediate activations or activation gradients between
model stages. Such vulnerabilities underscore the need for
robust mechanisms in a decentralized PP training setting.
Traditional Byzantine-tolerant methods, designed to prevent simpler DP threat models (Malinovsky et al.,
2024; Gorbunov et al., 2022; Mhamdi et al., 2018) fail to
address the cascading failures induced by partitioned model
execution in pipeline parallel (Lu et al., 2024).

065
066
067
068
069
070
071
072
073
In this paper, we provide the first comprehensive exploration
of secure and verifiable PP decentralized training by identifying
and addressing vulnerabilities unique to this setting. We formalize a suite of malicious attacks tailored
for this setting, for which traditional checkpoint-based veri-
fication is ineffective (Arun et al., 2025). To counter this,
we propose a lightweight verification mechanism using veri-
fier nodes, trusted intermediaries placed between stages, that continuously monitor computational
integrity without requiring full model replication or impeding training throughput.

074
075
076
077
078
079
080
081
082
Our method, called SENTINEL, implements a momentum-based anomaly detection system that tracks
exponential moving averages (EMAs) of activations and gradients across pipeline stages. At each
stage of the model, verifier nodes compute statistical divergence metrics between observed signals
and their EMA baselines. Deviations exceeding adaptively calibrated thresholds, determined via
inter-quartile range (IQR) analysis, are flagged for potential malicious activity. This lightweight veri-
fication introduces minimal computational overhead while enabling early detection of both gradient
and activation tampering attacks. Empirical evaluations demonstrate that our system successfully
detects and mitigates various attacks in decentralized PP setups scaling beyond hundreds of workers,
maintaining training integrity and convergence stability despite the presence of malicious participants.

083
The primary contributions of our work are summarized as follows:

- 084 • We present the first comprehensive study of vulnerabilities unique to decentralized training with
085 hybrid data-pipeline parallelism, and introduce a suite of training-interruption attacks that serve
086 as benchmarks for evaluating the security of future systems.
- 087 • We propose a lightweight verification method, dubbed SENTINEL, that leverages momentum-
088 based monitoring at verifier nodes. Our theoretical analysis demonstrates that undetected
089 malicious workers have a negligible impact on the convergence properties (see Fig. 1).
- 090 • We perform extensive experiments in distributed settings involving hundreds of workers, vali-
091 dating the effectiveness of our verification framework in mitigating malicious behaviors within
092 realistic decentralized training scenarios by achieving consistently high ($> 90\%$) F1 scores.
- 093 • We integrate our method with SWARM parallelism to demonstrate its remarkable versatility in
094 real-world decentralized training ecosystems.

096 2 PROBLEM STATEMENT

098 In this section, we outline our hybrid DP-PP architecture and threat model for decentralized training,
099 focusing on malicious worker behavior. Additional vulnerabilities are detailed in App. B.

101 2.1 THREAT MODEL

103 We consider a distributed pipeline parallel neural network composed of p stages and n worker
104 nodes (see Fig. 2b & c). The network outputs $\mathbf{y} = F(\mathbf{x}) = f_p \circ f_{p-1} \circ \dots \circ f_1(\mathbf{x})$. At itera-
105 tion t the parameters are $\boldsymbol{\theta}_t = (\boldsymbol{\theta}_t^{(1)}, \dots, \boldsymbol{\theta}_t^{(p)}) \in \mathbb{R}^{D_{\text{total}}}$, where each stage s has parameters
106 $\boldsymbol{\theta}_t^{(s)} \in \mathbb{R}^{D_s}$ and $D_{\text{total}} = \sum_{s=1}^p D_s$. Each stage s is replicated across d_s worker nodes oper-
107 ating in parallel, thus $n = \sum_{s=1}^p d_s$. Parameters $\boldsymbol{\theta}^{(s)}$ are shared across stage replicas, however

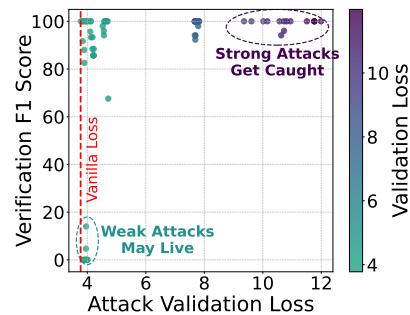


Figure 1: Scatter plot of F1-scores vs. validation loss for more than 75 different attack setups. Strong attacks (higher loss) are caught more often (high F1-score), while weak attacks may slip through. Our verification method thus catches the most harmful attacks that would disrupt training.

that continuously monitor computational integrity without requiring full model replication or impeding training throughput.

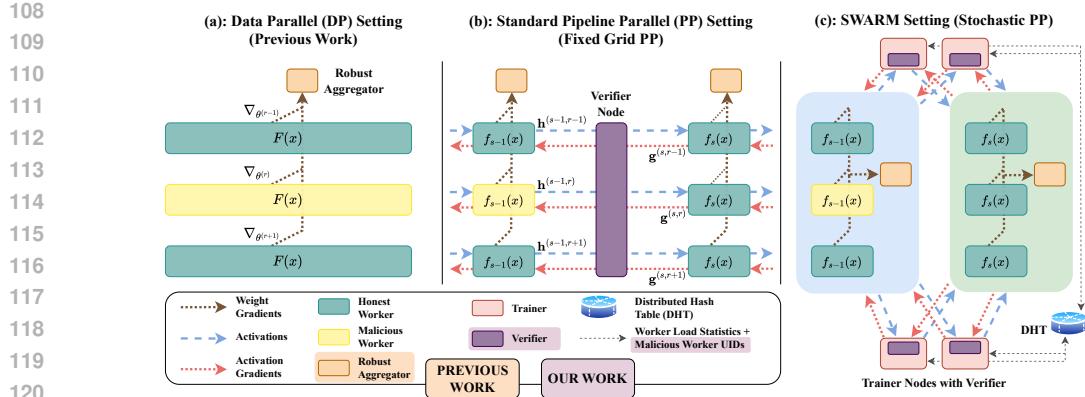


Figure 2: **Distributed threat models.** (a) In DP, workers hold full model replicas and only send *weight gradients*. **Traditional Byzantine-tolerant methods consider this case and use robust aggregation.** (b & c) The threat model considered in this paper (see Sec. 2.1). In PP, workers hold individual layers, and send intermediate *activations* h and *activation gradients* g , thus corruptions directly affect other workers. (b) In the standard setting there is a fixed grid of pipeline stages and data parallel replicas, and communication is routed through our verifier nodes. (c) In the SWARM setting designed for decentralized setups, data is stochastically routed by trainer nodes. Workers send their computations (h and g) to trainers, who then route them to an available worker in the next stage. Our proposed verifiers can seamlessly be added on top of these trainer nodes (App. G).

each replica processes a distinct mini-batch. Thus, worker (s, r) computes replica-specific activations $h^{(s,r)} = f_s(h^{(s-1,r)}; \theta^{(s)}) \in \mathbb{R}^m$ on the forward pass, and replica-specific activation gradients $g^{(s,r)} = \nabla_{h^{(s-1,r)}} \mathcal{L}(\theta)$ on the backward pass where $\mathcal{L}(\theta)$ is the training loss.

In current decentralized PP training frameworks such as SWARM (Ryabinin et al., 2023), workers exchange activations and activation gradients between pipeline stages without verification. Unlike federated learning (McMahan et al., 2017) where attacks primarily target weight gradient poisoning, PP setups are vulnerable to training-interruption attacks, where malicious workers can silently disrupt training by sending corrupted signals. This is particularly challenging because corruptions in early-stage workers *may only become apparent in workers of later layers*: errors can amplify due to model non-linearities and only surpass detection thresholds after several stages, allowing attackers to avoid detection and potentially flagging honest workers due to this *cascading effect*. For a comprehensive comparison between the DP and PP setup, please refer to Q0 in App. A.

To address these vulnerabilities, we introduce dedicated “verifier nodes” as trusted intermediaries for lightweight verification (see Fig. 2b). These verifier nodes intercept and validate all training signals exchanged between stages, operating efficiently even on CPU hardware with minimal overhead. Their key advantage is localizing malicious behavior at specific stages, preventing cascading corruption effects inherent to pipeline parallelism. In frameworks like SWARM(Ryabinin et al., 2023), verifier nodes are a simple extension to “trainer nodes” that handle orchestration between workers (see Fig. 2c, App. G). Formally, our threat model with verifier nodes is defined as follows:

Definition 1 (Data and Pipeline Parallel Threat Model). *Consider a neural network trained in a distributed system with the PP setup described above. We position trusted “verifier” nodes between consecutive stages, through which all communication is routed and thus can be validated, and our goal is to detect and exclude malicious participants. Explicitly, let $B_s \subset \{1, 2, \dots, d_s\}$ be the subset of malicious worker nodes at stage s , with the fraction of malicious workers defined as $\gamma_s = |B_s|/d_s$, and we assume $\gamma_s < 1/2$. Our objective is to detect the malicious subset B_s .*¹

In our threat model, malicious actors can employ various strategies to disrupt training. Next, we introduce these attacks based on their computational requirements and potential impact on the training.

Attack Variants. Let x represent any activation or gradient vector potentially manipulated by a compromised node, with \hat{x} denoting its manipulated version. We consider the following attacks:

- **Constant Attack:** The attacker submits a constant vector without performing the assigned computation, i.e., $\hat{x} = c$, such as vectors of negative ones or zeros.

¹We assume the first and last layers are managed by honest workers, as they process user data and compute the loss and thus require protection. We also assume that the malicious workers cannot collude with each other.

- 162 • **Random Value Attack:** The attacker replaces vector elements with values randomly drawn
163 from a standard normal distribution: $\hat{\mathbf{x}} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$.
- 164 • **Scaling Attack:** The attacker scales the vector by a factor α , i.e., $\hat{\mathbf{x}} = \alpha \mathbf{x}$.
- 165 • **Random Sign Attack:** The attacker flips each vector element's sign with probability p .
- 166 • **Bias Addition Attack:** The attacker introduces random noise to the vector: $\hat{\mathbf{x}} = \mathbf{x} + \boldsymbol{\epsilon}$, where
167 $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$. For stealthier attacks, σ can be chosen to match the original vector's magnitude.
- 168 • **Delay Attack:** The attacker use past values: $\hat{\mathbf{x}}_t = \mathbf{x}_{t-k}$ where k denotes the delay steps
- 169 • **Invisible Noise Attack:** Inspired by the ALIE attack (Baruch et al., 2019), the attacker replaces
170 benign values with statistically subtle boundary values: $\hat{\mathbf{x}} = \boldsymbol{\mu} + z_{\max}(\boldsymbol{\sigma} \odot \boldsymbol{\epsilon})$, where $\boldsymbol{\mu}$ is the
171 original vector's mean, $\boldsymbol{\sigma}$ its element-wise standard deviation, $z_{\max} = \sqrt{2} \cdot \text{erfinv}(2p - 1)$, with
172 $p = 1 - \alpha$ being the quantile threshold, erfinv is the inverse error function, and $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$.
- 173

175 3 MOMENTUM-BASED VERIFICATION OF WORKER NODES

177 Vanilla pipeline parallelism remains vulnerable to corrupted communications in both forward and
178 backward propagation. When malicious workers inject corrupted activations or activation gradients,
179 the effects can cascade through the network, potentially compromising model convergence with
180 minimal detectability. Naïve methods like full computation duplication (Rajput et al., 2019; Lu et al.,
181 2024) guarantee detection but reduce training throughput, while random sampling verification fails as
182 some attacks that can damage training within just a few iterations (see Fig. 7).

183 We introduce SENTINEL: a lightweight, statistically principled verification mechanism that leverages
184 EMA of activations and their gradients to establish reliable reference points for detecting anomalous
185 behavior. We design an algorithm to adaptively set the thresholds of our anomaly detection tests
186 using the IQR. Under relaxed assumptions, we analytically prove that undetected corrupted worker
187 nodes under our verification framework have negligible impact on final model convergence.

189 3.1 PROPOSED METHOD

191 **Motivation.** The key insight driving our approach is that in healthy distributed training scenarios,
192 each worker's activations and gradients should exhibit statistical consistency with the overall
193 population. Existing work duplication methods such as Lu et al. (2024) would require significant
194 computational resources as they need to allocate half of their resources for work verification. EMAs
195 offer three critical advantages as the foundation for our proposed lightweight verification:

- 196 • **Computational Efficiency:** Computing and updating EMA statistics requires minimal computa-
197 tion and memory overhead ($\mathcal{O}(m)$ complexity where m is the activation/gradient size), making
198 it suitable for resource-constrained verifier nodes.
- 199 • **Temporal Smoothing:** The EMA naturally smooths out mini-batch noise while capturing the
200 underlying distribution of legitimate worker outputs, creating a robust reference point.
- 201 • **Adaptivity to Training Dynamics:** As training distributions shift, the EMA automatically
202 adjusts to these shifts while remaining resistant to abrupt deviations from malicious workers.

203 Thus, we design SENTINEL to contain four key components: (1) using EMAs as statistical reference
204 points, (2) selecting appropriate distance measures for deviation detection, (3) implementing adaptive
205 thresholds for anomaly detection, and (4) handling cascading effects in the distributed architecture.
206 Below, we elaborate on each of these components. For a detailed step-by-step overview of our
207 approach, please refer to Alg. 1 in the Appendix. Furthermore, we refer the interested reader to
208 App. G for an overview of how SENTINEL gets integrated in SWARM (Ryabinin et al., 2023).

209 **Exponential Moving Average as Reference Point.** We leverage the EMA of activations and
210 gradients at each layer as a statistical reference point to detect deviations. Since the EMA serves as a
211 robust approximation of the expected value (Robbins & Monro, 1951), it is effective for detecting
212 anomalies. For non-malicious actors, we only expect small deviations since the optimization trajectory
213 typically remains smooth. Formally, each verifier node maintains a running EMA of activations:

$$214 \mathbf{m}_t^{(s)}(\mathbf{h}) = \beta_h \mathbf{m}_{t-1}^{(s)}(\mathbf{h}) + (1 - \beta_h) \frac{1}{d_s} \sum_{r=1}^{d_s} \mathbf{h}_t^{(s,r)}, \quad (1)$$

216 where $\mathbf{m}_t^{(s)}(\mathbf{h})$ denotes the EMA of activations at stage s , and $\beta_h \in [0, 1]$ is the decay rate. A
 217 **similar equation is used to capture the EMA of gradients $\mathbf{m}_t^{(s)}(\mathbf{g})$ with decay rate $\beta_g \in [0, 1]$.**
 218 To establish a reliable initial estimate, we employ a “warm-up” phase with only honest workers,
 219 during which the EMA statistics are collected and stabilized before verification begins (for additional
 220 explanation, please see Q4 in App. A.) After the warm-up period, each time workers submit new
 221 signals, the verifier conducts a lightweight statistical test comparing these signals with the established
 222 EMA. The deviation determines whether a worker is flagged as malicious. Formally, for activation
 223 $\mathbf{h}_t^{(s,r)}$ submitted by worker (s, r) , the verifier calculates:

$$\Gamma_t^{(s,r)} := \Omega\left(\mathbf{h}_t^{(s,r)}, \mathbf{m}_{t-1}^{(s)}(\mathbf{h})\right), \quad (2)$$

225 where $\Omega(\cdot, \cdot)$ is a suitable distance measure. A worker is flagged as malicious if $\Gamma_t^{(s,r)} > \tau$, and we
 226 skip updating the EMA to maintain verification integrity. We use a similar detector for $\mathbf{g}_t^{(s,r)}$.
 227

228 **Choice of Distance Measure Ω .** The distance function critically impacts verification sensitivity.
 229 Rather than using a single metric, we employ a collection of metrics \mathcal{M} , including *mean absolute*
 230 *difference*, *normalized Euclidean distance*, *sliced Wasserstein distance*, and *sign flip ratio*, to robustly
 231 detect various attack types. **These metrics capture different aspects of distributional shifts that may**
 232 **indicate malicious behavior:**

- 233 • **Mean Absolute Difference (L_1)** measures the average absolute deviation between signals:
 $\Omega_{L_1}(\mathbf{h}_t^{(s,r)}, \mathbf{m}_{t-1}^{(s)}(\mathbf{h})) = \mathbb{E}[\|\mathbf{h}_t^{(s,r)} - \mathbf{m}_{t-1}^{(s)}(\mathbf{h})\|_1]$, where $\mathbb{E}[\cdot]$ denotes the expectation over all
 234 features.
- 235 • **Normalized Euclidean Distance (L_2)** computes the squared difference between whitened
 236 representations: $\Omega_{L_2}(\mathbf{h}_t^{(s,r)}, \mathbf{m}_{t-1}^{(s)}(\mathbf{h})) = \mathbb{E}[\|\bar{\mathbf{h}}_t^{(s,r)} - \bar{\mathbf{m}}_{t-1}^{(s)}(\mathbf{h})\|^2]$, where $\bar{\mathbf{x}}$ denotes the
 237 whitened (z-scored) version of \mathbf{x} .
- 238 • **Sign Flip Ratio (SFR)** quantifies the fraction of coordinates with opposing signs, bounded in
 239 $[0, 1]$: $\Omega_{SFR}(\mathbf{h}_t^{(s,r)}, \mathbf{m}_{t-1}^{(s)}(\mathbf{h})) = \mathbb{E}[1(\text{sign}(\mathbf{h}_t^{(s,r)}) \neq \text{sign}(\mathbf{m}_{t-1}^{(s)}(\mathbf{h})))]$.
- 240 • **Sliced Wasserstein Distance (SW)** approximates the Wasserstein distance through random
 241 projections, i.e., $\Omega_{SW}(\mathbf{h}_t^{(s,r)}, \mathbf{m}_{t-1}^{(s)}(\mathbf{h})) = \mathbb{E}_{\mathbf{u} \sim \mathcal{S}^{d-1}}[W_1((\mathbf{h}_t^{(s,r)})_{\#} \mathbf{u}, (\mathbf{m}_{t-1}^{(s)}(\mathbf{h}))_{\#} \mathbf{u})]$, where
 242 W_1 is the 1-Wasserstein distance, and $(\cdot)_{\#}$ denotes the projection onto the random unit vector
 243 \mathbf{u} from the unit sphere \mathcal{S}^{d-1} .

244 We track all these metrics for each stage throughout training. The same metrics are used to measure
 245 the deviations between submitted gradients $\mathbf{g}_t^{(s,r)}$ and their momentum $\mathbf{m}_{t-1}^{(s)}(\mathbf{g})$. A worker is flagged
 246 if it exceeds the threshold for any metric. The diverse set of metrics provides robustness against
 247 various attack vectors, as different attacks may manifest in different statistical properties. Neural
 248 network-based distance measures could also be a promising candidate but left for future work. For an
 249 ablation study on the impact of each metric on the final performance, please refer to App. F.3.

250 **Automatic Threshold τ Updates.** Each distance metric requires a threshold τ to flag anomalous
 251 behavior. During the initial “warm-up” period mentioned earlier, we collect valid deviations for each
 252 metric at every stage and compute the IQR. We then use Tukey’s fences (Tukey, 1977) to establish
 253 valid deviation bounds. In particular, let \mathcal{H}_l^s denote the history of valid deviations over the past l
 254 iterations collected from all d_s workers at stage s . Let q_1 , q_2 , and q_3 be the 25th, 50th (median), and
 255 75th percentiles of these deviations. We define our test statistic as:

$$\text{if } |\Gamma_t^{(s,r)} - q_2| \geq k(q_3 - q_1) \Rightarrow \text{flag node } (s, r) \text{ as malicious} \quad (3)$$

256 where $k = 1.5$ is the conventional choice (Tukey, 1977). We choose to adaptively adjust k through
 257 an iterative process that widens or narrows thresholds to maintain a chosen false positive rate (e.g., $<$
 258 1%). This dynamic threshold continuously incorporates new benign deviations into the historical
 259 window, enabling the verification system to automatically adapt to evolving data changes throughout
 260 training (see Fig. 12). The details of this method is given in Alg. 5 in the Appendix.

261 **Handling Cascading Effects.** In our PP-based distributed architecture, corrupted activations from
 262 an early stage can affect downstream workers, potentially causing misclassification of honest nodes.
 263 We address this with two complementary mechanisms:

- 264 1. **Bottom-up Malicious-Node Identification:** When a worker at stage s is flagged as malicious,
 265 the verifier notifies all downstream verifiers to pause their deviation statistics for the affected

270 mini-batch and label subsequent anomalies as “tainted by upstream.” To maintain uninterrupted
 271 training, in the backward pass verifiers send the stored gradient EMA, enabling parameter updates
 272 without revealing any behavioral change. For more information see App. D.2 and Fig. 5.

273 **2. Violation Counter with Forgiveness:** Rather than banning a worker after one deviation, each
 274 verifier maintains a violation counter. Severe deviations ($\times 100$ above the threshold) result in
 275 immediate bans, while milder ones increment the counter by one. A worker is banned after c
 276 violations, but the counter decrements after $T_{\text{forgiveness}}$ consecutive clean steps, allowing recovery
 277 from transient anomalies. We use $c = 5$ and $T_{\text{forgiveness}} = 100$ in our experiments.

279 **3.2 THEORETICAL ANALYSIS**

281 To complement our practical approach, we provide theoretical guarantees under relaxed conditions,
 282 analyzing (1) the convergence behavior of the distributed training under bounded malicious perturba-
 283 tions, and (2) the conditions under which our system can maintain an honest majority at each stage.

285 **Theorem 1** (Convergence Under Bounded Perturbations (informal)). *Consider a distributed
 286 training setup that utilizes PP to split the model layers across workers and uses momentum-based
 287 verification to verify each worker’s contribution in the forward or backward pass.^a Also, assume
 288 that less than half of workers at each stage are malicious (i.e., $\gamma_s < 1/2$) and we use a fixed
 289 threshold τ for worker verification using Eq. (2).^b Under such relaxed conditions, training with
 290 non-convex loss functions optimized with momentum SGD converges to a neighborhood of a
 291 stationary point where the size of this neighborhood is directly proportional to τ .*

293 Theorem 1 states malicious workers who evade detection by keeping perturbations below threshold τ
 294 can only cause the final solution to deviate from the optimal solution by an amount proportional to τ .
 295 The formal theorem is given in App. E.1. Please also visit Q8 in the FAQ (App. A).

296 Recall that SENTINEL relies on the assumption that fewer than half of workers at any stage are
 297 malicious (i.e., $\gamma_s < 1/2$). Next, we quantify the conditions under which this assumption holds with
 298 high probability, given a total budget of malicious workers. The proof is given in App. E.2.

300 **Lemma 1** (Honest Majority Guarantee). *Consider our distributed training system with p pipeline
 301 stages, each replicated across d worker nodes. Let b be the total number of malicious workers,
 302 and $\epsilon \in (0, 1)$ be a small positive constant. If workers are assigned to each stage randomly and
 303 $b \leq dp/2 - p\sqrt{d/2 \ln(p/\epsilon)}$, then with probability at least $1 - \epsilon$ every pipeline stage has strictly
 304 fewer than $d/2$ malicious workers.*

306 **4 RELATED WORK**

309 Our momentum-based verification approach intersects three primary research directions that have
 310 largely evolved independently. For a more comprehensive review of related work, see App. C.

311 **Decentralized LLM Training** has emerged as a democratizing force in AI development. While
 312 frameworks like Tasklets (Yuan et al., 2022) or SWARM (Ryabinin et al., 2023) have made significant
 313 advances in communication efficiency and fault tolerance for non-malicious failures, they remain
 314 vulnerable to adversarial participants. Our work aims to address these vulnerabilities.

315 **Byzantine Robustness in Machine Learning** has traditionally focused on federated learning contexts
 316 where each worker computes complete model updates. Classic approaches like Krum (Blanchard
 317 et al., 2017), Bulyan (Mhamdi et al., 2018), and CENTEREDCLIP (Karimireddy et al., 2021) rely on
 318 comparing full gradients across workers, a fundamental mismatch with pipeline parallel architectures
 319 where each worker computes only a fraction of the model. SENTINEL is specifically designed for the
 320 unique constraints of pipeline parallelism.

321 ^aWe do not consider dishonest activity during the “all-reduce” operation for syncing parameter gradi-
 322 ents between DP replicas (which is the setting that all prior Byzantine-tolerant literature address). This is a
 323 complimentary axis and one can utilize any prior Byzantine-tolerant work.

^bThis is to relax our conditions. In practice we set this threshold automatically each iteration using the IQR.

Table 1: Attack detection performance for Llama-3-0.6B on C4 dataset.

Mode	Attack	Sentinel (Ours)					No Verif.
		Pr. (%) \uparrow	Re. (%) \uparrow	F1 (%) \uparrow	Det. Speed \downarrow	Val. Loss \downarrow	
-	None (Vanilla)	100.0	100.0	100.0	N/A	3.819	3.821
Activation	Scaling ($\alpha = -1$)	100.0	100.0	100.0	6.38	3.824	4.109
	Random Value	100.0	100.0	100.0	6.48	3.827	7.778
	Delay (100-steps)	88.9	100.0	94.1	13.21	3.841	7.675
	Bias Addition	84.6	91.7	88.0	14.57	3.830	3.892
	Invisible Noise (99%)	100.0	100.0	100.0	6.48	3.826	7.682
Gradient	Scaling ($\alpha = -1$)	0.0	0.0	0.0	N/A	3.893	3.893
	Random Value	100.0	100.0	100.0	1.0	3.818	9.595
	Delay (100-steps)	100.0	100.0	100.0	7.33	3.826	10.157
	Bias Addition	100.0	100.0	100.0	1.0	3.828	10.813
	Invisible Noise (99%)	100.0	79.2	88.4	211.0	3.943	4.176

Table 2: Detection performance for training Llama-3-0.6B against mixed attacks.

DATASET	SENTINEL (OURS)					VANILLA
	PR. (%) \uparrow	RE. (%) \uparrow	F1 (%) \uparrow	DET. SPEED \downarrow	VAL. LOSS \downarrow	VAL. LOSS \downarrow
COMMONCRAWL	83.7	92.3	87.8	78.14	3.831	3.821
FINEWEB	81.8	92.3	86.7	66.00	3.827	3.840
OPENWEBTEXT	91.9	87.2	89.5	52.70	3.784	3.778

Secure Distributed Systems principles inform our verifier node architecture, which draws inspiration from trusted intermediaries in distributed computing. While preliminary work by Lu et al. (2024) identified potential vulnerabilities in pipeline parallel training, their redundancy-based solution would significantly reduce throughput, negating the primary benefit of distributed training. Our lightweight verification mechanism provides robust security guarantees with minimal computational overhead.

5 EXPERIMENTAL RESULTS

In this section, we present our experimental results. Unless stated otherwise, we use a 0.6B-parameter Llama-3 (Dubey et al., 2024; Liang et al., 2025) model (16 layers, 32 attention heads, 1024 hidden dimension and context length) distributed across 128 workers in a 8×16 data-pipeline parallel mesh. We use AdamW (Loshchilov & Hutter, 2019) with initial learning rate $6e-4$ as our optimizer and train our models on FineWeb (FW) (Penedo et al., 2024), OpenWebText (OW) (Gokaslan et al., 2019), and Common Crawl (C4) (Raffel et al., 2020) datasets for 5k steps. We randomly designate 25% of workers at each pipeline stage as malicious (2:6 malicious vs. honest ratio), with only 25% of these activated simultaneously to soften our “no-collusion” assumption. Training begins with a 1k-step warm-up period before verification is activated. Based on validation runs on vanilla case, we set $\beta_h = 0.9$ and $\beta_g = 0.8$ for activation and gradient verification. Finally, for our adaptive threshold we use a window of past 100 steps. We relax these assumptions through various ablation studies to study their impact. Detailed experimental settings and extended results are provided in App. F.

Metrics. We evaluate our verification method using precision (Pr), recall (Re), and F1-score to measure effectiveness in detecting malicious workers while minimizing the false positives. To quantify detection efficiency, we report *detection speed* as the average number of iterations between the start of malicious behavior till the malicious worker gets banned. We also compare convergence rates across methods using the average loss at the last training step on a held-out validation set.

Verification Performance. We trained Llama-3-0.6B models on the C4 dataset with and without our verification mechanism and report the results in Tab. 1 (for more comprehensive results on C4 and other datasets, please see App. F.) Our experiments yield three key findings:

1. In pipeline parallelism, activation attacks are as threatening as gradient attacks, but their risk has been neglected in the Byzantine-tolerant literature.
2. Attack effectiveness varies significantly between activation and gradient domains. The same technique can severely disrupt training when targeting activations but have minimal impact when applied to gradients (e.g., invisible noise attacks).
3. Our EMA verification method achieves high F1-scores across attack types. When attacks produce negligible deviations and evade detection, their impact on convergence remains limited, confirming our theoretical analysis in Theorem 1 that undetected attacks can only shift parameters to a neighborhood of the optimum. Fig. 1 also confirms this relationship through a scatter plot of F1-scores against validation loss of more than 75 different attack setups from Tabs. 12 to 14.

378 Table 3: Activation attack detection performance for large-scale Llama-3 training on C4 dataset.
379

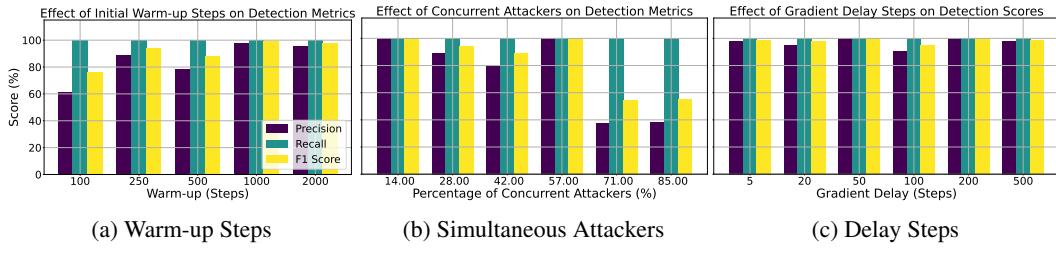
380 SETUP	381 ATTACK	382 SENTINEL (OURS)			
		383 PR. (%) ↑	384 RE. (%) ↑	385 F1 (%) ↑	386 DET. SPEED ↓
387 0.6B w/ 16 × 16 MESH	388 Random Value	389 100.0	390 100.0	391 100.0	392 7.96
	393 Delay (100-steps)	394 85.7	395 100.0	396 92.3	397 14.39
	398 Bias Addition	399 100.0	400 25.6	401 40.8	402 65.15
	403 Invisible Noise (99%)	404 100.0	405 100.0	406 100.0	407 3.981
408 1.2B w/ 8 × 8 MESH	409 Random Value	410 100.0	411 100.0	412 100.0	413 3.898
	414 Delay (100-steps)	415 37.5	416 100.0	417 54.5	418 67.0
	419 Bias Addition	420 0.0	421 0.0	422 0.0	423 N/A
	424 Invisible Noise (99%)	425 100.0	426 100.0	427 100.0	428 3.727

388 Table 4: **Detection performance for training alternative models against mixed activation attacks.**
389

390 ARCHITECTURE	391 SENTINEL (OURS)				392 VANILLA
	393 PR. (%) ↑	394 RE. (%) ↑	395 F1 (%) ↑	396 VAL. LOSS ↓	397 VAL. LOSS ↓
398 LLAMA-4-0.4B	399 73.5	400 92.3	401 81.8	402 3.617	403 3.628
404 DEEPSEEK-V3-1B	405 94.6	406 97.2	407 95.9	408 3.421	409 3.393
410 LLAMA-3-4B	411 94.9	412 68.6	413 79.6	414 3.714	415 3.668

394 Table 5: **Detection performance for training Llama-3-0.6B against mixed attacks.**
395

396 METHOD	397 ATTACKED TRAINING				398 VANILLA
	399 PR. (%) ↑	400 RE. (%) ↑	401 F1 (%) ↑	402 VAL. LOSS ↓	403 VAL. LOSS ↓
404 SENTINEL + Krum	405 93.6	406 80.6	407 86.6	408 3.873	409 3.855
410 SENTINEL + Bulyan	411 85.3	412 80.6	413 82.9	414 3.883	415 3.855

406 Figure 3: Ablation studies on the effect of various elements on the verification performance.
407

408 **Mixed Attacks.** While previous experiments examined individual attack types in isolation, real-world scenarios typically involve adversaries employing various attack strategies simultaneously across different nodes. To evaluate our approach under these more realistic conditions, we conduct experiments with the Llama-3-0.6B model in a distributed training environment where 37.5% of workers per stage are malicious (specifically, 3 malicious versus 5 honest workers per stage in an 8 × 16 mesh). All attackers begin their malicious activities simultaneously, with each attacker randomly assigned both an attack mode (activation or gradient manipulation) and a specific attack strategy from Sec. 2. Tab. 2 summarizes our experimental results across three datasets. The findings demonstrate that our verification method successfully identifies most detrimental attacks, resulting in validation loss metrics comparable to non-attacked baseline models. This robustness against heterogeneous attack vectors highlights the effectiveness of our approach in securing pipeline parallel-based LLM training against sophisticated adversarial scenarios.

409 **Large-scale Experiments.** To validate our approach at scale, we conduct experiments in two other 410 settings: (1) a 16 × 16 mesh configuration with 256 total workers, and (2) a larger 1.2B parameter 411 model trained on an 8 × 8 mesh with 64 workers. In both scenarios, we maintained 37.5% malicious 412 workers per stage (6:10 malicious-to-honest ratio for the 16 × 16 mesh and 3:5 for the 8 × 8 mesh), 413 with 25% of attackers active during each attack round. Results in Tab. 3 (see Tab. 15 for extended 414 results) show that our verification mechanism effectively preserves training integrity across all attack 415 scenarios, with validation losses comparable to non-attacked baselines. These results were achieved 416 without extensive hyperparameter tuning compared to our 0.6B setting, though improvements are 417 needed to reduce false positives in activation delay attacks on the 1.2B model.

418 **Ablation Studies.** Next, we study how key factors affect our verification method in detecting 419 malicious workers. We evaluate using our strongest attacks, 99% invisible noise activation attacks 420 (Studies 1 & 2), 100-step delayed gradient attacks (Study 3), and a mix of all activation attacks (Study 421 4 & 5), with a consistent ratio of 3 malicious to 5 honest workers per stage (37.5% malicious).

432 1. **Initial Warm-up Period:** Our method requires an initial warm-up phase to ensure that training
 433 has reached a stable point. Fig. 3a shows that while early detection achieves high recall, precision
 434 is initially low due to insufficiently robust thresholds because of training volatility. After roughly
 435 1k steps, precision stabilizes as the verification method establishes reliable bounds.

436 2. **Attacker Collusion:** When malicious workers coordinate their attacks, detection becomes more
 437 challenging. As shown in Fig. 3b, our verification method maintains effectiveness with up to 60%
 438 collusion among malicious workers. Beyond this threshold, false positives increase significantly,
 439 which could adversely impact precision.

440 3. **Gradient Delay Impact:** Fig. 3c demonstrates our method’s robustness to various delay lengths
 441 in gradient attacks. Even with minimal delays, where malicious gradients closely resemble
 442 current legitimate gradients, our verification method maintains high detection rates.

443 4. **Alternative Architectures:** To demonstrate the transferability of SENTINEL across transformer
 444 architectures, we train two Mixture-of-Expert (MoE) (Shazeer et al., 2017) architectures, namely
 445 Llama-4-0.4B (MetaAI, 2025) and DeepSeek-V3-1B (DeepSeek-AI, 2024), on the FineWeb-
 446 EDU (Penedo et al., 2024) dataset. We assume that 33% of the adversaries send their malicious
 447 vectors simultaneously as a form of collusion. We train these models as well as their vanilla
 448 baselines till 5000 steps. To demonstrate the transferability of our selected hyperparameters, we
 449 use the same hyperparameters used for the Llama-3-0.6B experiments for our verification. As
 450 shown in Tab. 4, SENTINEL can successfully be applied to these MoE-based models, showcasing
 451 its remarkable transferability across architectures. For more details, please see App. F.

452 5. **Large-scale Models:** Finally, we show that increasing model size from medium scale to large
 453 scale can be done without major performance degradation. To this end, we extend our Llama-
 454 3-0.6B experiment by increasing the hidden dimension and number of layers to bring the total
 455 parameter count to around 4B. We then train this model using a stringent total batch-size of 96 to
 456 fit within our available computational resources and train for 5k steps on the FineWeb dataset.
 457 As seen in Tab. 4, SENTINEL protects the training from interruption attacks and delivers a
 458 model which has a validation loss close to the vanilla, non-attacked baseline (0.04 difference
 459 only) while keeping the detection accuracy around 80%. This highlights that our findings are
 460 scalable to larger models with wider hidden dimension and deeper layers.

461 For more results on the impact of EMA, longer runs, datasets, etc. please see App. F.2.

462 **Adaptive Attacks.** In this section, we
 463 investigate an adaptive attack that knows
 464 how SENTINEL is verifying its signals.
 465 In particular, let us assume that the
 466 malicious workers maintain an EMA of sig-
 467 nals sent to subsequent layers, using the
 468 same β as the verifier node. After collecting sufficient EMA samples, the attack sends drifted
 469 activations/gradients biased toward a predetermined target direction, using stale momentum esti-
 470 mates to create consistent bias while appearing statistically legitimate. The attacker constructs its
 471 attacks according to $\hat{x} = m_{t-\delta} + \alpha \cdot (t_{\text{drift}} - m_{t-\delta}) / \|m_{t-\delta}\| + \epsilon$ where $m_{t-\delta}$ is the stale EMA
 472 momentum, α is the drift rate, t_{drift} is the predetermined drift target, $\epsilon \sim \mathcal{N}(0, \sigma^2)$ is Gaussian noise
 473 for stealth, and $\delta = \left\lceil \frac{\log(0.1)}{\log(\beta)} \right\rceil$ is the delay factor. We run this attack for activation and activation
 474 gradients for training a Llama-3-0.6B on C4 dataset. For activation attack, the verifier/attacker uses
 $\beta = 0.9$, while for gradient attack, they use $\beta = 0.8$.

475 Running this attack using our settings from Tab. 1 (25% malicious workers at each stage), we get
 476 the results in Tab. 6. As seen, the adaptive EMA attack can be destructive without verification,
 477 SENTINEL detects and mitigates it perfectly, validating our resilience against adaptive attacks that
 478 have a knowledge of our defense. This is because the attacker’s EMA would only comprise part of
 479 the total true EMA, and assuming an honest majority, this would not be sufficient to interrupt training.

480 **Integration with DP Defenses.** As elaborated throughout the paper, SENTINEL targets verifying the
 481 signals sent between stages in PP. Unlike the PP axis, we pointed out that existing Byzantine-tolerant
 482 literature exclusively target attacks that happen during gradient averaging (a.k.a. all-reduce). As such,
 483 these two are complementary axes that operate on different dimensions (please see Q0 in App. A).
 484 To demonstrate the complementary nature of these two axes, we use two commonly used robust
 485 aggregators, particularly Krum (Blanchard et al., 2017) and Bulyan (Mhamdi et al., 2018), to replace
 the vanilla all-reduce between workers. We then train our Llama-3-0.6B against mixed attacks with

486 3:5 malicious to honest ratio for 5k steps on FineWeb. As seen in Tab. 5, addition of these methods
 487 on the orthogonal DP axis has no significant impact on the convergence of the model under mixed
 488 attacks, supporting our claim that the PP and DP solutions are complementary.
 489

490 5.1 EXTENDING SENTINEL TO SWARM PARALLELISM 491

492 Finally, we adapt SENTINEL to verify worker node signals in a realistic SWARM (Ryabinin et al.,
 493 2023) experiment. SWARM parallelism provides a fault-tolerant distributed training ecosystem
 494 powered by the Hivemind framework (Ryabinin & Gusev, 2020). It comprises of worker nodes
 495 distributed across both DP and PP coordinates. At each stage, a pool of workers process batches
 496 of data, with the coordination managed through trainer nodes that are responsible for stochastically
 497 routing activations in the forward pass and activation gradients in the backward pass. From this
 498 standpoint, SWARM parallelism is akin to a stochastic DP/PP mesh in comparison to the fixed setting
 499 that we have considered so far. As discussed in Sec. 2.1, trainer nodes are in a natural position to be
 500 extended as verifier nodes when augmented with SENTINEL.

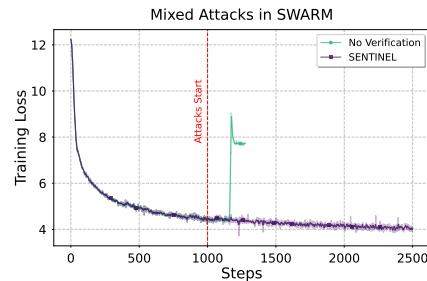
501 For this experiment, we train our Llama-3-0.6B
 502 model across a distributed SWARM configuration
 503 with 128 workers (8×16 mesh). We employ 32
 504 trainer nodes with verification capability to train
 505 our model on FineWeb-EDU, a curated subset of
 506 FineWeb. Since trainers do not have P2P communica-
 507 tion, each maintain independent EMAs with a single
 508 synchronization point at the end of the warm-up.

509 We evaluate robustness against a mixture of random
 510 attacks by designating 37.5% of workers at each stage
 511 (except the first and last two) as malicious, maintain-
 512 ing a 3:5 malicious-to-honest ratio with 15% col-
 513 lusion where attackers activate simultaneously. As
 514 shown in Fig. 4, the presence of malicious work-
 515 ers significantly disrupts training convergence in the
 516 absence of verification. However, SENTINEL suc-
 517 cessfully maintains the integrity of the training, enabling
 518 the training to continue without interruption. This result in a production-grade environment demon-
 519 strates the practical applicability of our approach for securing real-world decentralized training
 520 ecosystems. For implementation details and full results of SENTINEL in SWARM, see App. G.

521 6 CONCLUSION 522

523 In this paper, we investigated security vulnerabilities in decentralized, pipeline parallel networks,
 524 showing how malicious workers can corrupt activations and activation gradients exchanged between
 525 pipeline stages. To guide future research, we introduced a suite of training-interruption attacks as
 526 benchmarks for evaluating decentralized training security. Our key contribution is SENTINEL, a
 527 lightweight momentum-based verification mechanism that utilizes trusted verifier nodes to maintain
 528 EMAs of transmitted signals (activations and activation gradients) as statistical reference points. We
 529 further developed an IQR-based adaptive thresholding strategy to automatically calibrate detection
 530 sensitivity. We complement our approach by theoretical analysis and real-world integration for
 531 decentralized training using SWARM. Through extensive experiments with models up to 1.2B
 532 parameters distributed across hundreds of workers, we demonstrated its effectiveness in maintaining
 533 training integrity with consistently high F1 scores ($> 85\%$) across various attack scenarios.

534 **Limitations:** While our verification method effectively detects the attack types presented, it may not
 535 generalize to all possible adversarial strategies. Future work should explore better adaptive detection
 536 mechanisms that require less manual tuning, possibly neural networks for anomaly detection (Pang
 537 et al., 2021). Additionally, our approach addresses inter-stage attacks specific to pipeline parallelism,
 538 but decentralized training remains vulnerable to other threats including backdoor attacks (Li et al.,
 539 2024), membership inference (Shokri et al., 2017), and all-reduce gradient poisoning attacks (Gor-
 bunov et al., 2022) each presenting distinct challenges for truly open collaborative training.



509 Figure 4: Loss when training Llama-3-0.6B mod-
 510 els using SWARM (Ryabinin et al., 2023) with 128
 511 distributed workers on preemptible AWS instances.
 512 Workers employ various activation/gradient ma-
 513 nipulation attacks to disrupt training. While in
 514 the absence of verification training gets disrupted,
 515 SENTINEL can successfully protect training from
 516 divergence.

540 REPRODUCIBILITY STATEMENT
541542 We provide a detailed, step-by-step pseudo-code of our methodology in Algs. 1 to 5 and 7. Ad-
543 ditionally, we give a detailed overview of the hyper-parameters used in our empirical evaluations
544 in App. F.1. We are planning to release the code upon acceptance of the paper.
545546 REFERENCES
547548 Thalaiyasingam Ajanthan, Sameera Ramasinghe, Yan Zuo, Gil Avraham, and Alexander Long. Nes-
549 terov method for asynchronous pipeline parallel optimization. In *Proceedings of the International
550 Conference on Machine Learning (ICML)*, 2025.551 Arasu Arun, Adam St. Arnaud, Alexey Titov, Brian Wilcox, Viktor Kolobaric, Marc Brinkmann,
552 Oguzhan Ersoy, Ben Fielding, and Joseph Bonneau. Verde: Verification via refereed delegation for
553 machine learning programs. *CoRR*, abs/2502.19405, 2025.555 Gilad Baruch, Moran Baruch, and Yoav Goldberg. A little is enough: Circumventing defenses for
556 distributed learning. In *Proceedings of the Annual Conference on Neural Information Processing
557 Systems (NeurIPS)*, pp. 8632–8642, 2019.558 Xiao Bi, Deli Chen, Guanting Chen, Shanhua Chen, Damai Dai, Chengqi Deng, Honghui Ding,
559 Kai Dong, Qiushi Du, Zhe Fu, Huazuo Gao, Kaige Gao, Wenjun Gao, Ruiqi Ge, Kang Guan, Daya
560 Guo, Jianzhong Guo, Guangbo Hao, Zhewen Hao, Ying He, Wenjie Hu, Panpan Huang, Erhang Li,
561 Guowei Li, Jiashi Li, Yao Li, Y. K. Li, Wenfeng Liang, Fangyun Lin, Alex X. Liu, Bo Liu, Wen
562 Liu, Xiaodong Liu, Xin Liu, Yiyuan Liu, Haoyu Lu, Shanghao Lu, Fuli Luo, Shirong Ma, Xiaotao
563 Nie, Tian Pei, Yishi Piao, Junjie Qiu, Hui Qu, Tongzheng Ren, Zehui Ren, Chong Ruan, Zhangli
564 Sha, Zhihong Shao, Junxiao Song, Xuecheng Su, Jingxiang Sun, Yaofeng Sun, Minghui Tang,
565 Bingxuan Wang, Peiyi Wang, Shiyu Wang, Yaohui Wang, Yongji Wang, Tong Wu, Y. Wu, Xin Xie,
566 Zhenda Xie, Ziwei Xie, Yiliang Xiong, Hanwei Xu, R. X. Xu, Yanhong Xu, Dejian Yang, Yuxiang
567 You, Shuiping Yu, Xingkai Yu, B. Zhang, Haowei Zhang, Lecong Zhang, Liyue Zhang, Mingchuan
568 Zhang, Minghua Zhang, Wentao Zhang, Yichao Zhang, Chenggang Zhao, Yao Zhao, Shangyan
569 Zhou, Shunfeng Zhou, Qihao Zhu, and Yuheng Zou. DeepSeek LLM: scaling open-source language
570 models with longtermism. *CoRR*, abs/2401.02954, 2024. doi: 10.48550/ARXIV.2401.02954. URL
571 <https://doi.org/10.48550/arXiv.2401.02954>.572 Battista Biggio, Blaine Nelson, and Pavel Laskov. Poisoning attacks against support vector machines.
573 In *Proceedings of the International Conference on Machine Learning (ICML)*, 2012.574 Peva Blanchard, El Mahdi El Mhamdi, Rachid Guerraoui, and Julien Stainer. Machine learning with
575 adversaries: Byzantine tolerant gradient descent. In *Proceedings of the Annual Conference on
576 Neural Information Processing Systems (NeurIPS)*, pp. 119–129, 2017.577 Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal,
578 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel
579 Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler,
580 Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott
581 Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya
582 Sutskever, and Dario Amodei. Language models are few-shot learners. In *Proceedings of the
583 Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2020.584 DeepSeek-AI. Deepseek-v3 technical report. *CoRR*, abs/2412.19437, 2024. URL <https://doi.org/10.48550/arXiv.2412.19437>.585 DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu,
586 Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu,
587 Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao
588 Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan,
589 Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao,
590 Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding,
591 Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang
592 Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong,
593

594 Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao,
 595 Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang,
 596 Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang,
 597 Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L.
 598 Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhua Chen, Shengfeng Ye, Shiyu Wang,
 599 Shuiping Yu, Shufeng Zhou, Shuteng Pan, and S. S. Li. DeepSeek-R1: Incentivizing reasoning
 600 capability in llms via reinforcement learning. *CoRR*, abs/2501.12948, 2025.

601 Arthur Douillard, Qixiang Feng, Andrei A. Rusu, Rachita Chharia, Yani Donchev, Adhiguna
 602 Kuncoro, Marc’Aurelio Ranzato, Arthur Szlam, and Jiajun Shen. Diloco: Distributed low-
 603 communication training of language models. *CoRR*, abs/2311.08105, 2023.

604 Arthur Douillard, Yanislav Donchev, Keith Rush, Satyen Kale, Zachary Charles, Zachary Gar-
 605 rett, Gabriel Teston, Dave Lacey, Ross McIlroy, Jiajun Shen, Alexandre Ramé, Arthur Szlam,
 606 Marc’Aurelio Ranzato, and Paul Barham. Streaming diloco with overlapping communication:
 607 Towards a distributed free lunch. *CoRR*, abs/2501.18512, 2025.

608 Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha
 609 Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn,
 610 Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston
 611 Zhang, Aurélien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Rozière, Bethany Biron,
 612 Bin Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris
 613 McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton
 614 Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny Livshits, David
 615 Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes,
 616 Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip
 617 Radenovic, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Graeme
 618 Nail, Grégoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu,
 619 Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel M. Kloumann, Ishan Misra, Ivan
 620 Evtimov, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet
 621 Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi,
 622 Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph
 623 Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Kartikeya Upasani,
 624 Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, and et al. The Llama 3 herd of models.
 625 *CoRR*, abs/2407.21783, 2024.

626 Clement Fung, Chris J. M. Yoon, and Ivan Beschastnikh. The limitations of federated learning in
 627 sybil settings. In *Proceedings of the International Symposium on Research in Attacks, Intrusions*
 628 and *Defenses (RAID)*, 2020.

629 Aaron Gokaslan, Vanya Cohen, Ellie Pavlick, and Stefanie Tellex. Openwebtext corpus. <http://Skylion007.github.io/OpenWebTextCorpus>, 2019.

630 Eduard Gorbunov, Alexander Borzunov, Michael Diskin, and Max Ryabinin. Secure distributed
 631 training at scale. In *Proceedings of the International Conference on Machine Learning (ICML)*, pp.
 632 7679–7739, 2022.

633 Eduard Gorbunov, Samuel Horváth, Peter Richtárik, and Gauthier Gidel. Variance reduction is an
 634 antidote to byzantines: Better rates, weaker assumptions and communication compression as a
 635 cherry on the top. In *Proceedings of the International Conference on Learning Representations*
 636 (ICLR), 2023.

637 Lie He, Sai Praneeth Karimireddy, and Martin Jaggi. Byzantine-robust learning on heterogeneous
 638 datasets via resampling. *CoRR*, abs/2006.09365, 2020.

639 Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *The collected*
 640 *works of Wassily Hoeffding*, pp. 409–426, 1994.

641 Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza
 642 Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom
 643 Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy,
 644 Simon Osindero, Karen Simonyan, Erich Elsen, Jack W. Rae, Oriol Vinyals, and Laurent Sifre.
 645 Training compute-optimal large language models. *CoRR*, abs/2203.15556, 2022.

648 Yanping Huang, Youlong Cheng, Ankur Bapna, Orhan Firat, Dehao Chen, Mia Xu Chen, HyoukJoong
 649 Lee, Jiquan Ngiam, Quoc V. Le, Yonghui Wu, and Zhifeng Chen. GPipe: Efficient training of
 650 giant neural networks using pipeline parallelism. In *Proceedings of the Annual Conference on*
 651 *Neural Information Processing Systems (NeurIPS)*, pp. 103–112, 2019.

652 Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris
 653 Bamford, Devendra Singh Chaplot, Diego de Las Casas, Emma Bou Hanna, Florian Bressand,
 654 Gianna Lengyel, Guillaume Bour, Guillaume Lample, Lélio Renard Lavaud, Lucile Saulnier, Marie-
 655 Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, Szymon Antoniak, Teven Le
 656 Scao, Théophile Gervet, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed.
 657 Mixtral of experts. *CoRR*, abs/2401.04088, 2024.

658 Sai Praneeth Karimireddy, Lie He, and Martin Jaggi. Learning from history for byzantine robust
 659 optimization. In *Proceedings of the International Conference on Machine Learning (ICML)*, pp.
 660 5311–5319, 2021.

662 Sai Praneeth Karimireddy, Lie He, and Martin Jaggi. Byzantine-robust learning on heterogeneous
 663 datasets via bucketing. In *Proceedings of the International Conference on Learning Representations*
 664 (*ICLR*), 2022.

666 Andrej Karpathy. NanoGPT. <https://github.com/karpathy/nanoGPT>, 2022.

667 Anastasia Koloskova, Nikita Doikov, Sebastian U. Stich, and Martin Jaggi. On convergence of
 668 incremental gradient for non-convex smooth functions. In *Proceedings of the International*
 669 *Conference on Machine Learning (ICML)*, 2024.

671 Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. ImageNet classification with deep convolutional
 672 neural networks. *Communications of the ACM*, 60(6):84–90, 2017.

674 Quentin Lhoest, Albert Villanova del Moral, Yacine Jernite, Abhishek Thakur, Patrick von Platen,
 675 Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, Joe Davison, Mario
 676 Šaško, Gunjan Chhablani, Bhavya Malik, Simon Brandeis, Teven Le Scao, Victor Sanh, Canwen
 677 Xu, Nicolas Patry, Angelina McMillan-Major, Philipp Schmid, Sylvain Gugger, Clément Delangue,
 678 Théo Matussière, Lysandre Debut, Stas Bekman, Pierrick Cistac, Thibault Goehringer, Victor
 679 Mustar, François Lagunas, Alexander Rush, and Thomas Wolf. Datasets: A community library for
 680 natural language processing. In *Proceedings of the Conference on Empirical Methods in Natural*
 681 *Language Processing (EMNLP): System Demonstrations*, November 2021.

682 Shen Li, Yanli Zhao, Rohan Varma, Omkar Salpekar, Pieter Noordhuis, Teng Li, Adam Paszke, Jeff
 683 Smith, Brian Vaughan, Pritam Damania, and Soumith Chintala. PyTorch distributed: Experiences
 684 on accelerating data parallel training. *Proceedings of the VLDB Endowment*, 13(12):3005–3018,
 685 2020.

686 Yige Li, Hanxun Huang, Yunhan Zhao, Xingjun Ma, and Jun Sun. BackdoorLLM: A comprehensive
 687 benchmark for backdoor attacks on large language models. *CoRR*, abs/2408.12798, 2024.

688 Wanchao Liang, Tianyu Liu, Less Wright, Will Constable, Andrew Gu, Chien-Chin Huang, Iris
 689 Zhang, Wei Feng, Howard Huang, Junjie Wang, Sanket Purandare, Gokul Nadathur, and Stratos
 690 Idreos. TorchTitan: One-stop PyTorch native solution for production ready LLM pretraining. In
 691 *Proceedings of the International Conference on Learning Representations (ICLR)*, 2025.

693 Yanli Liu, Yuan Gao, and Wotao Yin. An improved analysis of stochastic gradient descent with
 694 momentum. In *Proceedings of the Annual Conference on Neural Information Processing Systems*
 695 (*NeurIPS*), 2020.

697 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *Proceedings of the*
 698 *International Conference on Learning Representations (ICLR)*, 2019.

700 Lin Lu, Chenxi Dai, Wangcheng Tao, Binhang Yuan, Yanan Sun, and Pan Zhou. Position: Exploring
 701 the robustness of pipeline-parallelism-based decentralized training. In *Proceedings of the*
 702 *International Conference on Machine Learning (ICML)*, 2024.

702 Vien Mai and Mikael Johansson. Convergence of a stochastic gradient method with momentum for
 703 non-smooth non-convex optimization. In *Proceedings of the International Conference on Machine*
 704 *Learning (ICML)*, pp. 6630–6639, 2020.

705 Grigory Malinovsky, Peter Richtárik, Samuel Horváth, and Eduard Gorbunov. Byzantine robustness
 706 and partial participation can be achieved at once: Just clip gradient differences. In *Proceedings of*
 707 *the Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2024.

708 Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas.
 709 Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the*
 710 *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pp. 1273–1282, 2017.

711 MetaAI. The Llama 4 herd: The beginning of a new era of natively multimodal AI innovation, 2025.

712 El Mahdi El Mhamdi, Rachid Guerraoui, and Sébastien Rouault. The hidden vulnerability of
 713 distributed learning in byzantium. In *Proceedings of the International Conference on Machine*
 714 *Learning (ICML)*, pp. 3518–3527, 2018.

715 OpenAI. GPT-4 technical report. *CoRR*, abs/2303.08774, 2023.

716 Guansong Pang, Chunhua Shen, Longbing Cao, and Anton Van Den Hengel. Deep learning for
 717 anomaly detection: A review. *ACM computing surveys (CSUR)*, 54(2):1–38, 2021.

718 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor
 719 Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Z. Yang,
 720 Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang,
 721 Junjie Bai, and Soumith Chintala. PyTorch: An imperative style, high-performance deep learning library.
 722 In *Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2019.

723 Guilherme Penedo, Hynek Kydlíček, Loubna Ben allal, Anton Lozhkov, Margaret Mitchell, Colin
 724 Raffel, Leandro Von Werra, and Thomas Wolf. The fineweb datasets: Decanting the web for
 725 the finest text data at scale. In *Proceedings of the Annual Conference on Neural Information*
 726 *Processing Systems Datasets and Benchmarks Track (NeurIPS)*, 2024.

727 B. T. Polyak. Some methods of speeding up the convergence of iteration methods. *USSR Computational*
 728 *Mathematics and Mathematical Physics*, 1964.

729 Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language
 730 models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.

731 Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi
 732 Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text
 733 transformer. *Journal of Machine Learning Research (JMLR)*, 21:140:1–140:67, 2020.

734 Shashank Rajput, Hongyi Wang, Zachary Charles, and Dimitris S. Papailiopoulos. DETOX: A
 735 redundancy-based framework for faster and more robust gradient aggregation. In Hanna M.
 736 Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman
 737 Garnett (eds.), *Proceedings of the Annual Conference on Neural Information Processing Systems*
 738 (*NeurIPS*), pp. 10320–10330, 2019.

739 Sameera Ramasinghe, Thalaiyasingam Ajanthan, Gil Avraham, Yan Zuo, and Alexander Long.
 740 Protocol models: Scaling decentralized training with communication-efficient model parallelism.
 741 In *Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS)*,
 742 2025.

743 Ahmad Rammal, Kaja Gruntkowska, Nikita Fedin, Eduard Gorbunov, and Peter Richtárik. Commu-
 744 nication compression for byzantine robust learning: New efficient algorithms and improved rates.
 745 In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*,
 746 pp. 1207–1215, 2024.

747 Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical*
 748 *statistics*, pp. 400–407, 1951.

756 Max Ryabinin and Anton Gusev. Towards crowdsourced training of large neural networks using
 757 decentralized mixture-of-experts. In *Proceedings of the Annual Conference on Neural Information
 758 Processing Systems (NeurIPS)*, 2020.

759 Max Ryabinin, Eduard Gorbunov, Vsevolod Plokhotnyuk, and Gennady Pekhimenko. Moshpit
 760 SGD: communication-efficient decentralized training on heterogeneous unreliable devices. In
 761 *Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS)*, pp.
 762 18195–18211, 2021.

763 Max Ryabinin, Tim Dettmers, Michael Diskin, and Alexander Borzunov. SWARM parallelism: Train-
 764 ing large models can be surprisingly communication-efficient. In *Proceedings of the International
 765 Conference on Machine Learning (ICML)*, pp. 29416–29440, 2023.

766 Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc V. Le, Geoffrey E. Hinton,
 767 and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer.
 768 In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.

769 Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks
 770 against machine learning models. In *Proceedings of the IEEE Symposium on Security and Privacy
 771 (SP)*, pp. 3–18, 2017.

772 Yuandong Tian, Yiping Wang, Beidi Chen, and Simon S Du. Scan and snap: Understanding training
 773 dynamics and token composition in 1-layer transformer. In *Proceedings of the Annual Conference
 774 on Neural Information Processing Systems (NeurIPS)*, 2023.

775 John W. Tukey. *Exploratory data analysis*. Addison-Wesley series in behavioral science: quantitative
 776 methods. Addison-Wesley, 1977.

777 Stephen M. Walker II. Everything we know about gpt-4, 2023. URL <https://klu.ai/blog/gpt-4-11m>. Accessed on May 2025.

778 An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li,
 779 Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin
 780 Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang,
 781 Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tingyu Xia,
 782 Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu
 783 Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report. *CoRR*, abs/2412.15115, 2024.

784 Binhang Yuan, Yongjun He, Jared Davis, Tianyi Zhang, Tri Dao, Beidi Chen, Percy Liang, Christopher
 785 Ré, and Ce Zhang. Decentralized training of foundation models in heterogeneous environments.
 786 In *Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS)*,
 787 2022.

788 Ruiqi Zhang, Spencer Frei, and Peter L. Bartlett. Trained transformers learn linear models in-context.
 789 *Journal of Machine Learning Research (JMLR)*, 25:49:1–49:55, 2024.

790 Yanli Zhao, Andrew Gu, Rohan Varma, Liang Luo, Chien-Chin Huang, Min Xu, Less Wright, Hamid
 791 Shojanazeri, Myle Ott, Sam Shleifer, Alban Desmaison, Can Balioglu, Pritam Damania, Bernard
 792 Nguyen, Geeta Chauhan, Yuchen Hao, Ajit Mathews, and Shen Li. PyTorch FSDP: experiences
 793 on scaling fully sharded data parallel. *Proceedings of the VLDB Endowment*, 16(12):3848–3860,
 2023.

794 Tongtian Zhu, Wenhao Li, Can Wang, and Fengxiang He. DICE: Data influence cascade in decen-
 795 tralized learning. In *Proceedings of the International Conference on Learning Representations
 796 (ICLR)*, 2025.

797

798

799

800

801

802

803

804

805

806

807

808

809

APPENDIX

The appendix is organized as follows:

- In App. A, we present answers to several common questions, hoping to clarify misconceptions that might arise when interpreting our work and its limitations.
- App. B provides a holistic view of vulnerabilities related to decentralized training that are beyond the scope of the current work.
- We provide an extended version of the related work to SENTINEL in App. C that was omitted from the main paper due to space limitations.
- We present our methodology details in App. D.
- App. E contains our theoretical analysis of SENTINEL, including convergence analysis under relaxed assumptions and conditions under which random worker assignment could result in an honest majority at each stage.
- In App. F.2, we provide detailed experimental results and additional ablations.
- Finally, in App. G, we present a step-by-step integration of SENTINEL with SWARM parallelism (Ryabinin et al., 2023). We provide extensive investigation into seamless integration with trainer nodes that coordinate training in SWARM, followed by real-world experiments training LLMs over 128 untrusted worker nodes employing malicious attacks.

A FREQUENTLY ASKED QUESTIONS (FAQs)

Q0: What are the fundamental differences between data parallel (DP) and pipeline parallel (PP) settings that make existing Byzantine tolerant literature inapplicable to this work? We provide a clear comparison between DP and PP settings, highlighting why existing Byzantine tolerant literature does not apply to our work:

Table 7: Data Parallel vs. Pipeline Parallel Comparison

Aspect	Data Parallel (Prior Work)	Pipeline Parallel (Our Work)
Model Distribution	Full model replica per worker	Model split across workers (layers/stages)
Data Distribution	Different batches per worker	Same batch flows through pipeline
Communication Pattern	Parameter gradients aggregated	Activations/gradients passed sequentially
Byzantine Threat	Corrupted parameter gradients	Corrupted inter-stage activations/gradients
Detection Target	Malicious gradient contributions	Malicious activation/gradient transmissions
Aggregation Method	Robust gradient aggregation	Sequential verification at each stage
Literature Focus	Robust aggregators	No non-trivial prior verification exists

Thus, the key distinctions are:

- **Prior Byzantine tolerant literature:** Secures the DP axis by developing robust aggregation methods for parameter gradients from multiple model replicas
- **Our work:** Secures the orthogonal PP axis by verifying activations and activation gradients transmitted between sequential pipeline stages using SENTINEL.

Hence, existing Byzantine tolerant works do NOT apply to the PP axis because:

1. **No aggregation possible:** In PP, activations from different workers cannot be aggregated (they represent different layers processing different data batch).
2. **Sequential dependency:** Each stage depends on the previous stage’s output, making robust aggregation impossible.
3. **Different threat model:** Malicious workers corrupt intermediate representations rather than final parameter updates.
4. **Verification vs. Aggregation:** Our verifiers monitor communication channels rather than aggregate multiple contributions.

This fundamental difference explains why our threat model and verification approach are necessarily different from “typical Byzantine tolerant” literature.

864 **Q1: Why would a model owner deliver part of their model to an untrusted entity to train?**
 865 The computational resources required for training LLMs are becoming increasingly unsustainable.
 866 As reported by Brown et al. (2020), training a single GPT-3 175B model requires up to 3.6K Petaflop-
 867 days, incurring a total cost of \$4M in AWS pricing. In the absence of big corporations delivering open
 868 source models, decentralized training provides an alternative solution for training such models in a
 869 democratized environment where models can be trained openly and participants are reimbursed based
 870 on their contributions. As discussed in (Yuan et al., 2022), consumer device GPUs are becoming
 871 increasingly available worldwide, many of which are underutilized. “If we could make use of these
 872 devices in a decentralized open-volunteering paradigm for foundation model training, this would be a
 873 **revolutionary alternative** to the expensive solutions offered by data centers.”

874 **Q2: Is assuming 25-37.5% malicious workers realistic? Why would anyone trust such a system?**
 875 We understand the intuitive concern about our 25-37.5% malicious worker percentages. However, note
 876 that these assumptions are standard practice in the Byzantine fault tolerance literature, as evidenced
 877 by recent work in this area outlined in Tab. 8. Importantly, our work serves as a preventative security
 878 mechanism. The goal is to deter malicious behavior by demonstrating robust detection capabilities,
 879 not to operate under the assumption that such high percentages will necessarily occur in practice.
 880 As Fig. 3b demonstrates, when malicious nodes are at lower levels, our performance approaches
 881 nearly 100% detection accuracy. The worst-case analysis ensures system reliability even under
 882 extreme adversarial conditions.

883 Table 8: Malicious Worker Percentages in Byzantine Fault Tolerance Literature

884 Reference	885 Venue	886 Malicious Nodes (%)
886 Mhamdi et al. (2018)	887 ICML 2018	888 47.37
887 Gorbunov et al. (2022)	888 ICML 2022	889 43.75
888 Blanchard et al. (2017)	889 NeurIPS 2017	890 33.00
889 Karimireddy et al. (2021)	890 ICML 2021	891 30.55
890 Malinovsky et al. (2024)	891 NeurIPS 2024	892 25.00
891 Gorbunov et al. (2023)	893 ICLR 2023	894 20.00
892 Karimireddy et al. (2022)	895 ICLR 2022	896 20.00
893 Rammal et al. (2024)	897 AISTATS 2024	898 18.75

899 **Q3: How practical is the integration of verifier nodes in distributed frameworks?** In real-world
 900 systems deploying our algorithm (such as SWARM), trainer nodes are responsible for transmitting
 901 activations/activation gradients between layers. We propose modifying these existing trainer nodes to
 902 perform verification, essentially obtaining this security functionality “for free” since they already have
 903 access to all signals passing between layers. Crucially, trainer nodes represent a centrally controlled
 904 role in distributed systems as they are managed by the network coordinator rather than volunteers,
 905 making it economically practical to maintain control over them. Since their primary responsibility
 906 is coordinating signal transmission, these are lightweight CPU-based nodes with minimal cost. For
 907 example, consider training a 4B parameter LLM where each layer/stage requires roughly 18GB
 908 of GPU VRAM. On AWS, each worker would require a g5.2xlarge EC2 instance with 24GB
 909 VRAM at approximately \$1.212 per hour. In contrast, trainer nodes can be bundled with 8 nodes
 910 per c5a.8xlarge instance (32vCPUs, 4vCPU per trainer) at \$1.232 per hour (i.e., \$0.154 per
 911 trainer node per hour). This represents a significant cost reduction compared to worker instances, and
 912 since trainers are centrally controlled, it is both economically and operationally feasible for network
 913 coordinators to bear these costs while maintaining security and reliability. If trainer roles must be
 914 delegated to volunteers, proper authentication mechanisms would be required to prevent malicious
 915 behavior, which we leave to future work. For a more in-depth discussion on implementing SENTINEL
 916 in SWARM, please see App. G.

917 **Q4: Why do you assume a “warm-up” period with only honest workers? Does this contradict the**
 918 **spirit of distributed training?** In distributed environments, it is common practice to start training
 919 in a controlled environment until training reaches a stabilized point before allowing public workers to
 920 join. This approach is important not only for our verification method, but also ensures that training
 921 is stable before public participation begins. This is NOT against the spirit of distributed training:
 922 consider a 40-layer LLM intended for decentralized training. Initially, we deploy a single replica
 923 across 40 nodes under our control, but training throughput is limited. After the initial warm-up phase,

918 we replicate the pipeline across 7 additional replicas, bringing our total DP workers to 320, where
 919 only 12.5% are controlled by the model owner and 87.5% are public workers. Note that maintaining
 920 at least one trusted worker at each stage is crucial not only for initial warm-up but also due to system
 921 reliability as the volunteer nodes (280 of them in our example) might drop mid-run. We need at least
 922 one reliable pipeline to ensure training continuity when volunteers drop.
 923

924 **Q5: Does the N -to-1 communication to verifiers create a bandwidth bottleneck?** There is no
 925 bandwidth bottleneck. In model parallel implementations such as SWARM, coordination between
 926 stages occurs through trainer nodes that send and receive activations/gradients. Since each trainer
 927 node already observes all signals passing from one stage to another, we leverage it as our trusted
 928 verifier node. As we discuss in detail in App. G, SWARM utilizes multiple trainer nodes to streamline
 929 the data through the stages one-by-one. Thus, each trainer can run their own verification using
 930 SENTINEL simultaneously to other trainers sending signals to the workers. For technical details on
 931 trainer node operations, see the Hivemind (Ryabinin & Gusev, 2020) library.
 932

933 **Q6: How does the verification mechanism handle false positives?** The violation counter with a
 934 forgiveness mechanism is a clever way to handle transient anomalies and avoid unfairly banning honest
 935 workers. The training curves are given in Figs. 8 and 16 in App. F. As seen, the training/validation
 936 curves show no visible impact on training dynamics from transiently replacing submitted signals
 937 with EMA values. This makes intuitive sense: when only a few nodes undergo the EMA replacement
 938 phase, the remaining workers in the DP setup continue to submit useful signals that guide training
 939 effectively.
 940

941 **Q7: Why does the effectiveness of training-interruption attacks vary across different attack
 942 versions and targets?** When comparing activation-based attacks against gradient-based attacks,
 943 two key factors explain the effectiveness differences:
 944

- **Magnitude differences:** Activation values are orders of magnitude larger than activation gradients. Therefore, methods that modify activations can achieve larger perturbation magnitudes during attacks, resulting in more successful disruption.
- **Propagation scope:** Activation manipulation at layer $1 \leq \ell \leq L$ affects the forward pass for all subsequent layers ($[\ell + 1, \ell + 2, \dots, L]$) and the backward pass of all layers. In contrast, gradient manipulation at layer ℓ only affects the gradients of preceding layers $[1, 2, \dots, \ell - 1]$. Consequently, if an attacker is positioned in the middle of the network, manipulating activations has broader impact on the entire training process.

952 **Q8: How does your convergence rate compare to well-known lower bounds from Byzantine-
 953 tolerant literature?** Our convergence guarantee provides an accurate bound given our assumptions.
 954 We note a key distinction in our setting: prior Byzantine-tolerant literature considers data parallel
 955 training where malicious actors modify “parameter gradients”. Our work addresses the *orthogonal*
 956 pipeline parallel axis where activations and gradients “between layers” are shared and require
 957 verification. These two axes are complementary: securing both pipeline parallel and data parallel
 958 axes is important in decentralized settings, and this work focuses on the former. Therefore, prior
 959 guarantees from Byzantine-tolerant literature are not directly comparable.
 960

961 **Q9: How does your method compare to the prior work by Lu et al. (2024)?** Lu et al. (2024) used
 962 a naïve approach of assigning one duplicate replica per worker for verification. For instance, with
 963 320 worker nodes this requires splitting them into two groups: 160 workers performing computation
 964 and 160 workers replicating their work for verification. While this approach achieves 100% F1-score
 965 on all attacks, it operates at HALF the true distributed network throughput. Our solution does NOT
 966 replicate volunteer node work, achieving twice the training speed. Conceptually, on OpenWebText
 967 against mixed attacks we will have:
 968

969 **Q10: Does removing nodes from the training runs degrade the throughput?** Note that in
 970 real-world eco-systems such as SWARM (Ryabinin et al., 2023), we are not using a fixed mesh
 971 anymore. Workers join and leave the SWARM as they wish, and within each pipeline stage, there are
 972

972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
Table 9: Comparison with prior work on OpenWebText mixed attacks

METHOD	METRICS			
	PR. (%) ↑	RE. (%) ↑	F1 (%) ↑	TPS ↑
DUPLICATE WORK (Lu et al., 2024)	100.0	100.0	100.0	6483
SENTINEL (OURS)	91.9	87.2	89.5	12966

numerous workers that are serving that stage. This combined with stochastic routing in SWARM (please see the details in App. G.3) ensures that workers would reach their maximum utilization. Thus, in real-world systems where we have an abundance of workers serving each stage where kicking workers out would not necessarily degrade throughput.

B VULNERABILITIES OF DECENTRALIZED TRAINING USING DATA AND PIPELINE PARALLELISM

In decentralized settings used for collaborative training, the verification of participant workers is essential to maintain the integrity, security, and overall effectiveness of the training process (Lu et al., 2024). Verification acts as a critical quality-control measure, ensuring that each participant meaningfully contributes to the collective training effort. Without adequate verification mechanisms, malicious actors can infiltrate the SWARM, potentially compromising its integrity. Such attackers might disrupt the collaborative training, degrade its efficiency, or illegitimately benefit by accessing the trained model without genuinely contributing.

Below, we categorize common malicious behaviors that could arise in decentralized collaborative training scenarios. These categories are not mutually exclusive, as attackers may employ several tactics simultaneously. Nevertheless, this classification provides a structured overview of the key threats:

- **Training Disruption (Denial-of-Service or DoS):** Attackers intentionally impede or halt the training process. This can occur through dropping essential updates, introducing malicious data designed to break communication protocols, or overwhelming the system with excessive or irrelevant submissions.
- **Free-Riding or Minimal Effort Contributions:** Participants contribute minimal computational effort or data yet aim to reap the benefits of the collective process, such as accessing the final model, receiving rewards, or boosting their reputation (Zhu et al., 2025). Common tactics include submitting trivial updates or strategically remaining inactive until training nears completion.
- **Model Poisoning and Backdoor Attacks:** Malicious actors provide adversarial updates designed to introduce subtle vulnerabilities or targeted misbehaviors in the resulting model (Li et al., 2024). Typically concealed under normal operational conditions, these backdoors or compromised models trigger malicious outcomes only under specific, pre-defined scenarios.
- **Privacy Violations (Data Extraction or Inference Attacks):** Attackers exploit gradients, activations, or other shared information during training to infer sensitive or private information from other participants’ datasets, thereby breaching confidentiality and compromising user privacy (Shokri et al., 2017).
- **Reputation or Credit Manipulation:** Participants deliberately falsify or exaggerate their contributions (for instance, by generating seemingly high-quality updates) to unjustly obtain greater rewards, enhanced reputation, or tokens. This form of manipulation undermines the fairness of the system and distorts trust among honest contributors.

In this paper, our primary focus is mitigating threats associated with training disruption. Ensuring the identification and exclusion of malicious participants who submit harmful or disruptive updates is critical. Failure to address these threats effectively would prevent the swarm from achieving model convergence and producing a reliable, functional final model.

1026 **C EXTENDED RELATED WORK**
10271028 Secure distributed training has gained significant attention with the proliferation of decentralized
1029 machine learning systems. Our work builds upon several research threads while addressing unique
1030 challenges posed by pipeline parallelism in LLM training.
10311032 **Decentralized Training Frameworks.** Decentralized training has emerged as a promising approach
1033 for democratizing AI capabilities. Ryabinin and Gusev (Ryabinin & Gusev, 2020) introduced the
1034 Hivemind framework, enabling mixture-of-experts models to be trained in a decentralized fashion.
1035 Building on this foundation, (Ryabinin et al., 2021) proposed Moshpit SGD, a communication-
1036 efficient algorithm for training on heterogeneous and unreliable devices. In parallel, (Yuan et al.,
1037 2022) introduced Tasklets, a system for decentralized training in heterogeneous environments that
1038 adapts to varying network conditions and compute capabilities. SWARM parallelism (Ryabinin
1039 et al., 2023) further enhanced this approach by combining pipeline and data parallelism to enable
1040 training of models significantly larger than those possible with previous decentralized methods.
1041 While these frameworks prioritize fault tolerance against non-malicious failures, they generally lack
1042 protection against adversarial participants which is a critical vulnerability in open decentralized
1043 training environments.
10441045 **Byzantine-Resilient Distributed Training.** Byzantine fault tolerance in distributed learning has
1046 been extensively studied in the context of federated learning (McMahan et al., 2017) and data-parallel
1047 training (Li et al., 2020). (Blanchard et al., 2017) introduced Krum, the first Byzantine-tolerant
1048 aggregation rule for distributed SGD that could withstand arbitrary gradient manipulations from com-
1049 promised workers. This was followed by more sophisticated approaches including Bulyan (Mhamdi
1050 et al., 2018), median-based aggregation (Baruch et al., 2019), and clipping-based methods (He et al.,
2020; Malinovsky et al., 2024).1051 A fundamentally different approach called CENTEREDCLIP was proposed by (Karimireddy et al.,
1052 2021), who leveraged historical gradient information to detect anomalous updates – conceptually
1053 similar to our momentum-based verification but applied specifically to gradient aggregation. Recent
1054 work by Rammal et al. (Rammal et al., 2024) demonstrated that communication compression could
1055 be effectively combined with Byzantine-robust learning, achieving improved convergence rates while
1056 maintaining security guarantees.
10571058 While these methods provide strong theoretical guarantees, they primarily target scenarios where
1059 workers compute complete gradients independently, making them ill-suited for pipeline parallel
1060 configurations like SWARM where intermediate activations are communicated between stages.
1061 Furthermore, these approaches often involve comparing gradients across workers, which would
1062 necessitate parameter replication across stages, contradicting pipeline parallel’s objective of enabling
1063 training of models too large to fit on a single device.
10641065 **Security in Pipeline Parallel Architectures.** Security considerations specific to pipeline parallel
1066 training have received limited attention compared to other distributed paradigms. (Lu et al., 2024)
1067 recently presented a position paper exploring robustness challenges in pipeline parallelism-based
1068 decentralized training, highlighting activation-based attacks as a critical concern. Their work, however,
1069 focused primarily on identifying vulnerabilities rather than proposing comprehensive verification
1070 solutions. The redundancy-based approach proposed by (Lu et al., 2024) and (Rajput et al., 2019)
1071 could, in principle, be adapted to decentralized pipeline parallel settings. However, these approaches
1072 would introduce significant computational overhead (due to duplicating computations across workers)
1073 which would greatly diminish the scalability benefits of decentralized training.
10741075 **D DETAILS OF MOMENTUM-BASED VERIFICATION**
10761077 **D.1 MOMENTUM-BASED VERIFICATION ALGORITHMS**
10781079 In this section, we present our detailed algorithm for worker verification in decentralized training.
Alg. 1 outlines our end-to-end verification mechanism for this setting. We present the algorithm
chronologically as training progresses. The verifier nodes perform all verification operations, while
worker nodes are solely responsible for computing activations during the forward pass and their
1080

1080
 1081
 1082
 1083
 1084
 1085
 1086
 1087
 1088
 1089
 1090
 1091
 1092
 1093
 1094
 1095
 1096
 1097
 1098
 1099
 1100
 1101
 1102
 1103
 1104
 1105
 1106
 1107
 1108
 1109
 1110
 1111
 1112
 1113
 1114
 1115
 1116
 1117
 1118
 1119
 1120
 1121
 1122
 1123
 1124
 1125
 1126
 1127
 1128
 1129
 1130
 1131
 1132
 1133

respective gradients during the backward pass. Algorithms 2 and 3 detail the verification procedures for both forward and backward passes, respectively. Alg. 4 presents our approach for mitigating cascading effects as described in Section 3.1. Finally, Alg. 5 specifies our adaptive IQR threshold setting methodology for each metric in our approach.

Algorithm 1 Momentum-based Verification for SWARM Parallelism

Require: Parameters $\beta_h, \beta_g \in (0, 1)$, violation threshold c , forgiveness period $T_{\text{forgiveness}}$, set of metrics \mathcal{M}

1: Initialize $\mathbf{m}_0^{(s)}(\mathbf{h}) = \mathbf{0}$, $\mathbf{m}_0^{(s)}(\mathbf{g}) = \mathbf{0}$, $v_r^{(s)} = 0$, $B_s = \emptyset$, $\mathcal{H}_l^{(s)} = \emptyset$, $\mathcal{G}_l^{(s)} = \emptyset$ for all s, r
 2: Initialize $\mathcal{T} = \emptyset$ // Initialize global tainted set

3: // Warm-up phase to establish baseline statistics

4: **for** $t = 1$ to T_{warmup} **do**

5: **for** $s \in \{1, 2, \dots, p\}$ **do**

6: Collect $\mathbf{h}_t^{(s,r)}$ and $\mathbf{g}_t^{(s,r)}$ from all workers $r \in \{1, \dots, d\}$

7: Compute $\Gamma_t^{(s,r)} = \Omega(\mathbf{h}_t^{(s,r)}, \mathbf{m}_{t-1}^{(s)}(\mathbf{h})) \forall r$, add to $\mathcal{H}_l^{(s)}$

8: Compute $\Gamma_t^{(s,r)} = \Omega(\mathbf{g}_t^{(s,r)}, \mathbf{m}_{t-1}^{(s)}(\mathbf{g})) \forall r$, add to $\mathcal{G}_l^{(s)}$

9: Update momentum $\mathbf{m}_t^{(s)}(\mathbf{h})$ and $\mathbf{m}_t^{(s)}(\mathbf{g})$ using Eq. (1)

10: **end for**

11: **end for**

12: // Main training phase with verification

13: **for** $t = T_{\text{warmup}} + 1$ to T_{total} **do**

14: $\mathcal{T}_t = \emptyset$ // Initialize tainted set for current iteration

15: // Step 1: Forward Pass and Activation Verification

16: **for** $s \in \{1, \dots, p\}$ **do**

17: $\mathcal{T}_t^{(s)} \leftarrow \text{ACTIVATIONVERIFICATION}(s, t, \mathbf{m}_{t-1}^{(s)}(\mathbf{h}), \mathcal{H}_l^{(s)}, B_s, v_r^{(s)}, \mathcal{M}, c, T_{\text{forgiveness}}, \mathcal{T}_t)$

18: $\mathcal{T}_t \leftarrow \mathcal{T}_t \cup \mathcal{T}_t^{(s)}$ // Accumulate tainted workers

19: $R_{\text{clean}} = \{r : (t, s, r) \notin \mathcal{T}_t\}$

20: $\mathbf{m}_t^{(s)}(\mathbf{h}) = \beta_h \mathbf{m}_{t-1}^{(s)}(\mathbf{h}) + (1 - \beta_h) \frac{1}{|R_{\text{clean}}|} \sum_{r \in R_{\text{clean}}} \mathbf{h}_t^{(s,r)}$

21: **end for**

22: **for** $s \in \{p, p-1, \dots, 1\}$ **do**

23: // Step 2: Backward Pass and Gradient Verification

24: $\mathcal{T}_t^{(s)} \leftarrow \text{GRADIENTVERIFICATION}(s, t, \mathbf{m}_{t-1}^{(s)}(\mathbf{g}), \mathcal{G}_l^{(s)}, B_s, v_r^{(s)}, \mathcal{M}, c, T_{\text{forgiveness}}, \mathcal{T}_t)$

25: $\mathcal{T}_t \leftarrow \mathcal{T}_t \cup \mathcal{T}_t^{(s)}$ // Accumulate tainted workers

26: // Step 3: Gradient Replacement for Tainted Workers

27: $\mathbf{g}_t^{(s)} \leftarrow \text{GRADIENTREPLACEMENT}(s, t, \mathbf{m}_{t-1}^{(s)}(\mathbf{g}), \mathbf{g}_t^{(s)}, \mathcal{T}_t)$

28: $R_{\text{clean}} = \{r : (t, s, r) \notin \mathcal{T}_t\}$

29: $\mathbf{m}_t^{(s)}(\mathbf{g}) = \beta_g \mathbf{m}_{t-1}^{(s)}(\mathbf{g}) + (1 - \beta_g) \frac{1}{|R_{\text{clean}}|} \sum_{r \in R_{\text{clean}}} \mathbf{g}_t^{(s,r)}$

30: **end for**

31: $\mathcal{T} \leftarrow \mathcal{T} \cup \mathcal{T}_t$ // Accumulate tainted entries across iterations

32: **end for**

33: **return** \mathcal{T} // Return the complete set of tainted worker-stage-iteration tuples

1134
 1135
 1136
 1137
 1138
 1139
 1140
 1141
 1142
 1143
 1144

Algorithm 2 ACTIVATIONVERIFICATION

1145
 1146 **Require:** $s, t, \mathbf{m}_{t-1}^{(s)}(\mathbf{h}), \mathcal{H}_l^{(s)}, B_s, v_r^{(s)}, \mathcal{M}, c, T_{\text{forgiveness}}, \mathcal{T}_t$
 1147 1: $\mathcal{T}_t^{(s)} = \emptyset$ // Initialize tainted set for current stage
 1148 2: Truncate $\mathcal{H}_l^{(s)} \leftarrow \mathcal{H}_l^{(s)}[-l : \text{end}]$
 1149 3: Collect $\mathbf{h}_t^{(s,r)}$ from all $r \in \{1, \dots, d\} \setminus B_s$
 1150 4: **for** $r \in \{1, \dots, d\} \setminus B_s$ not in \mathcal{T}_t **do**
 1151 5: Compute metrics: $\Gamma_t^{(s,r,i)} = \Omega_i(\mathbf{h}_t^{(s,r)}, \mathbf{m}_{t-1}^{(s)}(\mathbf{h}))$ for $i \in \mathcal{M}$
 1152 6: **if** $\exists i \in \mathcal{M} : |\Gamma_t^{(s,r,i)} - q_2^{(s,i)}| \geq k_{\text{tukey}}^{(s,i)}(q_3^{(s,i)} - q_1^{(s,i)})$ **then**
 1153 7: $v_r^{(s)} \leftarrow v_r^{(s)} + 1$ // Increment violation counter
 1154 8: $\mathcal{T}_t^{(s)} \leftarrow \mathcal{T}_t^{(s)} \cup \{(t, s, r)\}$ // Mark as tainted in current stage
 1155 9: **if** $v_r^{(s)} \geq c$ or $\Gamma_t^{(s,r,i)} \gg k_{\text{tukey}}^{(s,i)}(q_3^{(s,i)} - q_1^{(s,i)})$ **then**
 1156 10: $B_s \leftarrow B_s \cup \{r\}$ // Ban worker
 1157 11: Notify stages $s' > s$ to flag affected mini-batches
 1158 12: **end if**
 1159 13: **else**
 1160 14: Add $\Gamma_t^{(s,r,i)}$ to $\mathcal{H}_l^{(s,i)}$ $\forall i \in \mathcal{M}$
 1161 15: $v_r^{(s)} \leftarrow \max(0, v_r^{(s)} - 1)$ if $T_{\text{forgiveness}}$ consecutive clean steps
 1162 16: **end if**
 1163 17: **end for**
 1164 18: Update IQR statistics and adjust $k_{\text{tukey}}^{(s,i)}$ $\forall i \in \mathcal{M}$ using Alg. 5
 1165 19: **if** $|\mathcal{T}_t^{(s)}| > 0.5 \cdot (d - |B_s|)$ **then**
 1166 20: $\mathcal{T}_t^{(s)} \leftarrow \emptyset$ // Clear if more than 50% flagged (natural shift, see App. D.3)
 1167 21: **end if**
 1168 22: **return** $\mathcal{T}_t^{(s)}$ // Return tainted workers for this stage

1169
 1170
 1171
 1172
 1173
 1174
 1175
 1176
 1177
 1178
 1179
 1180
 1181
 1182
 1183
 1184
 1185
 1186
 1187

1188

1189

Algorithm 3 GRADIENT VERIFICATION

1190

1191

Require: $s, t, \mathbf{m}_{t-1}^{(s)}(\mathbf{g}), \mathcal{G}_l^{(s)}, B_s, v_r^{(s)}, \mathcal{M}, c, T_{\text{forgiveness}}, \mathcal{T}_t$

1192

1: $\mathcal{T}_t^{(s)} = \emptyset$ // Initialize tainted set for current stage

1193

2: Truncate $\mathcal{G}_l^{(s)} \leftarrow \mathcal{G}_l^{(s)}[-l : \text{end}]$

1194

3: $\text{Tainted}_{\text{downstream}} = \{r : (t, s', r) \in \mathcal{T}_t \text{ for some } s' > s\}$ // Workers tainted in downstream stages

1195

4: **for** $r \in \{1, \dots, d\} \setminus B_s$ **do**

1196

5: **if** $r \in \text{Tainted}_{\text{downstream}}$ **then**

1197

6: $\mathcal{T}_t^{(s)} \leftarrow \mathcal{T}_t^{(s)} \cup \{(t, s, r)\}$ // Mark as tainted by downstream

1198

7: **else**

1199

8: Collect $\mathbf{g}_t^{(s,r)}$ from worker r

1200

9: Compute metrics: $\Gamma_t^{(s,r,i)} = \Omega_i(\mathbf{g}_t^{(s,r)}, \mathbf{m}_{t-1}^{(s)}(\mathbf{g}))$ for $i \in \mathcal{M}$

1201

10: **if** $\exists i \in \mathcal{M} : |\Gamma_t^{(s,r,i)} - q_2^{(s,i)}| \geq k_{\text{tukey}}^{(s,i)}(q_3^{(s,i)} - q_1^{(s,i)})$ **then**

1202

11: $v_r^{(s)} \leftarrow v_r^{(s)} + 1$ // Increment violation counter

1203

12: $\mathcal{T}_t^{(s)} \leftarrow \mathcal{T}_t^{(s)} \cup \{(t, s, r)\}$ // Mark as tainted

1204

13: **if** $v_r^{(s)} \geq c$ or $\Gamma_t^{(s,r,i)} \gg k_{\text{tukey}}^{(s,i)}(q_3^{(s,i)} - q_1^{(s,i)})$ **then**

1205

14: $B_s \leftarrow B_s \cup \{r\}$ // Ban worker

1206

15: Notify stages $s' < s$ to flag affected mini-batches

1207

16: **end if**

1208

17: **else**

1209

18: Add $\Gamma_t^{(s,r,i)}$ to $\mathcal{G}_l^{(s,i)}$ $\forall i \in \mathcal{M}$

1210

19: $v_r^{(s)} \leftarrow \max(0, v_r^{(s)} - 1)$ if $T_{\text{forgiveness}}$ consecutive clean steps

1211

20: **end if**

1212

21: **end if**

1213

22: **end for**

1214

23: Update IQR statistics and adjust $k_{\text{tukey}}^{(s,i)}$ $\forall i \in \mathcal{M}$ using Alg. 5

1215

24: **if** $|\mathcal{T}_t^{(s)}| > 0.5 \cdot (d - |B_s|)$ **then**

1216

25: $\mathcal{T}_t^{(s)} \leftarrow \emptyset$ // Clear if more than 50% flagged (natural shift, see App. D.3)

1217

26: **end if**

1218

27: **return** $\mathcal{T}_t^{(s)}$ // Return tainted workers for this stage

1219

1220

1221

Algorithm 4 GRADIENT REPLACEMENT

1222

1223

Require: $s, t, \mathbf{m}_{t-1}^{(s)}(\mathbf{g}), \mathbf{g}_t^{(s)}, \mathcal{T}_t$

1224

1: **for** $r \in \{1, \dots, d\}$ **do**

1225

2: **if** $(t, s, r) \in \mathcal{T}_t$ **then**

1226

3: $\mathbf{g}_t^{(s,r)} \leftarrow \mathbf{m}_{t-1}^{(s)}(\mathbf{g})$ // Replace gradient with momentum

1227

4: **end if**

1228

5: **end for**

1229

6: **return** $\mathbf{g}_t^{(s)}$ // Return updated gradients

1230

1231

1232

1242
1243
1244
1245
1246

Algorithm 5 Adaptive IQR Threshold Adjustment

1247 **Require:** History window $\mathcal{H}_l^{(s,i)}$ (or $\mathcal{G}_l^{(s,i)}$ for gradients) for stage s and metric $i \in \mathcal{M}$, initial
1248 multiplier k_0 , target false positive rate α , growth factor $\gamma_g > 1$, shrink factor $\gamma_s < 1$, maximum
1249 iterations N_{\max} , minimum distance multipliers Λ
1250
1251
1252
1253 1: *// Calculate initial statistics*
1254 2: $q_1, q_2, q_3 \leftarrow$ 25th, 50th, 75th percentiles of $\mathcal{H}_l^{(s,i)}$
1255 3: $\text{IQR} \leftarrow \max(q_3 - q_1, \epsilon)$ *// Ensure non-zero IQR with small ϵ*
1256 4: $k \leftarrow k_0$ *// Initialize with previous multiplier value*
1257 5: $\tau_{\text{lower}} \leftarrow q_2 - k \cdot \text{IQR}$
1258 6: $\tau_{\text{upper}} \leftarrow q_2 + k \cdot \text{IQR}$
1259 7: $\text{FP-rate} \leftarrow$ fraction of $\mathcal{H}_l^{(s,i)}$ outside $[\tau_{\text{lower}}, \tau_{\text{upper}}]$
1260
1261
1262 8: *// Widen thresholds if false positive rate too high*
1263 9: $\text{iter} \leftarrow 0$
1264 10: **while** $\text{FP-rate} > \alpha$ and $\text{iter} < N_{\max}$ **do**
1265 11: $k \leftarrow k \cdot \gamma_g$ *// Grow multiplier*
1266 12: $\tau_{\text{lower}} \leftarrow q_2 - k \cdot \text{IQR}$
1267 13: $\tau_{\text{upper}} \leftarrow q_2 + k \cdot \text{IQR}$
1268 14: $\text{FP-rate} \leftarrow$ fraction of $\mathcal{H}_k^{(s,i)}$ outside $[\tau_{\text{lower}}, \tau_{\text{upper}}]$
1269 15: $\text{iter} \leftarrow \text{iter} + 1$
1270
1271 16: **end while**
1272
1273
1274 17: *// Narrow thresholds if false positive rate too low*
1275 18: $\text{iter} \leftarrow 0$
1276 19: **while** $\text{FP-rate} \ll \alpha$ and $\text{iter} < N_{\max}$ **do**
1277 20: $k \leftarrow k \cdot \gamma_s$ *// Shrink multiplier*
1278 21: $\tau_{\text{lower}} \leftarrow q_2 - k \cdot \text{IQR}$
1279 22: $\tau_{\text{upper}} \leftarrow q_2 + k \cdot \text{IQR}$
1280 23: $\text{FP-rate} \leftarrow$ fraction of $\mathcal{H}_l^{(s,i)}$ outside $[\tau_{\text{lower}}, \tau_{\text{upper}}]$
1281 24: $\text{iter} \leftarrow \text{iter} + 1$
1282
1283 25: **end while**
1284
1285 26: *// Enforce minimum distance from median based on metric type (optional)*
1286 27: $\lambda \leftarrow \Lambda[i]$ *// Get multiplier for current metric*
1287 28: $d_{\min} \leftarrow |q_2| \cdot \lambda$ *// Minimum threshold distance*
1288 29: $\tau_{\text{lower}} \leftarrow \min(\tau_{\text{lower}}, q_2 - d_{\min})$
1289 30: $\tau_{\text{upper}} \leftarrow \max(\tau_{\text{upper}}, q_2 + d_{\min})$
1290 31: **return** $\tau_{\text{lower}}, \tau_{\text{upper}}, k$
1291
1292
1293
1294
1295

1296
1297

D.2 ON HANDLING THE CASCADING EFFECT

1298
1299
1300
1301
1302
1303
1304

Pipeline parallelism exhibits distinct architectural characteristics compared to traditional federated learning (McMahan et al., 2017) approaches. One key challenge is what we term the “cascading effect” which occurs exclusively in pipeline parallelism. During forward propagation, a single node submitting malicious activations can contaminate all subsequent activations, potentially causing downstream verifier nodes to incorrectly flag benign nodes as malicious (see Fig. 5a). This phenomenon occurs similarly during backward propagation as depicted in Fig. 6b. The cascading effect could significantly increase false positive detection rates, making it critical to address this challenge.

1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315

To mitigate this issue, our approach (described in Sec. 3.1) implements inter-node communication protocols. Specifically, verifier nodes maintain a “tainted” list tracking upstream nodes identified as potentially malicious which they communicate with subsequent verifiers to prevent them from updating their EMAs and falsely flagging nodes affected by an attacker downstream. During backward propagation, all nodes sharing the same data parallel rank as the compromised node receive gradient momentum instead of actual gradients. Throughout this verification process, worker nodes continue processing data at a consistent pace, ensuring no node detects unusual behavioral patterns in the network.

1316
1317
1318
1319

The cascading effect manifests in both propagation directions, as malicious behavior can target either activation or gradient signals. We address this bidirectional vulnerability with two corresponding mitigation strategies:

1. When activations are compromised, all affected nodes receive gradient momentum (see Fig. 5);
2. When malicious behavior occurs during backward propagation, all downstream nodes switch to activation gradient (see Fig. 6).

1320
1321
1322
1323

While an alternative approach could involve sending zero vectors as gradients, this would effectively stall training in the affected pipe.² We leave exploration of appropriate gradient signals for the tainted segment to future work.

1324
1325

D.3 ON NATURAL DISTRIBUTION SHIFT

1326
1327
1328
1329

Beyond malicious attacks, legitimate distribution shifts can occur naturally during training due to evolving data characteristics or model dynamics (Tian et al., 2023; Zhang et al., 2024). In such cases, multiple worker nodes at the same pipeline stage may simultaneously exhibit statistical deviations that would normally trigger malicious detection, despite all nodes behaving honestly.

1330
1331
1332
1333
1334
1335

To distinguish between natural distribution shifts and coordinated attacks, we implement a consensus-based approach at the verifier level. When more than 50% of nodes at a given pipeline stage are flagged as potentially malicious, the verifier attributes this to a natural distribution shift rather than malicious behavior. This threshold leverages the *honest majority assumption*: coordinated attacks involving more than half the nodes would violate our security model, making such scenarios indistinguishable from legitimate system-wide changes.

1336
1337
1338
1339
1340
1341
1342
1343
1344
1345

Upon detecting a natural distribution shift, the system responds as follows:

1. Training continues normally without malicious mitigation protocols.
2. The cascading effect mechanism described in App. D.2 is not activated.
3. Nodes update their EMA statistics to adapt to the new data distribution.

1346
1347
1348
1349

This consensus mechanism ensures that legitimate distributional changes do not trigger unnecessary verification overhead or training disruptions. However, refining this approach for asynchronous training environments (Ajanthan et al., 2025), where nodes may experience distribution shifts at different times, remains an important direction for future work.

²We will demonstrate that in our SWARM implementation in App. G, this choice does not have an impact on convergence as SWARM utilizes a stochastic routing.

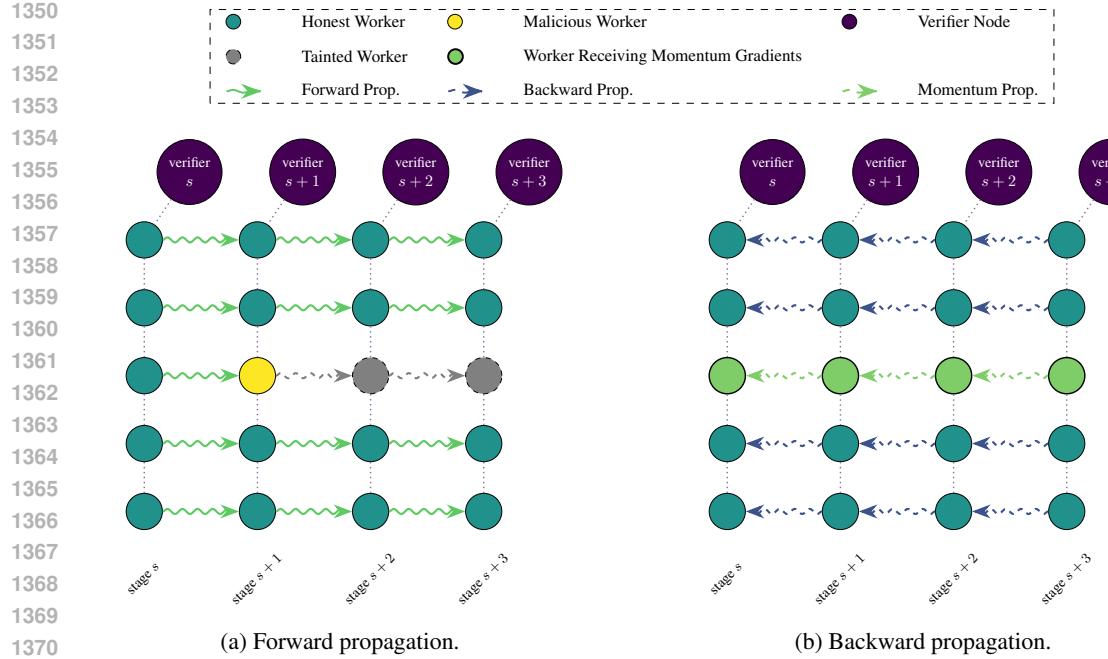


Figure 5: Verification protocol for handling compromised workers during distributed training. During forward propagation, a worker at stage $s + 1$ is detected as potentially compromised (shown in yellow). The verifier nodes continue forwarding activations to subsequent stages without alerting downstream workers to avoid disrupting the pipeline. During backward propagation, instead of propagating gradients computed by the compromised worker, verifier nodes substitute gradient momentum values to maintain training stability. Communication flows through verifier nodes between consecutive pipeline stages, though direct worker-to-worker arrows are shown for visual clarity.

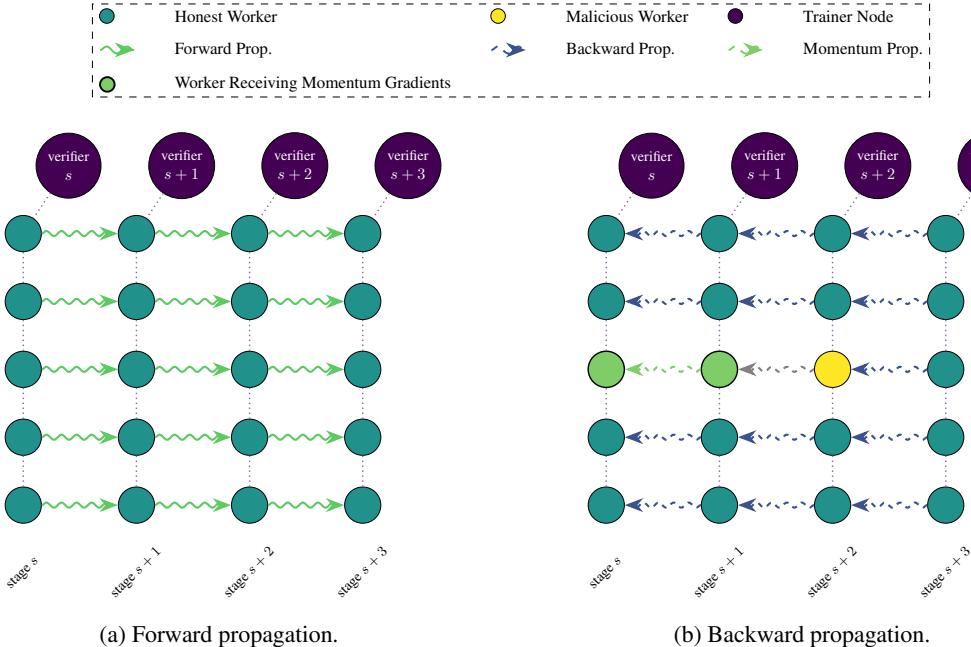


Figure 6: Verification protocol for handling compromised workers during gradient propagation. During backward propagation, a worker at stage $s + 2$ is detected as potentially compromised (shown in yellow). To prevent propagation of tainted gradients, verifier nodes substitute gradient momentum values for all workers in preceding stages ($s + 1, s, \dots$) instead of forwarding the corrupted gradients. This ensures training stability while maintaining the pipeline flow without alerting downstream workers to the compromise.

1404 **E THEORETICAL GUARANTEES AND THEIR PROOFS**
 1405

1406 **E.1 CONVERGENCE ANALYSIS**
 1407

1408 In this section, we present our full convergence analysis. Note that the bounds that we derive are in
 1409 no means the tightest possible bounds. Instead, our aim is to establish a mathematical connection
 1410 between our momentum-based verification and training dynamics.

1411 **E.1.1 MALICIOUS DETECTION BOUNDS**
 1412

1413 Our first goal is to establish bounds on the maximum perturbation a malicious worker can introduce
 1414 without being detected. We begin by analyzing how momentum smoothing affects the global deviation
 1415 in activation vectors.³

1416 **Lemma 2** (Momentum Smoothing Bounds the Global Deviation). *Let the activation vector
 1417 momentum at stage s and iteration t be updated by*

$$1418 \mathbf{m}_t^{(s)} = \beta_h \mathbf{m}_{t-1}^{(s)} + (1 - \beta_h) \left(\frac{1}{d} \sum_{r=1}^d \mathbf{h}_t^{(s,r)} \right), \quad 0 \leq \beta < 1, \quad (4)$$

1419 where d represents the number of worker replicas at each stage. Assume:

- 1420 1. A fraction $\gamma_s < \frac{1}{2}$ of the workers are malicious, with $B_s \subset \{1, 2, \dots, d\}$ denoting the
 1421 subset of malicious workers.
- 1422 2. A malicious worker adds a vector perturbation $\delta_t^{(s,r)}$ satisfying

$$1423 \|\delta_t^{(s,r)}\| \leq \varepsilon. \quad (5)$$

1424 Then the deviation in the momentum caused by the malicious perturbations obeys

$$1425 \|\Delta \mathbf{m}_t^{(s)}\| \leq \gamma_s \varepsilon, \quad (6)$$

1426 where $\Delta \mathbf{m}_t^{(s)}$ is the difference between the momentum computed with the malicious perturbations
 1427 and the momentum computed using only the unperturbed (honest) activations.

1428 *Proof.* Let $H_s = \{1, 2, \dots, d\} \setminus B_s$ denote the set of honest workers at stage s . Since at most a
 1429 fraction γ_s of workers are malicious:

$$1430 |B_s| = \gamma_s d, \quad |H_s| = (1 - \gamma_s) d. \quad (7)$$

1431 We can express each worker's activation vector as

$$1432 \mathbf{h}_t^{(s,r)} = \mathbf{h}_{t,\text{nom}}^{(s,r)} + \mathbf{e}_t^{(s,r)}, \quad \text{where} \quad \mathbf{e}_t^{(s,r)} = \begin{cases} \delta_t^{(s,r)}, & r \in B_s, \\ \mathbf{0}, & r \in H_s. \end{cases} \quad (8)$$

1433 The nominal (unperturbed) average activation can be written as:

$$1434 \bar{\mathbf{h}}_t^{(s)} = \frac{1}{d} \sum_{r=1}^d \mathbf{h}_{t,\text{nom}}^{(s,r)}. \quad (9)$$

1435 For the observed (perturbed) average, we have:

$$1436 \hat{\mathbf{h}}_t^{(s)} = \frac{1}{d} \sum_{r=1}^d (\mathbf{h}_{t,\text{nom}}^{(s,r)} + \mathbf{e}_t^{(s,r)}) \\ 1437 = \bar{\mathbf{h}}_t^{(s)} + \frac{1}{d} \sum_{r \in B_s} \delta_t^{(s,r)}. \quad (10)$$

1438 ³Even though we present our theory for the activation manipulation, our results are easily extendable to the
 1439 gradient manipulation as well.

1458 For the momentum terms, we can write:

$$\begin{aligned} 1459 \quad \mathbf{m}_{t,\text{obs}}^{(s)} &= \beta_h \mathbf{m}_{t-1,\text{obs}}^{(s)} + (1 - \beta_h) \hat{\mathbf{h}}_t^{(s)}, \\ 1460 \quad \mathbf{m}_{t,\text{nom}}^{(s)} &= \beta_h \mathbf{m}_{t-1,\text{nom}}^{(s)} + (1 - \beta_h) \bar{\mathbf{h}}_t^{(s)}. \end{aligned} \quad (11)$$

1463 The deviation in the momentum at iteration t is:

$$\begin{aligned} 1464 \quad \Delta \mathbf{m}_t^{(s)} &= \mathbf{m}_{t,\text{obs}}^{(s)} - \mathbf{m}_{t,\text{nom}}^{(s)} \\ 1465 \quad &= \beta_h (\mathbf{m}_{t-1,\text{obs}}^{(s)} - \mathbf{m}_{t-1,\text{nom}}^{(s)}) + (1 - \beta_h) (\hat{\mathbf{h}}_t^{(s)} - \bar{\mathbf{h}}_t^{(s)}) \\ 1466 \quad &= \beta_h \Delta \mathbf{m}_{t-1}^{(s)} + (1 - \beta_h) \frac{1}{d} \sum_{r \in B_s} \boldsymbol{\delta}_t^{(s,r)}. \end{aligned} \quad (12)$$

1470 Assuming $\Delta \mathbf{m}_0^{(s)} = \mathbf{0}$, we can solve this recurrence relation:

$$\Delta \mathbf{m}_t^{(s)} = (1 - \beta_h) \sum_{j=1}^t \beta_h^{t-j} \frac{1}{d} \sum_{r \in B_s} \boldsymbol{\delta}_j^{(s,r)}. \quad (13)$$

1473 Taking the norm and applying the triangle inequality, we have:

$$\begin{aligned} 1474 \quad \|\Delta \mathbf{m}_t^{(s)}\| &\leq (1 - \beta_h) \sum_{j=1}^t \beta_h^{t-j} \frac{1}{d} \sum_{r \in B_s} \|\boldsymbol{\delta}_j^{(s,r)}\| \\ 1475 \quad &\leq (1 - \beta_h) \sum_{j=1}^t \beta_h^{t-j} \frac{|B_s|}{d} \varepsilon \\ 1476 \quad &= (1 - \beta_h) \gamma_s \varepsilon \sum_{j=1}^t \beta_h^{t-j} \\ 1477 \quad &= (1 - \beta_h) \gamma_s \varepsilon \frac{1 - \beta_h^t}{1 - \beta_h} \\ 1478 \quad &\leq (1 - \beta_h) \gamma_s \varepsilon \frac{1}{1 - \beta_h} \\ 1479 \quad &= \gamma_s \varepsilon. \end{aligned} \quad (14)$$

1480 which establishes the stated bound. For the case where we consider only the most recent iteration's
1481 effect (equivalent to initializing $\mathbf{m}_{t-1,\text{obs}}^{(s)} = \mathbf{m}_{t-1,\text{nom}}^{(s)}$), we have:

$$\begin{aligned} 1482 \quad \|\Delta \mathbf{m}_t^{(s)}\| &= \left\| (1 - \beta) \frac{1}{d} \sum_{r \in B_s} \boldsymbol{\delta}_t^{(s,r)} \right\| \\ 1483 \quad &\leq (1 - \beta) \frac{|B_s|}{d} \varepsilon \\ 1484 \quad &= (1 - \beta) \gamma_s \varepsilon. \end{aligned} \quad (15)$$

1485 \square

1486 Lemma 2 establishes a key property of momentum-based smoothing: it naturally attenuates the impact
1487 of malicious perturbations. This attenuation is proportional to the fraction of malicious workers γ_s ,
1488 demonstrating that smaller malicious coalitions have less impact on the global state. This result is
1489 critical for understanding how effectively the system can contain malicious influence.

1490 Building on this foundation, we now analyze how these bounded perturbations affect our detection
1491 statistics:

1512

Lemma 3 (Test Statistic Deviation). *For a metric function Ω , assume that the detector computes*

$$\Gamma_t^{(s,r)} = \|\Omega(\mathbf{h}_t^{(s,r)}, \mathbf{m}_{t-1}^{(s)}) - \Omega_{\text{ref}}^{(s)}\| \quad (16)$$

1513

where $\Omega_{\text{ref}}^{(s)}$ is a reference statistic computed by the trusted trainer nodes (e.g., our median based reference statistic).

1514

Assume for every worker replica $r \in \{1, 2, \dots, d\}$:

1515

- *Activation perturbation: $\|\delta_t^{(s,r)}\| \leq \varepsilon$.*
- *Momentum update: $\mathbf{m}_t^{(s)} = \beta_h \mathbf{m}_{t-1}^{(s)} + (1 - \beta_h) \frac{1}{d} \sum_{r=1}^d \mathbf{h}_t^{(s,r)}$ with at most a fraction γ_s malicious workers (from Lemma 2)*

1516

$$\|\Delta \mathbf{m}_t^{(s)}\| \leq \gamma_s \varepsilon. \quad (17)$$

1517

- *Lipschitz continuity of Ω : For any inputs \mathbf{x}, \mathbf{y} and perturbations δ_x, δ_y :*

1518

$$\|\Omega(\mathbf{x} + \delta_x, \mathbf{y} + \delta_y) - \Omega(\mathbf{x}, \mathbf{y})\| \leq L_\Omega (\|\delta_x\| + \|\delta_y\|), \quad (18)$$

1519

where L_Ω is the Lipschitz constant of Ω .

1520

Define the (possibly known) baseline gap as:

1521

$$\delta^{\text{base}} = \Omega(\mathbf{h}_{t,\text{honest}}^{(s,r)}, \mathbf{m}_{t-1,\text{honest}}^{(s)}). \quad (19)$$

1522

Then, the test statistic satisfies:

1523

$$\boxed{\Gamma_t^{(s,r)} \leq \delta^{\text{base}} + L_\Omega \varepsilon + L_\Omega \gamma_s \varepsilon} \quad (20)$$

1524

1525

1526

1527

1528

1529

1530

1531

1532

1533

1534

1535

1536

1537

1538

1539

Proof. For an activation vector with malicious perturbation $\delta_t^{(s,r)}$, using the Lipschitz property of Ω , we have:

1540

$$\begin{aligned} \|\Omega(\mathbf{h}_t^{(s,r)}, \mathbf{m}_{t-1}^{(s)}) - \Omega(\mathbf{h}_{t,\text{honest}}^{(s,r)}, \mathbf{m}_{t-1}^{(s)})\| &\leq L_\Omega \|\delta_t^{(s,r)}\| \\ &\leq L_\Omega \varepsilon. \end{aligned} \quad (21)$$

1541

1542

1543

1544

1545

From Lemma 2, we know that the momentum vector deviation is bounded by $\|\Delta \mathbf{m}_{t-1}^{(s)}\| \leq \gamma_s \varepsilon$. Thus, applying the Lipschitz property again:

1546

1547

1548

$$\begin{aligned} \|\Omega(\mathbf{h}_{t,\text{honest}}^{(s,r)}, \mathbf{m}_{t-1}^{(s)}) - \Omega(\mathbf{h}_{t,\text{honest}}^{(s,r)}, \mathbf{m}_{t-1,\text{honest}}^{(s)})\| &\leq L_\Omega \|\Delta \mathbf{m}_{t-1}^{(s)}\| \\ &\leq L_\Omega \gamma_s \varepsilon. \end{aligned} \quad (22)$$

1549

1550

We can now decompose the test statistic using the triangle inequality:

1551

1552

1553

1554

1555

1556

1557

1558

1559

1560

1561

1562

1563

1564

1565

$$\begin{aligned} \Gamma_t^{(s,r)} &= \|\Omega(\mathbf{h}_t^{(s,r)}, \mathbf{m}_{t-1}^{(s)}) - \Omega_{\text{ref}}^{(s)}\| \\ &= \|\Omega(\mathbf{h}_t^{(s,r)}, \mathbf{m}_{t-1}^{(s)}) - \Omega(\mathbf{h}_{t,\text{honest}}^{(s,r)}, \mathbf{m}_{t-1}^{(s)}) \\ &\quad + \Omega(\mathbf{h}_{t,\text{honest}}^{(s,r)}, \mathbf{m}_{t-1}^{(s)}) - \Omega(\mathbf{h}_{t,\text{honest}}^{(s,r)}, \mathbf{m}_{t-1,\text{honest}}^{(s)}) \\ &\quad + \Omega(\mathbf{h}_{t,\text{honest}}^{(s,r)}, \mathbf{m}_{t-1,\text{honest}}^{(s)}) - \Omega_{\text{ref}}^{(s)}\| \\ &\leq \|\Omega(\mathbf{h}_t^{(s,r)}, \mathbf{m}_{t-1}^{(s)}) - \Omega(\mathbf{h}_{t,\text{honest}}^{(s,r)}, \mathbf{m}_{t-1}^{(s)})\| \\ &\quad + \|\Omega(\mathbf{h}_{t,\text{honest}}^{(s,r)}, \mathbf{m}_{t-1}^{(s)}) - \Omega(\mathbf{h}_{t,\text{honest}}^{(s,r)}, \mathbf{m}_{t-1,\text{honest}}^{(s)})\| \\ &\quad + \|\Omega(\mathbf{h}_{t,\text{honest}}^{(s,r)}, \mathbf{m}_{t-1,\text{honest}}^{(s)}) - \Omega_{\text{ref}}^{(s)}\| \\ &\leq L_\Omega \varepsilon + L_\Omega \gamma_s \varepsilon + \delta^{\text{base}} \end{aligned} \quad (23)$$

If the detector compensates for (or ignores) the baseline gap δ^{base} and raises an alarm when $\Gamma_t^{(s,r)} > \tau$, the additional deviation attributable only to malicious perturbations is:

$$\Gamma_{\text{pert}} \leq L_\Omega \varepsilon + L_\Omega \gamma_s \varepsilon = L_\Omega (1 + \gamma_s) \varepsilon, \quad (24)$$

1566 so a malicious worker can remain undetected provided:
 1567

$$1568 \quad L_\Omega(1 + \gamma_s)\varepsilon \leq \tau, \quad \Rightarrow \quad \boxed{\varepsilon \leq \frac{\tau}{L_\Omega(1 + \gamma_s)}} \quad (25)$$

1570 which completes the proof. \square
 1571

1573 Lemma 3 provides a crucial bound on the test statistic deviation under malicious perturbations. The
 1574 bound depends on two key factors: (1) the Lipschitz constant L_Ω of the test function and (2) the
 1575 fraction of malicious workers γ_s . The practical implication is that a malicious worker can remain
 1576 undetected only if its perturbation magnitude satisfies:

$$1577 \quad \varepsilon \leq \frac{\tau}{L_\Omega(1 + \gamma_s)} \quad (26)$$

1579 This establishes a direct relationship between the detection threshold τ and the maximum undetectable
 1580 perturbation magnitude. This equation demonstrates how tuning τ affects the security-performance
 1581 tradeoff: with lower thresholds we can provide stronger security guarantees at the potential cost of
 1582 increased false positives. This highlights the importance of setting an appropriate threshold for the
 1583 test statistic.

1585 E.1.2 GRADIENT PERTURBATION ANALYSIS

1587 Now that we have established bounds on undetectable activation perturbations, we analyze how these
 1588 perturbations propagate through the network to affect parameter gradients. This analysis is critical
 1589 for understanding the impact on training dynamics.

1591 **Lemma 4** (Per-stage Lipschitz constants). *Assume replica r of stage s implements a map*

$$1592 \quad \mathbf{h}^{(s,r)} = f_s(\mathbf{h}^{(s-1,r)}; \boldsymbol{\theta}^{(s)})$$

1593 *whose Jacobians satisfy*

$$1595 \quad \|\partial_{\boldsymbol{\theta}} f_s\| \leq L_{\boldsymbol{\theta}}^{(s)}, \quad \|\partial_{\mathbf{h}} f_s\| \leq L_f^{(s)}.$$

1596 *Then, the parameter gradient of stage s obeys*

$$1598 \quad \|\nabla_{\boldsymbol{\theta}^{(s)}} \mathcal{L}(\boldsymbol{\theta})\| \leq L_{\boldsymbol{\theta}}^{(s)} \frac{1}{d} \sum_{r=1}^d \|\mathbf{g}^{(s,r)}\|,$$

1601 *where $\mathbf{g}^{(s+1,r)}$ is the gradient with respect to the activation $\mathbf{h}^{(s,r)}$.*

1604 *Proof.* The loss $\mathcal{L}(\boldsymbol{\theta})$ depends on $\boldsymbol{\theta}^{(s)}$ only through the composition of stage maps:

$$1606 \quad \mathbf{h}^{(s,r)} = f_s(\mathbf{h}^{(s-1,r)}; \boldsymbol{\theta}^{(s)})$$

$$1607 \quad \mathbf{h}^{(s+1,r)} = f_{s+1}(\mathbf{h}^{(s,r)}; \boldsymbol{\theta}^{(s+1)})$$

$$1608 \quad \vdots$$

$$1610 \quad \mathbf{h}^{(p,r)} = f_p(\mathbf{h}^{(p-1,r)}; \boldsymbol{\theta}^{(p)})$$

1612 followed by a readout $\mathcal{L}_{\text{head}}(\mathbf{h}^{(p,r)})$.

1613 Applying the chain rule yields

$$1614 \quad \nabla_{\boldsymbol{\theta}^{(s,r)}} \mathcal{L}(\boldsymbol{\theta}) = \partial_{\boldsymbol{\theta}} f_s(\partial_{\mathbf{h}} f_{s+1}) \cdots (\partial_{\mathbf{h}} f_p) \nabla_{\mathbf{h}^{(p,r)}} \mathcal{L}_{\text{head}}. \quad (28)$$

1616 Here, each $\partial_{\boldsymbol{\theta}} f_s$ is evaluated at $(\mathbf{h}^{(s-1,r)}, \boldsymbol{\theta}^{(s)})$ and each $\partial_{\mathbf{h}} f_j$ at $(\mathbf{h}^{(j-1,r)}, \boldsymbol{\theta}^{(j)})$. Define

$$1617 \quad \mathbf{g}^{(s+1,r)} := (\partial_{\mathbf{h}} f_{s+1}) \cdots (\partial_{\mathbf{h}} f_p) \nabla_{\mathbf{h}^{(p,r)}} \mathcal{L}_{\text{head}}, \quad (29)$$

1619 so that $\mathbf{g}^{(s+1,r)}$ is precisely the activation gradient that enters replica r of stage s during
 back-propagation. The all-reduce operation aggregates all gradients from the d replicas of stage s

1620 before applying them, i.e.,
 1621

$$\nabla_{\theta^{(s)}}^{\text{agg}} \mathcal{L}(\theta) = \frac{1}{d} \sum_{r=1}^d \nabla_{\theta^{(s,r)}} \mathcal{L}(\theta) \quad (30)$$

1622 Taking Euclidean norms, applying the triangle inequality, and using sub-multiplicativity of the
 1623 operator norm yields

$$\begin{aligned} \|\nabla_{\theta^{(s)}}^{\text{agg}} \mathcal{L}(\theta)\| &\leq \frac{1}{d} \sum_{r=1}^d \|\nabla_{\theta^{(s,r)}} \mathcal{L}(\theta)\| \\ &\leq \frac{1}{d} \sum_{r=1}^d \|\partial_{\theta} f_s\| \|\mathbf{g}^{(s+1,r)}\| \\ &\leq L_{\theta}^{(s)} \frac{1}{d} \sum_{r=1}^d \|\mathbf{g}^{(s+1,r)}\|. \end{aligned} \quad (31)$$

1624 Note that the step in Eq. (30) follow the fact that the Lipschitz constant assumptions are uniform over
 1625 the data distribution. \square
 1626

1627 Lemma 4 characterizes how strongly the parameter gradients at each stage depend on activation
 1628 perturbations. The Lipschitz constants $L_{\theta}^{(s)}$ and $L_f^{(s)}$ quantify this relationship, providing a foundation
 1629 for understanding gradient sensitivity. These stage-specific Lipschitz constants are important because
 1630 they reveal which stages of the model are most vulnerable to malicious manipulation. Stages with
 1631 larger constants amplify perturbations more strongly, making them prime targets for attackers and
 1632 priority areas for enhanced monitoring.

1633 Building on these Lipschitz properties, we now quantify exactly how activation perturbations translate
 1634 to gradient perturbations:

1635 **Lemma 5** (Sensitivity of Parameter Gradient to Activation Perturbation). *Let an expected honest
 1636 replica in stage s be activation $\mathbf{h}^{(s,r)}$. Assume that a malicious worker replaces it by $\mathbf{h}^{(s,r)} + \delta$.
 1637 If the activation perturbations are small such that changing the input activation by δ perturbs
 1638 $\mathbf{g}^{(s+1,r)}$ through the local Jacobian only, then the change in the aggregated stage s parameter
 1639 gradient satisfies*

$$\|\Delta \nabla_{\theta^{(s)}}^{\text{agg}} \mathcal{L}(\theta)\| := \|\nabla_{\theta^{(s)}}^{\text{agg}} \mathcal{L}(\theta \mid \mathbf{h}^{(s,r)} + \delta) - \nabla_{\theta^{(s)}}^{\text{agg}} \mathcal{L}(\theta \mid \mathbf{h}^{(s,r)})\| \leq \frac{G_s}{d} \|\delta\|,$$

1640 where $G_s := L_{\theta}^{(s)} \left(\prod_{j \geq s+1} L_f^{(j)} \right)$. More generally, if a set B_s of $|B_s| = \gamma_s \cdot d$ malicious replicas
 1641 each injects a perturbation of norm at most $\|\delta\|$, then

$$\|\Delta \nabla_{\theta^{(s)}}^{\text{agg}}\| \leq \gamma_s \cdot G_s \cdot \varepsilon. \quad (32)$$

1642 *Proof.* Let us first consider a single replica r at stage s . For this replica, using Lemma 4 we can
 1643 write:

$$\nabla_{\theta^{(s,r)}} \mathcal{L}(\theta) = \partial_{\theta} f_s \mathbf{g}^{(s+1,r)}, \quad \text{s.t. } \mathbf{g}^{(s+1,r)} = \left(\prod_{j=s+1}^p \partial_{\mathbf{h}} f_j \right) \nabla_{\mathbf{h}^{(p,r)}} \mathcal{L}_{\text{head}}.$$

1644 Thus, assuming a linearization of the change in gradient signal under small input perturbation, we
 1645 can write:

$$\delta \mathbf{g}^{(s+1,r)} = \left(\prod_{j=s+1}^p \partial_{\mathbf{h}} f_j \right) \delta.$$

1646 Hence, for replica r we can write the sensitivity of the parameter gradient as:

$$\Delta \nabla_{\theta^{(s,r)}} \mathcal{L}(\theta) = (\partial_{\theta} f_s) \delta \mathbf{g}^{(s+1,r)} = \partial_{\theta} f_s \left(\prod_{j=s+1}^p \partial_{\mathbf{h}} f_j \right) \delta.$$

1674 Using sub-multiplicativity and the Lipschitz bounds, we have
 1675

$$\begin{aligned}
 1676 \quad \|\Delta \nabla_{\theta^{(s)}} \mathcal{L}(\theta)\| &\leq \|\partial_{\theta} f_s\| \left(\prod_{j=s+1}^p \|\partial_{\theta} f_j\| \right) \|\delta\| \\
 1677 \\
 1678 \quad &\leq L_{\theta}^{(s)} \left(\prod_{j \geq s+1} L_f^{(j)} \right) \|\delta\| \\
 1679 \\
 1680 \quad &= G_s \|\delta\| \\
 1681 \\
 1682
 \end{aligned} \tag{33}$$

1683 Since the stage update uses the aggregate gradients, we can write

$$1684 \quad \Delta \nabla_{\theta^{(s)}}^{\text{agg}} \mathcal{L}(\theta) = \frac{1}{d} \Delta \nabla_{\theta^{(s,r)}} \mathcal{L}(\theta).
 1685$$

1686 Hence, $\|\Delta \nabla_{\theta^{(s)}}^{\text{agg}} \mathcal{L}(\theta)\| \leq \frac{G_s}{d} \|\delta\|$. If $|B_s|$ replicas are corrupted, we would have

$$\begin{aligned}
 1687 \quad \|\Delta \nabla_{\theta^{(s)}}^{\text{agg}} \mathcal{L}(\theta)\| &\leq \frac{1}{d} \sum_{r \in B_s} \|\Delta \nabla_{\theta^{(s,r)}} \mathcal{L}(\theta)\| \\
 1688 \\
 1689 \quad &\leq \frac{|B_s|}{d} \cdot G_s \cdot \|\delta\| \\
 1690 \\
 1691 \quad &= \gamma_s \cdot G_s \cdot \|\delta\|,
 1692 \\
 1693
 \end{aligned} \tag{34}$$

1694 and the proof is complete. \square

1695 Lemma 5 provides the crucial link between activation perturbations and their impact on parameter
 1696 gradients. The amplification factor G_s represents how perturbations at stage s propagate through the
 1697 network during backpropagation. This factor depends on both the local parameter gradient sensitivity
 1698 ($L_{\theta}^{(s)}$) and the product of activation gradient sensitivities in subsequent stages ($\prod_{j \geq s+1} L_f^{(j)}$). This
 1699 result has important implications for robustness against malicious workers in pipeline-parallel training:
 1700

- 1701 1. Earlier stages (lower s) typically have larger amplification factors because perturbations
 1702 must propagate through more subsequent stages.
- 1703 2. Stages with larger parameter counts or complex activation patterns may have higher individual
 1704 Lipschitz constants.
- 1705 3. The fractional impact of malicious workers is reduced by the averaging effect of the all-
 1706 reduce operation, as captured by the γ_s factor.

1707 Combined with our detection bounds, we can now establish the maximum parameter gradient
 1708 perturbation that can be induced by undetected malicious workers:

$$1709 \quad \|\Delta \nabla_{\theta^{(s)}}^{\text{agg}} \mathcal{L}(\theta)\| \leq \gamma_s \cdot G_s \cdot \frac{\tau}{L_{\Omega}(1 + \gamma_s)} := \zeta \tag{35}$$

1710 This bound directly links detection thresholds to gradient perturbations, which will be essential for
 1711 our convergence analysis.

1712 E.1.3 CONVERGENCE UNDER PERTURBED GRADIENTS

1713 Having established bounds on gradient perturbations, we now analyze how these perturbations affect
 1714 the convergence properties of momentum-SGD. We consider general non-convex loss functions, but
 1715 our results can be easily extended to the strongly convex case.

1716 **Theorem 2** (Convergence of Momentum SGD under Smoothness for Convex and Non-convex
 1717 Cases with Perturbation and Noise). *Consider the balanced momentum update:*

$$\begin{aligned}
 1718 \quad \mathbf{v}_{t+1} &= \beta \mathbf{v}_t + (1 - \beta) \mathbf{g}_t, \\
 1719 \quad \theta_{t+1} &= \theta_t - \eta \mathbf{v}_{t+1}
 \end{aligned} \tag{36}$$

1720 where $\mathbf{g}_t = \nabla \mathcal{L}(\theta_t) + \zeta_t + \xi_t$, with a Lyapunov function $\Psi_t = \mathcal{L}(\theta_t) + c \|\mathbf{v}_t\|^2$ for some constant
 1721 $c > 0$.

1722 *Assume:*

1728
 1729 1. \mathcal{L} is L -smooth but potentially non-convex
 1730 2. \mathcal{L} is bounded below by \mathcal{L}^*
 1731 3. $\beta \in [0, 1]$ is the momentum parameter
 1732 4. $\eta > 0$ is the learning rate
 1733 5. ζ_t is a deterministic perturbation with maximum perturbation norm $\|\zeta_t\| \leq \zeta$, and ξ_t
 1734 is zero-mean noise with $\mathbb{E}[\|\xi_t\|^2] \leq \sigma^2$

1735 For any positive constants $\varepsilon_1, \varepsilon_2, \varepsilon_3$ conditioned on the past \mathcal{F}_t , we have:

$$\mathbb{E}[\Psi_{t+1} | \mathcal{F}_t] \leq \Psi_t - \alpha \|\nabla \mathcal{L}(\boldsymbol{\theta}_t)\|^2 + C_1 \|\mathbf{v}_t\|^2 + C_2 \|\zeta_t\|^2 + D\sigma^2 \quad (37)$$

1736 Where the constants are given by:

$$\begin{aligned} \alpha &= \eta(1 - \beta) \left(1 - \varepsilon_2 - \frac{\beta}{4\varepsilon_1(1 - \beta)} - \frac{2}{\eta} \left(\frac{\eta^2 L}{2} + c \right) \left(1 - \beta + \frac{\beta}{4\varepsilon_1} \right) \right) \\ C_1 &= \left(\eta\beta\varepsilon_1 + \left(\frac{\eta^2 L}{2} + c \right) \beta(\beta + 2(1 - \beta)(\varepsilon_1 + \varepsilon_3)) - c \right) \\ C_2 &= \left(\frac{\eta(1 - \beta)}{4\varepsilon_2} + 2 \left(\frac{\eta^2 L}{2} + c \right) (1 - \beta) \left(1 - \beta - \frac{\beta}{4\varepsilon_3} \right) \right) \\ D &= \left(\frac{\eta^2 L}{2} + c \right) (1 - \beta)^2. \end{aligned} \quad (38)$$

1737 If we choose appropriate values for $\varepsilon_1, \varepsilon_2, \varepsilon_3$ such that $\alpha > 0$ and $C_1 < 0$, and assume $\mathbf{v}_0 = \mathbf{0}$,
 1738 then the algorithm converges in expectation to a neighborhood of a stationary point:

$$\boxed{\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|\nabla \mathcal{L}(\boldsymbol{\theta}_t)\|^2] \leq \frac{\mathcal{L}_0 - \mathcal{L}^*}{\alpha T} + \frac{C_2 \zeta^2 + D\sigma^2}{\alpha}}, \quad (39)$$

1739 where $\mathcal{L}_0 := \mathcal{L}(\boldsymbol{\theta}_0)$ is our loss value at initialization.

1740 Proof. We begin by analyzing one-step progress with the Lyapunov potential function
 1741 $\Psi_t = \mathcal{L}(\boldsymbol{\theta}_t) + c\|\mathbf{v}_t\|^2$ inspired by (Liu et al., 2020; Mai & Johansson, 2020).

1742 **Evolution of the Loss Term.** By the L -smoothness of \mathcal{L} , we have:

$$\begin{aligned} \mathcal{L}(\boldsymbol{\theta}_{t+1}) &\leq \mathcal{L}(\boldsymbol{\theta}_t) + \langle \nabla \mathcal{L}(\boldsymbol{\theta}_t), \boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}_t \rangle + \frac{L}{2} \|\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}_t\|^2 \\ &= \mathcal{L}(\boldsymbol{\theta}_t) - \eta \langle \nabla \mathcal{L}(\boldsymbol{\theta}_t), \mathbf{v}_{t+1} \rangle + \frac{\eta^2 L}{2} \|\mathbf{v}_{t+1}\|^2 \end{aligned} \quad (40)$$

1743 Now, expanding $\Psi_{t+1} - \Psi_t = [\mathcal{L}(\boldsymbol{\theta}_{t+1}) - \mathcal{L}(\boldsymbol{\theta}_t)] + c[\|\mathbf{v}_{t+1}\|^2 - \|\mathbf{v}_t\|^2]$, we have:

$$\begin{aligned} \Psi_{t+1} - \Psi_t &= [\mathcal{L}(\boldsymbol{\theta}_{t+1}) - \mathcal{L}(\boldsymbol{\theta}_t)] + c[\|\mathbf{v}_{t+1}\|^2 - \|\mathbf{v}_t\|^2] \\ &= -\eta \langle \nabla \mathcal{L}(\boldsymbol{\theta}_t), \mathbf{v}_{t+1} \rangle + \frac{\eta^2 L}{2} \|\mathbf{v}_{t+1}\|^2 + c\|\mathbf{v}_{t+1}\|^2 - c\|\mathbf{v}_t\|^2 \\ &= -\eta \langle \nabla \mathcal{L}(\boldsymbol{\theta}_t), \mathbf{v}_{t+1} \rangle + \left(\frac{\eta^2 L}{2} + c \right) \|\mathbf{v}_{t+1}\|^2 - c\|\mathbf{v}_t\|^2. \end{aligned} \quad (41)$$

1744 Substitute $\mathbf{v}_{t+1} = \beta\mathbf{v}_t + (1 - \beta)\mathbf{g}_t$ (Polyak, 1964), then we have:

$$\begin{aligned} \Psi_{t+1} - \Psi_t &\leq -\eta \langle \nabla \mathcal{L}(\boldsymbol{\theta}_t), \mathbf{v}_{t+1} \rangle + \left(\frac{\eta^2 L}{2} + c \right) \|\mathbf{v}_{t+1}\|^2 - c\|\mathbf{v}_t\|^2 \\ &= -\eta \langle \nabla \mathcal{L}(\boldsymbol{\theta}_t), \beta\mathbf{v}_t + (1 - \beta)\mathbf{g}_t \rangle + \left(\frac{\eta^2 L}{2} + c \right) \|\beta\mathbf{v}_t + (1 - \beta)\mathbf{g}_t\|^2 - c\|\mathbf{v}_t\|^2 \\ &= -\eta\beta \langle \nabla \mathcal{L}(\boldsymbol{\theta}_t), \mathbf{v}_t \rangle - \eta(1 - \beta) \langle \nabla \mathcal{L}(\boldsymbol{\theta}_t), \mathbf{g}_t \rangle \\ &\quad + \left(\frac{\eta^2 L}{2} + c \right) (\beta^2 \|\mathbf{v}_t\|^2 + (1 - \beta)^2 \|\mathbf{g}_t\|^2 + 2\beta(1 - \beta) \langle \mathbf{v}_t, \mathbf{g}_t \rangle) - c\|\mathbf{v}_t\|^2 \end{aligned} \quad (42)$$

1782 Next, we bound individual terms.
 1783

1784

1785

1786

1787 **Bounding $-\eta\beta\langle\nabla\mathcal{L}(\boldsymbol{\theta}_t), \mathbf{v}_t\rangle$.** Using Young's inequality with parameter $\varepsilon_1 > 0$, we write:

$$1788 \quad 1789 \quad -\eta\beta\langle\mathbf{v}_t, \nabla\mathcal{L}(\boldsymbol{\theta}_t)\rangle \leq \eta\beta\varepsilon_1\|\mathbf{v}_t\|^2 + \frac{\eta\beta}{4\varepsilon_1}\|\nabla\mathcal{L}(\boldsymbol{\theta}_t)\|^2. \quad (43)$$

1790

1791

1792

1793

1794 **Bounding $-\eta(1-\beta)\langle\nabla\mathcal{L}(\boldsymbol{\theta}_t), \mathbf{g}_t\rangle$.** Since $\mathbf{g}_t = \nabla\mathcal{L}(\boldsymbol{\theta}_t) + \boldsymbol{\zeta}_t + \boldsymbol{\xi}_t$, we can write:

$$1795 \quad 1796 \quad -\eta(1-\beta)\langle\nabla\mathcal{L}(\boldsymbol{\theta}_t), \mathbf{g}_t\rangle = -\eta(1-\beta)\langle\nabla\mathcal{L}(\boldsymbol{\theta}_t), \nabla\mathcal{L}(\boldsymbol{\theta}_t) + \boldsymbol{\zeta}_t + \boldsymbol{\xi}_t\rangle \\ 1797 \quad \leq -\eta(1-\beta)\|\nabla\mathcal{L}(\boldsymbol{\theta}_t)\|^2 - \eta(1-\beta)\langle\nabla\mathcal{L}(\boldsymbol{\theta}_t), \boldsymbol{\zeta}_t\rangle \\ 1798 \quad \quad - \eta(1-\beta)\langle\nabla\mathcal{L}(\boldsymbol{\theta}_t), \boldsymbol{\xi}_t\rangle \\ 1799 \quad \leq -\eta(1-\beta)\|\nabla\mathcal{L}(\boldsymbol{\theta}_t)\|^2 + \eta(1-\beta)\varepsilon_2\|\nabla\mathcal{L}(\boldsymbol{\theta}_t)\|^2 + \frac{\eta(1-\beta)}{4\varepsilon_2}\|\boldsymbol{\zeta}_t\|^2 \\ 1800 \quad \quad - \eta(1-\beta)\langle\nabla\mathcal{L}(\boldsymbol{\theta}_t), \boldsymbol{\xi}_t\rangle \quad (44)$$

1801

1802

1803

1804

1805

1806

1807 **Bounding $\|\mathbf{g}_t\|^2$.** For this term, we write:

$$1808 \quad 1809 \quad \|\mathbf{g}_t\|^2 = \|\nabla\mathcal{L}(\boldsymbol{\theta}_t) + \boldsymbol{\zeta}_t + \boldsymbol{\xi}_t\|^2 \\ 1810 \quad = \|\nabla\mathcal{L}(\boldsymbol{\theta}_t) + \boldsymbol{\zeta}_t\|^2 + \|\boldsymbol{\xi}_t\|^2 + 2\langle\nabla\mathcal{L}(\boldsymbol{\theta}_t) + \boldsymbol{\zeta}_t, \boldsymbol{\xi}_t\rangle \\ 1811 \quad \leq 2\|\nabla\mathcal{L}(\boldsymbol{\theta}_t)\|^2 + 2\|\boldsymbol{\zeta}_t\|^2 + \|\boldsymbol{\xi}_t\|^2 + 2\langle\nabla\mathcal{L}(\boldsymbol{\theta}_t) + \boldsymbol{\zeta}_t, \boldsymbol{\xi}_t\rangle, \quad (45)$$

1812

1813

1814

1815

1816 **Bounding $\langle\mathbf{v}_t, \mathbf{g}_t\rangle$.** Expanding $\mathbf{g}_t = \nabla\mathcal{L}(\boldsymbol{\theta}_t) + \boldsymbol{\zeta}_t + \boldsymbol{\xi}_t$, we have:

$$1817 \quad 1818 \quad \langle\mathbf{v}_t, \mathbf{g}_t\rangle = \langle\mathbf{v}_t, \nabla\mathcal{L}(\boldsymbol{\theta}_t) + \boldsymbol{\zeta}_t + \boldsymbol{\xi}_t\rangle \\ 1819 \quad = \langle\mathbf{v}_t, \nabla\mathcal{L}(\boldsymbol{\theta}_t)\rangle + \langle\mathbf{v}_t, \boldsymbol{\zeta}_t\rangle + \langle\mathbf{v}_t, \boldsymbol{\xi}_t\rangle \\ 1820 \quad \leq \varepsilon_1\|\mathbf{v}_t\|^2 + \frac{1}{4\varepsilon_1}\|\nabla\mathcal{L}(\boldsymbol{\theta}_t)\|^2 + \varepsilon_3\|\mathbf{v}_t\|^2 + \frac{1}{4\varepsilon_3}\|\boldsymbol{\zeta}_t\|^2 + \langle\mathbf{v}_t, \boldsymbol{\xi}_t\rangle \quad (46)$$

1821

1822

1823

1824

1825 **Combining Terms.** Taking conditional expectation from Eq. (42) and substituting the previous
 1826 bounds, we have

$$1827 \quad 1828 \quad \mathbb{E}[\Psi_{t+1}|\mathcal{F}_t] - \Psi_t \leq \eta\beta\varepsilon_1\|\mathbf{v}_t\|^2 + \frac{\eta\beta}{4\varepsilon_1}\|\nabla\mathcal{L}(\boldsymbol{\theta}_t)\|^2 \\ 1829 \quad \quad - \eta(1-\beta)\|\nabla\mathcal{L}(\boldsymbol{\theta}_t)\|^2 + \eta(1-\beta)\varepsilon_2\|\nabla\mathcal{L}(\boldsymbol{\theta}_t)\|^2 + \frac{\eta(1-\beta)}{4\varepsilon_2}\|\boldsymbol{\zeta}_t\|^2 \\ 1830 \quad \quad + \left(\frac{\eta^2L}{2} + c\right)(\beta^2\|\mathbf{v}_t\|^2 + 2(1-\beta)^2\|\nabla\mathcal{L}(\boldsymbol{\theta}_t)\|^2 + 2(1-\beta)^2\|\boldsymbol{\zeta}_t\|^2 + (1-\beta)^2\sigma^2) \\ 1831 \quad \quad 2\left(\frac{\eta^2L}{2} + c\right)\beta(1-\beta)\left((\varepsilon_1 + \varepsilon_3)\|\mathbf{v}_t\|^2 + \frac{1}{4\varepsilon_1}\|\nabla\mathcal{L}(\boldsymbol{\theta}_t)\|^2 + \frac{1}{4\varepsilon_3}\|\boldsymbol{\zeta}_t\|^2\right) - c\|\mathbf{v}_t\|^2 \quad (47)$$

1832

1833

1834

1835

1836 **Collecting Terms and Setting Bounds.** After substituting all bounds and collecting terms, we
 1837 have:
 1838

$$\begin{aligned}
 \mathbb{E}[\Psi_{t+1} | \mathcal{F}_t] - \Psi_t &\leq \\
 &- \eta(1-\beta) \left(1 - \varepsilon_2 - \frac{\eta\beta}{4\varepsilon_1\eta(1-\beta)} - \frac{2}{\eta} \left(\frac{\eta^2 L}{2} + c \right) \left(1 - \beta + \frac{\beta}{4\varepsilon_1} \right) \right) \|\nabla \mathcal{L}(\boldsymbol{\theta}_t)\|^2 \\
 &+ \left(\eta\beta\varepsilon_1 + \left(\frac{\eta^2 L}{2} + c \right) \beta (\beta + 2(1-\beta)(\varepsilon_1 + \varepsilon_3)) - c \right) \|\mathbf{v}_t\|^2 \\
 &+ \left(\frac{\eta(1-\beta)}{4\varepsilon_2} + 2 \left(\frac{\eta^2 L}{2} + c \right) (1-\beta) \left(1 - \beta - \frac{\beta}{4\varepsilon_3} \right) \right) \|\zeta_t\|^2 \\
 &+ \left(\frac{\eta^2 L}{2} + c \right) (1-\beta)^2 \sigma^2
 \end{aligned} \tag{48}$$

1849 Define the following constants:
 1850

$$\begin{aligned}
 \alpha &= \eta(1-\beta) \left(1 - \varepsilon_2 - \frac{\beta}{4\varepsilon_1(1-\beta)} - \frac{2}{\eta} \left(\frac{\eta^2 L}{2} + c \right) \left(1 - \beta + \frac{\beta}{4\varepsilon_1} \right) \right) \\
 C_1 &= \left(\eta\beta\varepsilon_1 + \left(\frac{\eta^2 L}{2} + c \right) \beta (\beta + 2(1-\beta)(\varepsilon_1 + \varepsilon_3)) - c \right) \\
 C_2 &= \left(\frac{\eta(1-\beta)}{4\varepsilon_2} + 2 \left(\frac{\eta^2 L}{2} + c \right) (1-\beta) \left(1 - \beta - \frac{\beta}{4\varepsilon_3} \right) \right) \\
 D &= \left(\frac{\eta^2 L}{2} + c \right) (1-\beta)^2
 \end{aligned} \tag{49}$$

1862 **Establishing Convergence.** For convergence, we can set the variables such that $\alpha > 0$ and $C_1 < 0$.
 1863 The one-step progress in expectation becomes:
 1864

$$\mathbb{E}[\Psi_{t+1} | \mathcal{F}_t] \leq \Psi_t - \alpha \|\nabla \mathcal{L}(\boldsymbol{\theta}_t)\|^2 + C_1 \|\mathbf{v}_t\|^2 + C_2 \|\zeta_t\|^2 + D\sigma^2 \tag{50}$$

1866 When $C_1 < 0$, the term with $\|\mathbf{v}_t\|^2$ helps convergence. Taking the full expectation and summing
 1867 from $t = 0$ to $T - 1$:

$$\begin{aligned}
 \sum_{t=0}^{T-1} \mathbb{E}[\alpha \|\nabla \mathcal{L}(\boldsymbol{\theta}_t)\|^2] &\leq \mathbb{E}[\Psi_0] - \mathbb{E}[\Psi_T] + \sum_{t=0}^{T-1} (C_2 \|\zeta_t\|^2 + D\sigma^2) \\
 &\leq \mathbb{E}[\Psi_0] - \mathcal{L}^* + T(C_2 \zeta^2 + D\sigma^2)
 \end{aligned} \tag{51}$$

1873 Where we used $\mathcal{L}^* \leq \mathcal{L}(\boldsymbol{\theta}_t)$ and dropped the negative term with $C_1 < 0$. Substituting
 1874 $\Psi_0 = \mathcal{L}_0 + c\|\mathbf{v}_0\|^2 = \mathcal{L}_0$ and dividing by αT we have:
 1875

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|\nabla \mathcal{L}(\boldsymbol{\theta}_t)\|^2] \leq \frac{\mathcal{L}_0 - \mathcal{L}^*}{\alpha T} + \frac{C_2 \zeta^2 + D\sigma^2}{\alpha}. \tag{52}$$

1879 Thus, the average squared gradient norm converges to a neighborhood determined by the perturbation
 1880 magnitude ζ^2 and noise variance σ^2 . This result shows that momentum SGD with perturbation and
 1881 noise converges to a neighborhood of a stationary point in the non-convex smooth case. \square
 1882

1884 Theorem 2 provides convergence guarantees for momentum-SGD in the non-convex setting, which
 1885 is particularly relevant for deep learning applications like LLMs. Instead of convergence to a
 1886 neighborhood of the optimum, we provide guarantees on the average gradient norm, a standard
 1887 measure for non-convex optimization. The bound depends directly on the perturbation magnitude,
 1888 establishing that even in non-convex settings, controlling malicious perturbations through effective
 1889 detection mechanisms is crucial for ensuring convergence to stationary points.

1890 E.1.4 UNIFIED ANALYSIS: DETECTION-CONVERGENCE RELATIONSHIP
18911892 We now unify our results to establish a comprehensive relationship between detection thresholds and
1893 convergence guarantees. This unified perspective provides us with clear guidance on the security-
1894 performance tradeoff.
18951896 **Theorem 3** (Convergence Guarantees for Distributed Training with Malicious Workers). *Let
1897 a distributed training system with p stages, each having d worker replicas, where a fraction
1898 $\gamma_s < \frac{1}{2}$ of workers at stage s are malicious. Using the momentum-based verification with a test
1899 statistic detector with threshold τ , let us assume we will use balanced momentum SGD with
1900 parameter $\beta \in [0, 1)$ and learning rate $\eta > 0$ for optimization.*
1901 Assume:1902 1. The test statistic function Ω is Lipschitz continuous with constant L_Ω ,
1903 2. Each stage s implements a map $\mathbf{h}^{(s,r)} = f_s(\mathbf{h}^{(s-1,r)}; \theta^{(s)})$ with Jacobian bounds
1904 $\|\partial_\theta f_s\| \leq L_\theta^{(s)}$ and $\|\partial_\mathbf{h} f_s\| \leq L_f^{(s)}$,
1905 3. The loss function \mathcal{L} is L -smooth and bounded below by \mathcal{L}^* ,
1906 4. The stochastic gradient includes zero-mean noise with variance bounded by σ^2 .
19071908 Then the following results hold:
19091910 • **Detection Evasion Bound.** For a malicious worker to remain undetected by the test statistic
1911 detector:
1912

1913
$$\varepsilon \leq \frac{\tau}{L_\Omega(1 + \gamma_s)}. \quad (53)$$

1914 • **Parameter Gradient Perturbation.** The maximum parameter gradient perturbation that
1915 can be induced by undetected malicious workers at stage s is:
1916

1917
$$\|\Delta \nabla_{\theta^{(s)}}^{\text{agg}}\| \leq \gamma_s \cdot G_s \cdot \frac{\tau}{L_\Omega(1 + \gamma_s)} \quad (54)$$

1918 where $G_s = L_\theta^{(s)} \left(\prod_{j \geq s+1} L_f^{(j)} \right)$ represents the amplification factor for perturbations.
19191920 • **Convergence Bounds.** Under the maximum undetected perturbation and assuming
1921 $\zeta := \gamma_s \cdot G_s \cdot \frac{\tau}{L_\Omega(1 + \gamma_s)}$, for non-convex loss we have:
1922

1923
$$\boxed{\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|\nabla \mathcal{L}(\theta_t)\|^2] \leq \frac{\mathcal{L}_0 - \mathcal{L}^*}{\alpha T} + \frac{C_2 \zeta^2 + D \sigma^2}{\alpha}} \quad (55)$$

1924

1925 where constants are as defined in Theorem 2.
19261927 *Proof.* We prove the theorem by connecting the results from Lemmas 2 and 5 and Theorem 2. \square
1928

1929

1930 E.1.5 RECOVERING WELL-KNOWN LOWER-BOUNDS FOR SGD CONVERGENCE FROM
1931 THEOREM 3
19321933 To evaluate our convergence theorem's validity, we examine whether it generalizes to common
1934 non-convex optimization bounds. Consider vanilla SGD without malicious perturbations: setting
1935 $\beta = 0$ (relaxing momentum SGD to SGD), $\zeta = 0$ (no malicious perturbation), and substituting into
1936 our coefficients from Eq. (38) with $c \rightarrow 0^+$:
1937

1938
$$\alpha \approx \mathcal{O}(\eta) \quad (56)$$

1939
$$D \approx \mathcal{O}(\eta^2 L),$$

1940

1941 where η is our learning rate (see Eq. (36)).
1942

1943

1944

Substituting these coefficients into our convergence theorem, we have:

1945

1946

1947

1948

1949

1950

1951

1952

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|\nabla \mathcal{L}(\theta_t)\|^2] \leq \frac{\mathcal{L}_0 - \mathcal{L}^*}{\alpha T} + \frac{C_2 \zeta^2 + D\sigma^2}{\alpha} \quad (57)$$

$$\stackrel{(a)}{\leq} \frac{\mathcal{L}_0 - \mathcal{L}^*}{\alpha T} + \frac{D\sigma^2}{\alpha} \quad (58)$$

$$\stackrel{(b)}{\leq} \mathcal{O}\left(\frac{\mathcal{L}_0 - \mathcal{L}^*}{\eta T} + \eta L\sigma^2\right) \quad (59)$$

1953

where (a) assumes no perturbation, and (b) uses the derived constants.

1954

1955

This matches the classical SGD bound from Koloskova et al. (2024):

1956

1957

1958

“SGD, Ex. 3.1 Since $\sigma^2 \leq \tau\sigma_{\text{SGD}}^2$ (see Table 2), and using that $\tau = \Theta(1/L\gamma)$ the convergence rate in Theorem 5.1 converts to

1959

1960

1961

$$\frac{1}{T} \sum_{t=0}^T \mathbb{E}\|\nabla f(x_t)\|^2 \leq \mathcal{O}\left(\frac{F_0}{\gamma T} + L\gamma\sigma_{\text{SGD}}^2\right),$$

1962

1963

with $\gamma \leq \frac{1}{8\sqrt{3}L}$ [and where $F_0 = f(\mathbf{x}_0) - f^*$]. This recovers classical convergence rate of SGD for non-convex functions (up to constants)."

1964

1965

This bound exactly matches what we derive from our convergence theorem.

1966

1967

E.1.6 THEORY IMPLICATIONS

1968

1969

Our unified analysis reveals several important implications for designing malicious-tolerant SWARM verification:

1970

1971

1972

1973

1974

- **Security-Convergence Tradeoff:** The detection threshold τ directly impacts the convergence guarantees through its effect on the maximum undetected perturbation. Lower thresholds provide stronger security guarantees but may increase false positives and potentially slow convergence due to unnecessary worker exclusion.
- **Byzantine Fraction Impact:** As shown in Eq. (54), the maximum parameter deviation is proportional to $\frac{\gamma_s}{1+\gamma_s}$, where γ_s represents the fraction of malicious workers at stage s . This monotonically increasing function implies that larger malicious fractions allow more severe parameter deviations, highlighting the importance of maintaining an honest majority.
- **Detector Sensitivity:** A detector Ω with larger Lipschitz constant L_Ω reduces the parameter gradient deviation. In practical terms, employing a more sensitive detection function constrains the potential impact of malicious workers by allowing them less room for undetected perturbation.
- **Stage Vulnerability:** Stages with higher values of the amplification factor G_s are more vulnerable to malicious perturbations. In typical neural network architectures, this often means that earlier layers (which affect all subsequent computation) have greater vulnerability. This suggests that security resources should be prioritized to monitor these critical stages more closely.
- **Momentum as Robustness against Malicious Behavior:** Higher momentum values β naturally reduce the impact of per-iteration perturbations by placing more weight on the historical gradient estimates. This provides an inherent form of robustness that complements explicit detection mechanisms. The optimal momentum value therefore depends not only on optimization dynamics but also on security considerations.
- **Adaptive Detection Thresholds:** Our analysis suggests that detection thresholds could be optimally set differently for each stage based on their amplification factors G_s . Stages with higher amplification factors should use stricter thresholds to maintain consistent convergence guarantees across the model. We leave this for future work.

1994

1995

1996

1997

E.2 PROOF OF HONEST MAJORITY GUARANTEE

This section provides the complete proof of Lemma 1 from Sec. 3.2.

1998

1999

2000

2001

2002

2003

2004

2005

2006

2007

2008

2009

2010

2011

2012

2013

2014

2015

2016

2017

Lemma 1 (Honest Majority Guarantee). *Consider our distributed training system with p pipeline stages, each replicated across d worker nodes. Let b be the total number of malicious workers, and $\epsilon \in (0, 1)$ be a small positive constant. If workers are assigned to each stage randomly and*

$$b \leq \frac{dp}{2} - p\sqrt{\frac{d}{2} \ln\left(\frac{p}{\epsilon}\right)}, \quad (60)$$

then with probability at least $1 - \epsilon$ every pipeline stage has strictly fewer than $d/2$ malicious workers.

In other words, Lemma 1 makes a connection between the initial pool of malicious workers and the conditions under which a “random worker assignment” to each stage would preserve the per-stage honest majority assumption needed by SENTINEL. This is important as usually we do not know apriori if a worker is Byzantine or not, and since usually a random worker assignment is used to allocate workers, we need to ensure that the honest majority assumption is preserved.

Proof. Let B_s denote the set of malicious worker nodes at stage s . We model the assignment of malicious nodes as follows: each worker node has probability $q = b/n = b/(dp)$ of being malicious, independently of other nodes.

For any stage s , when stages are randomly assigned to workers, the number of malicious nodes $|B_s|$ follows a binomial distribution with parameters d and q :

$$|B_s| \sim \text{Binomial}(d, q), \quad \mathbb{E}[|B_s|] = qd. \quad (61)$$

Our goal is to ensure that, with high probability, every stage s has $|B_s| < d/2$. Using Hoeffding’s inequality (Hoeffding, 1994) for sums of independent Bernoulli random variables, we have

$$\Pr [|B_s| - \mathbb{E}[|B_s|] \geq t] \leq \exp\left(-\frac{2t^2}{d}\right). \quad (62)$$

Setting $t = d/2 - qd = d(1/2 - q)$, we obtain

$$\Pr [|B_s| \geq d/2] \leq \exp\left(-2d\left(\frac{1}{2} - q\right)^2\right). \quad (63)$$

Applying the union bound across all p stages, the probability that at least one stage has a majority of malicious workers is bounded by

$$\Pr [\exists s : |B_s| \geq d/2] \leq p \cdot \exp\left(-2d\left(\frac{1}{2} - q\right)^2\right). \quad (64)$$

For this probability to be at most ϵ , we require

$$p \cdot \exp\left(-2d\left(\frac{1}{2} - q\right)^2\right) \leq \epsilon. \quad (65)$$

Taking the natural logarithm of both sides and solving for q , we get

$$\frac{1}{2} - q \geq \sqrt{\frac{\ln(p/\epsilon)}{2d}}. \quad (66)$$

Substituting $q = b/n = b/(dp)$ and solving for b , we obtain the maximum allowable number of malicious nodes:

$$b_{\max} = \left(\frac{1}{2} - \sqrt{\frac{\ln(p/\epsilon)}{2d}}\right) dp = \frac{dp}{2} - p\sqrt{\frac{d}{2} \ln\left(\frac{p}{\epsilon}\right)}. \quad (67)$$

Therefore, if the total number of malicious nodes b is at most b_{\max} , then with probability at least $1 - \epsilon$, all pipeline stages will have strictly fewer than $d/2$ malicious worker nodes. \square

2052 This theoretical result demonstrates that as our system scales with more replicas per stage, it becomes
 2053 increasingly robust against adversarial workers, approaching the theoretical limit of tolerating up to
 2054 half of all workers being malicious.
 2055

2056 F EXTENDED EXPERIMENTAL RESULTS

2058 This section provides comprehensive experimental details supporting our main findings. We first
 2059 describe the experimental setup and hyper-parameters, followed by an extended version of our results.
 2060

2061 F.1 DETAILED EXPERIMENTAL SETTINGS

2063 **Experimental Infrastructure.** To simulate a heterogeneous distributed environment, we developed
 2064 our experiments using the TorchTitan (Liang et al., 2025) framework built on PyTorch (Paszke
 2065 et al., 2019).⁴ Our setup employed pipeline parallelism where each transformer layer corresponds
 2066 to a pipeline stage, with data parallel replicas serving as workers within each stage. We con-
 2067 ducted experiments across three scales using 1-3 compute nodes, each equipped with 8 NVIDIA
 2068 A100-SXM4-40GB GPUs, resulting in configurations with 64, 128, and 256 total workers. Ablation
 2069 studies were performed using 8 NVIDIA A100-SXM4-80GB GPUs. **For our large scale experiment**
 2070 **with Llama-3-4B, we used 8 NVIDIA H100-HBM3-80GB GPUs with activation checkpointing.**

2071 **Model and Training Configuration.** We evaluated decoder-only Llama-3 (Dubey et al., 2024)
 2072 models with varying architectural configurations. Complete training hyper-parameters are provided
 2073 in Tab. 10. Additionally, we also evaluate the performance of our approach on NanoGPT (Karpathy,
 2074 2022) as a representative GPT2 (Radford et al., 2019) architecture, **and Llama-4-0.4B (MetaAI, 2025)**
 2075 **and DeepSeek-V3-1B (DeepSeek-AI, 2024) as representative Mixture-of-Expert models.** Unless
 2076 stated otherwise, we train all model configurations for 5000 steps.
 2077

2078 **Verification Settings.** SENTINEL deploys dedicated verifier nodes that intercept inter-stage com-
 2079 munications and monitor both forward activations and backward activation gradients for anomalous
 2080 behavior. We permanently ban workers after $c = 5$ violations. Moreover, we implement an adaptive
 2081 IQR thresholding mechanism to detect outliers, with parameters detailed in Tab. 11.⁵ These detection
 2082 thresholds were calibrated empirically through preliminary experiments on each configuration, though
 2083 additional hyper-parameter tuning may yield further improvements.

2084 As pointed out in the paper, we assume that the first layer and final two layers of each model are
 2085 operated by honest nodes. This assumption is crucial for two reasons:

- 2087 1. these layers control the primary data and gradient flow during forward and backward propagation,
 2088 making them essential for training stability, and
- 2089 2. they represent key attack surfaces for adversaries seeking to compromise model integrity through
 2090 data poisoning/backdoor attacks (Li et al., 2024) (via the input layer) or label manipulation (Big-
 2091 gio et al., 2012; Fung et al., 2020) (via the output layers).

2092 Without this honest-node assumption, malicious actors could easily circumvent intermediate verifica-
 2093 tion by corrupting inputs or outputs directly.
 2094

2095 **Datasets.** We conduct experiments on three large-scale text corpora: CommonCrawl (C4) (Raffel
 2096 et al., 2020), FineWeb (FW) (Penedo et al., 2024), and OpenWebText (OW) (Gokaslan et al., 2019),
 2097 all obtained through Hugging Face Datasets Streaming API (Lhoest et al., 2021). To enable model
 2098 evaluation, we construct a held-out validation set comprising 100k samples that remains separate
 2099 from training data throughout the process. During validation phases, we sample batches of 160
 2100 samples per data-parallel worker from this validation set for evaluation.

2101
 2102
 2103
 2104
⁴We integrate the NanoGPT following the implementation of Karpathy (2022) into TorchTitan.
 2105
⁵For our experiments on MoE-based models, we use the same settings as the Llama-3-0.6B model.

Table 10: Model architectures and training configuration.

MODEL CONFIGURATION	NANO GPT-0.25B	LLAMA-3-0.6B	LLAMA-3-1.2B	LLAMA-3-4B	LLAMA-4-0.4B	DEEPSPEECH-V3-1B
MODEL ARCHITECTURE						
Parameters	278,364,672	574,391,296	1,224,247,296	3,984,989,184	443,957,760	967,989,760
Hidden Dimensions	768	1024	2048	3072	512	1024
Number of Layers (Stages)	12	16	8	22	16	15
Attention Heads	12	32	32	32	32	16
Key-Value Heads	–	8	8	8	8	–
Number of Experts	–	–	–	–	8	8
Number of Shared Experts	–	–	–	–	1	2
Top-K (for MoE)	–	–	–	–	1	6
RoPE Theta	–	500000	500000	500000	500000	500000
FFN Dimension Multiplier	–	–	1.3	1.3	1.3	1.3
Multiple Of	–	–	1024	1024	1024	1024
RoPE Theta	–	500000	500000	500000	500000	500000
DISTRIBUTED SETUP						
Data Parallel Dimension	8	8 (16 for 16×16 mesh)	8	8	8	8
Pipeline Parallel Dimension	12	16	22	22	16	15
Total Workers	96	128 (256 for 16×16 mesh)	64	176	128	90
OPTIMIZER						
Type	AdamW	AdamW	AdamW	AdamW	AdamW	AdamW
Learning Rate	6e-4	6e-4	6e-4	6e-4	4e-3	2.2e-4
Epsilon	1e-8	1e-8	1e-8	1e-8	1e-15	1e-8
LEARNING RATE SCHEDULER						
Warmup Steps	100	100	100	500	600	600
Decay Ratio	0.8	0.8	0.8	0.8	–	0.8
Decay Type	Linear	Linear	Linear	Linear	Linear	Cosine
Minimum LR	0.0	0.0	0.0	0.0	0.1	0.1
TRAINING						
Worker Batch Size	8	12	10	12	16	16
Global Batch Size	64	96 (192 for 16×16 mesh)	80	96	128	128
Sequence Length	1024	1024	1024	1024	1024	1024
Gradient Clipping	1.0	1.0	1.0	1.0	1.0	1.0

2160
 2161
 2162
 2163
 2164
 2165
 2166
 2167
 2168
 2169
 2170
 2171
 2172
 2173
 2174
 2175
 2176
 2177
 2178
 2179
 2180
 2181
 2182
 2183
 2184
 2185
 2186
 2187
 2188
 2189
 2190
 2191
 2192
 2193
 2194
 2195
 2196
 2197
 2198
 2199
 2200
 2201
 2202
 2203
 2204
 2205
 2206
 2207
 2208
 2209
 2210
 2211
 2212
 2213

Table 11: Hyper-parameters of the proposed verification approach (Alg. 1).

PARAMETER	NANO GPT			LLAMA-3		
	8 × 12 MESH, 0.25B		8 × 16 MESH, 0.6B	16 × 16 MESH, 0.6B		8 × 8 MESH, 1.2B
	ACTIVATION	GRADIENT	ACTIVATION	GRADIENT	ACTIVATION	GRADIENT
Momentum Beta (β)	0.90	0.80	0.90	0.80	0.90	0.80
Initial IQR Multiplier (k_0)	1.5	3.0	1.5	3.0	1.5	3.0
Target FP-Rate (α)	0.01	0.001	0.0001	0.001	0.01	0.001
Adaptive Max Iterations (N_{\max})	10	10	10	10	10	10
Adaptive Grow Factor (γ_g)	1.1	1.01	1.1	1.01	1.1	1.01
Adaptive Shrink Factor (γ_s)	0.9	0.99	0.9	0.99	0.9	0.99
Adaptive Min Distance Multiplier (Λ)	0.35	0.15	0.15	0.05	0.25	0.05
Adaptive Epsilon (ε)	0.01	0.001	0.0005	0.0001	0.0001	0.0001

2214 F.2 DETAILED EXPERIMENTAL RESULTS
22152216 **Activation vs. Gradient Attack Analysis in Pipeline Parallelism.** Building on the subset of
2217 attacks presented in Sec. 5, we now provide comprehensive verification results across all activation
2218 and gradient attacks from Sec. 2 to fully characterize their behavior. Tab. 12 presents complete results
2219 on the C4 dataset, while Tabs. 13 and 14 demonstrate performance on FineWeb and OpenWebText
2220 datasets, respectively.2221 Our comprehensive evaluation reveals several important insights. First, activation manipulation
2222 poses an equally significant threat as gradient manipulation in distributed, pipeline parallel-based
2223 training, confirming that both attack vectors require careful consideration in Byzantine-tolerant
2224 systems. Second, when attacks evade detection, their deviation from baseline vanilla training remains
2225 negligible, directly supporting our theoretical analysis presented in Theorem 1. This consistency
2226 holds across all three datasets, demonstrating the robustness and versatility of our approach across
2227 diverse training scenarios.2228 Particularly noteworthy are attacks detected with a detection speed of 1.0, indicating that despite our
2229 forgiveness strategy introduced in Sec. 3.1, these attacks produce sufficiently substantial deviations
2230 to warrant immediate exclusion of the malicious worker. The training and validation loss curves
2231 in Figs. 7 to 10 further illustrate how our EMA verification approach effectively controls malicious
2232 behavior, maintaining performance close to vanilla baselines throughout training. These results
2233 collectively validate the effectiveness and generalizability of our proposed verification framework.2234 Table 12: Attack detection performance for Llama-3-0.6B on C4 dataset. Metrics shown include precision,
2235 recall, F1 score (all as percentages), average detection speed (in iterations), and validation loss.
2236

2237 MODE	2238 ATTACK	2239 SENTINEL (OURS)					2240 NO VERIF.
		2241 PR. (%) \uparrow	2242 RE. (%) \uparrow	2243 F1 (%) \uparrow	2244 DET. SPEED \downarrow	2245 VAL. LOSS \downarrow	
2246 -	2247 None (Vanilla)	2248 100.0	2249 100.0	2250 100.0	2251 N/A	2252 3.819	2253 3.821
2254 ACTIVATION MANIPULATION	2255 Constant (Zeros)	2256 100.0	2257 100.0	2258 100.0	2259 6.5	2260 3.809	2261 11.761
	2262 Constant (Ones)	2263 100.0	2264 100.0	2265 100.0	2266 6.33	2267 3.817	2268 7.778
	2269 Random Value	2270 100.0	2271 100.0	2272 100.0	2273 6.48	2274 3.827	2275 7.778
	2276 Scaling ($\alpha = -1$)	2277 100.0	2278 100.0	2279 100.0	2280 6.38	2281 3.824	2282 4.109
	2283 Random Sign (1%)	2284 100.0	2285 100.0	2286 100.0	2287 6.33	2288 3.825	2289 4.670
	2290 Random Sign (10%)	2291 100.0	2292 100.0	2293 100.0	2294 6.52	2295 3.822	2296 4.619
	2297 Random Sign (30%)	2298 88.9	2299 100.0	2300 94.1	2301 70.91	2302 3.841	2303 4.567
	2304 Delay (100-steps)	2305 88.9	2306 100.0	2307 94.1	2308 13.21	2309 3.841	2310 7.675
	2312 Bias Addition	2313 84.6	2314 91.7	2315 88.0	2316 14.57	2317 3.830	2318 3.892
	2320 Invisible Noise (90%)	2321 100.0	2322 100.0	2323 100.0	2324 6.48	2325 3.836	2326 7.675
2328 GRADIENT MANIPULATION	2329 Invisible Noise (95%)	2330 100.0	2331 100.0	2332 100.0	2333 6.52	2334 3.823	2335 7.677
	2337 Invisible Noise (99%)	2338 100.0	2339 100.0	2340 100.0	2341 6.48	2342 3.826	2343 7.682
	2345 Constant (Zeros)	2346 100.0	2347 100.0	2348 100.0	2349 6.42	2350 3.829	2351 3.942
	2353 Constant (Ones)	2354 88.9	2355 100.0	2356 94.1	2357 1.0	2358 3.816	2359 10.630
	2361 Random Value	2362 100.0	2363 100.0	2364 100.0	2365 1.0	2366 3.818	2367 9.595
	2369 Scaling ($\alpha = -1$)	2370 0.0	2371 0.0	2372 0.0	2373 N/A	2374 3.893	2375 3.893
	2377 Random Sign (1%)	2378 0.0	2379 0.0	2380 0.0	2381 N/A	2382 3.990	2383 3.982
	2385 Random Sign (10%)	2386 0.0	2387 0.0	2388 0.0	2389 N/A	2390 3.982	2391 3.994
	2393 Random Sign (30%)	2394 0.0	2395 0.0	2396 0.0	2397 N/A	2398 3.944	2399 3.933
	2401 Delay (100-steps)	2402 100.0	2403 100.0	2404 100.0	2405 7.33	2406 3.826	2407 10.157
2409 GRADIENT MANIPULATION	2410 Bias Addition	2411 100.0	2412 100.0	2413 100.0	2414 1.0	2415 3.828	2416 10.813
	2418 Invisible Noise (90%)	2419 100.0	2420 75.0	2421 85.7	2422 101.89	2423 3.968	2424 4.218
	2426 Invisible Noise (95%)	2427 100.0	2428 79.2	2429 88.4	2430 229.68	2431 3.954	2432 4.174
	2434 Invisible Noise (99%)	2435 100.0	2436 79.2	2437 88.4	2438 211.0	2439 3.943	2440 4.176

2268
 2269 Table 13: Attack detection performance for Llama-3-0.6B on FineWeb dataset. Metrics shown include precision,
 2270 recall, F1 score (all as percentages), average detection speed (in iterations), and validation loss.

2271 MODE	2272 ATTACK	2273 SENTINEL (OURS)					2274 NO VERIF.
		2275 PR. (%) \uparrow	2276 RE. (%) \uparrow	2277 F1 (%) \uparrow	2278 DET. SPEED \downarrow	2279 VAL. LOSS \downarrow	
2273 -	2274 None (Vanilla)	2275 100.0	2276 100.0	2277 100.0	2278 N/A	2279 3.818	2280 3.840
2281 ACTIVATION MANIPULATION	2282 Constant (Zeros)	2283 100.0	2284 100.0	2285 100.0	2286 6.43	2287 3.819	2288 11.761
	2289 Constant (Ones)	2290 100.0	2291 100.0	2292 100.0	2293 6.61	2294 3.814	2295 7.793
	2296 Random Value	2297 96.0	2298 100.0	2299 98.0	2300 6.46	2301 3.831	2302 7.793
	2303 Scaling ($\alpha = -1$)	2304 100.0	2305 100.0	2306 100.0	2307 6.29	2308 3.827	2309 4.121
	2310 Random Sign (1%)	2311 100.0	2312 100.0	2313 100.0	2314 6.42	2315 3.825	2316 4.693
	2317 Random Sign (10%)	2318 51.1	2319 100.0	2320 67.6	2321 3.58	2322 3.829	2323 4.716
	2324 Random Sign (30%)	2325 92.3	2326 100.0	2327 96.0	2328 7.5	2329 3.826	2330 4.564
	2331 Delay (100-steps)	2332 85.7	2333 100.0	2334 92.3	2335 14.83	2336 3.832	2337 7.692
	2338 Bias Addition	2339 86.4	2340 79.2	2341 82.6	2342 10.68	2343 3.828	2344 3.898
	2345 Invisible Noise (90%)	2346 100.0	2347 100.0	2348 100.0	2349 6.38	2350 3.824	2351 7.709
	2352 Invisible Noise (95%)	2353 100.0	2354 100.0	2355 100.0	2356 6.29	2357 3.824	2358 7.713
	2359 Invisible Noise (99%)	2360 100.0	2361 100.0	2362 100.0	2363 6.33	2364 3.829	2365 7.712
2366 GRADIENT MANIPULATION	2367 Constant (Zeros)	2368 100.0	2369 100.0	2370 100.0	2371 6.33	2372 3.815	2373 3.949
	2374 Constant (Ones)	2375 92.3	2376 100.0	2377 96.0	2378 1.14	2379 3.824	2380 10.726
	2381 Random Value	2382 100.0	2383 100.0	2384 100.0	2385 1.38	2386 3.831	2387 9.359
	2388 Scaling ($\alpha = -1$)	2389 0.0	2390 0.0	2391 0.0	2392 N/A	2393 3.888	2394 3.901
	2395 Random Sign (1%)	2396 0.0	2397 0.0	2398 0.0	2399 N/A	2400 3.989	2401 3.999
	2402 Random Sign (10%)	2403 0.0	2404 0.0	2405 0.0	2406 N/A	2407 3.990	2408 4.004
	2409 Random Sign (30%)	2410 0.0	2411 0.0	2412 0.0	2413 N/A	2414 3.941	2415 3.966
	2416 Delay (100-steps)	2417 100.0	2418 100.0	2419 100.0	2420 7.29	2421 3.817	2422 10.017
	2423 Bias Addition	2424 100.0	2425 100.0	2426 100.0	2427 1.14	2428 3.828	2429 10.573
	2430 Invisible Noise (90%)	2431 100.0	2432 75.0	2433 85.7	2434 52.94	2435 3.954	2436 4.212
	2437 Invisible Noise (95%)	2438 100.0	2439 75.0	2440 85.7	2441 140.89	2442 3.959	2443 4.217
	2444 Invisible Noise (99%)	2445 100.0	2446 75.0	2447 85.7	2448 209.56	2449 3.949	2450 4.197

2295
 2296 Table 14: Attack detection performance for Llama-3-0.6B on OpenWebText dataset. Metrics shown include
 2297 precision, recall, F1 score (all as percentages), average detection speed (in iterations), and validation loss.

2298 MODE	2299 ATTACK	2300 SENTINEL (OURS)					2301 NO VERIF.
		2302 PR. (%) \uparrow	2303 RE. (%) \uparrow	2304 F1 (%) \uparrow	2305 DET. SPEED \downarrow	2306 VAL. LOSS \downarrow	
2301 -	2302 None (Vanilla)	2303 100.0	2304 100.0	2305 100.0	2306 N/A	2307 3.773	2308 3.778
2309 ACTIVATION MANIPULATION	2310 Constant (Zeros)	2311 100.0	2312 100.0	2313 100.0	2314 6.33	2315 3.773	2316 11.761
	2317 Constant (Ones)	2318 100.0	2319 100.0	2320 100.0	2321 6.29	2322 3.779	2323 7.820
	2324 Random Value	2325 100.0	2326 100.0	2327 100.0	2328 6.29	2329 3.777	2330 7.821
	2331 Scaling ($\alpha = -1$)	2332 100.0	2333 100.0	2334 100.0	2335 6.29	2336 3.776	2337 4.016
	2338 Random Sign (1%)	2339 100.0	2340 100.0	2341 100.0	2342 6.21	2343 3.774	2344 4.614
	2346 Random Sign (10%)	2347 100.0	2348 100.0	2349 100.0	2350 6.29	2351 3.779	2352 4.578
	2354 Random Sign (30%)	2355 96.0	2356 100.0	2357 98.0	2358 25.62	2359 3.779	2360 4.524
	2363 Delay (100-steps)	2364 88.9	2365 100.0	2366 94.1	2367 10.74	2368 3.790	2369 7.701
	2372 Bias Addition	2373 91.7	2374 91.7	2375 91.7	2376 14.0	2377 3.782	2378 3.843
	2381 Invisible Noise (90%)	2382 100.0	2383 100.0	2384 100.0	2385 6.25	2386 3.777	2387 7.677
	2390 Invisible Noise (95%)	2391 100.0	2392 100.0	2393 100.0	2394 6.5	2395 3.783	2396 7.674
	2399 Invisible Noise (99%)	2400 100.0	2401 100.0	2402 100.0	2403 6.29	2404 3.780	2405 7.678
2406 GRADIENT MANIPULATION	2407 Constant (Zeros)	2408 100.0	2409 100.0	2410 100.0	2411 5.7	2412 3.775	2413 3.908
	2414 Constant (Ones)	2415 100.0	2416 100.0	2417 100.0	2418 1.0	2419 3.774	2420 10.722
	2423 Random Value	2424 100.0	2425 100.0	2426 100.0	2427 1.0	2428 3.780	2429 9.611
	2432 Scaling ($\alpha = -1$)	2433 0.0	2434 0.0	2435 0.0	2436 N/A	2437 3.847	2438 3.860
	2441 Random Sign (1%)	2442 12.1	2443 16.7	2444 14.0	2445 1516.33	2446 3.990	2447 3.957
	2450 Random Sign (10%)	2451 5.3	2452 4.2	2453 4.7	2454 2646.0	2455 3.967	2456 3.952
	2459 Random Sign (30%)	2460 0.0	2461 0.0	2462 0.0	2463 N/A	2464 3.901	2465 3.906
	2473 Delay (100-steps)	2474 100.0	2475 100.0	2476 100.0	2477 7.09	2478 3.771	2479 11.508
	2482 Bias Addition	2483 100.0	2484 100.0	2485 100.0	2486 1.0	2487 3.777	2488 10.947
	2491 Invisible Noise (90%)	2492 100.0	2493 87.5	2494 93.3	2495 92.95	2496 3.877	2497 4.178
	2500 Invisible Noise (95%)	2501 100.0	2502 91.7	2503 95.7	2504 136.42	2505 3.828	2506 4.098
	2509 Invisible Noise (99%)	2510 100.0	2511 87.5	2512 93.3	2513 64.42	2514 3.858	2515 4.123

2322

2323

2324

2325

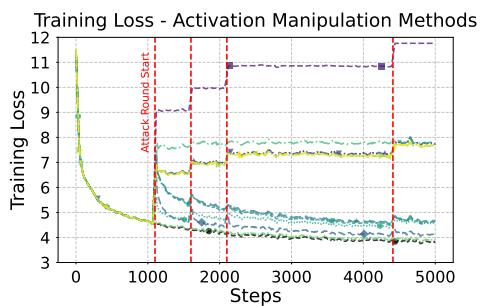
2326

2327

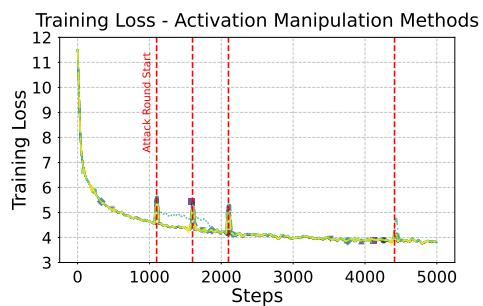
2328

2329

● Vanilla (No Attack) ▲ Random Value ▲ Random Sign (10%) ▲ Bias Addition ▲ Invisible Noise (95%)
 ■ Constant (Zeros) ▲ Scaling (-1) ▲ Random Sign (30%) ▲ Invisible Noise (90%) ▲ Invisible Noise (99%)
 ▲ Constant (Ones) ▲ Random Sign (1%) ▲ Delay (100 Steps)

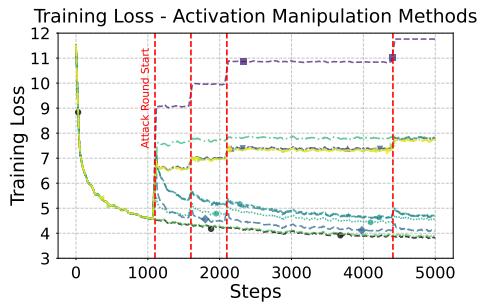


Without verification

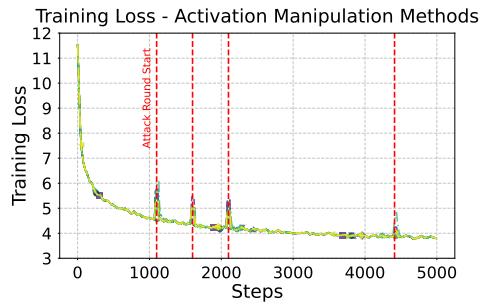


With verification

(a) C4 Dataset

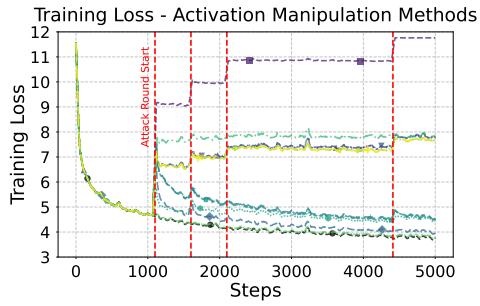


Without verification

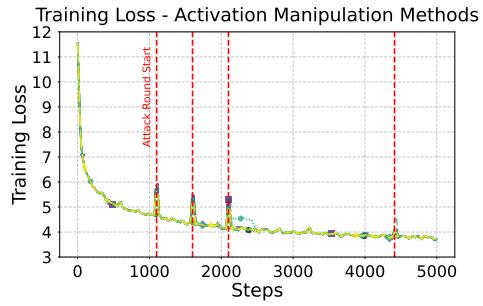


With verification

(b) FineWeb Dataset



Without verification



With verification

(c) OpenWebText Dataset

Figure 7: **Training loss** comparing our verification mechanism against baseline vanilla training under **activation manipulation attacks**. We evaluate on Llama-3-0.6B using three datasets (C4, FineWeb, and OpenWebText). Dotted vertical lines indicate attack initiation points where 6 randomly selected nodes begin submitting adversarial activations in coordinated Byzantine attacks. Our verification approach maintains stable convergence while the baseline suffers significant degradation under attack.

2366

2367

2368

2369

2370

2371

2372

2373

2376

2377

2378

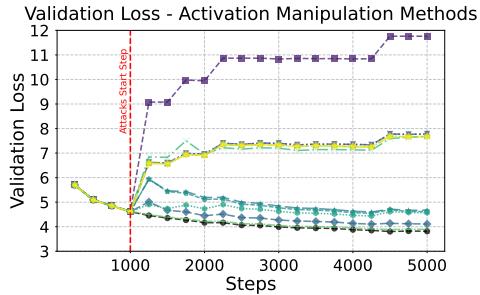
2379

2380

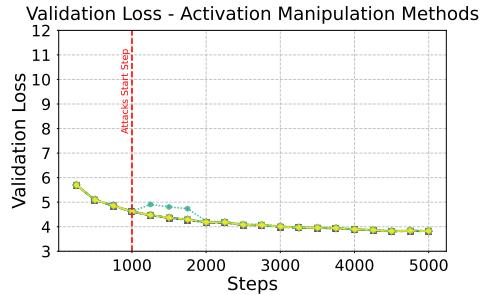
2381

2382

2383

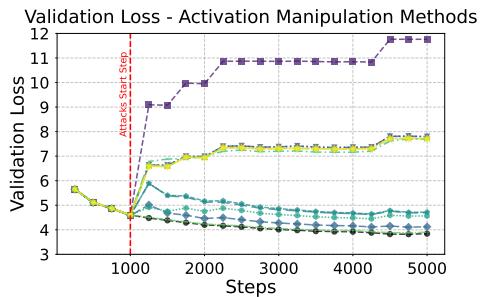


Without verification

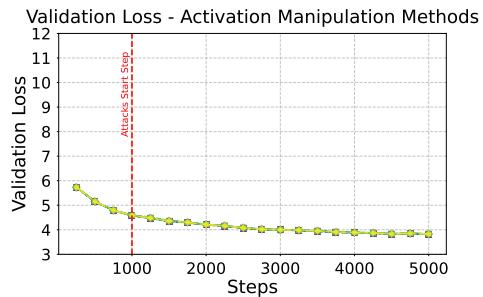


With verification

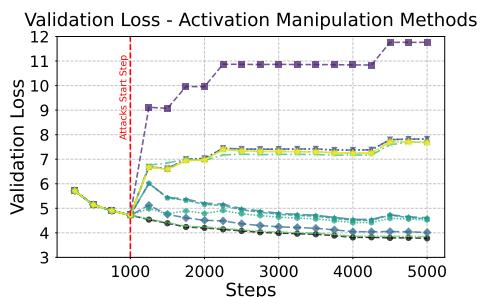
(a) C4 Dataset



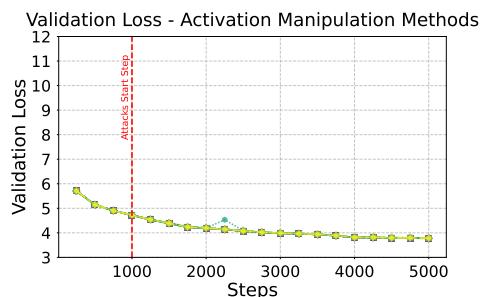
Without verification



(b) FineWeb Dataset



Without verification



(c) OpenWebText Dataset

Figure 8: **Validation loss** comparing our verification mechanism against baseline vanilla training under **activation manipulation attacks**. We evaluate Llama-3-0.6B across three datasets (C4, FineWeb, and OpenWebText). The dotted red line marks the transition from warm-up to the attack phase, where Byzantine nodes begin submitting adversarial activations. Our verification approach maintains stable validation performance while without it, the baseline shows significant degradation post-attack.

2420

2421

2422

2423

2424

2425

2426

2427

2428

2429

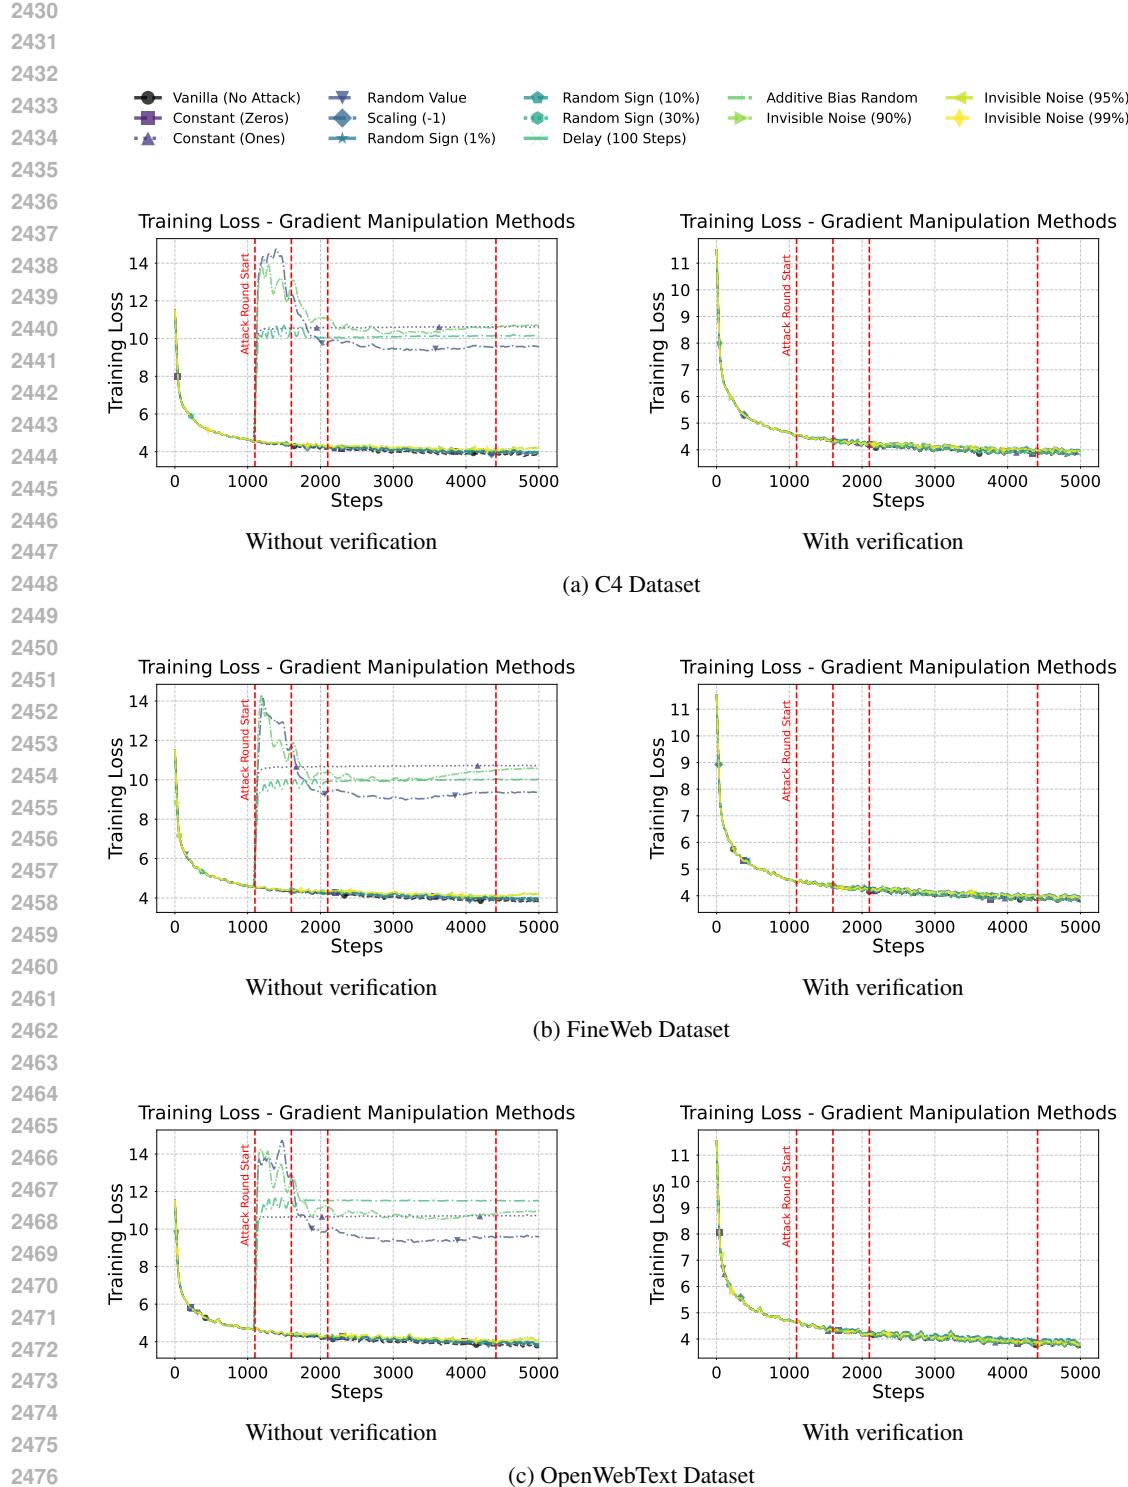


Figure 9: **Training loss** comparing our verification mechanism against baseline vanilla training under **gradient manipulation attacks**. We evaluate on Llama-3-0.6B using three datasets (C4, FineWeb, and OpenWebText). Dotted vertical lines indicate attack initiation points where 6 randomly selected nodes begin submitting adversarial gradients in coordinated Byzantine attacks. Our verification approach maintains stable convergence while the baseline suffers significant degradation under attack.

2484

2485

2486

2487

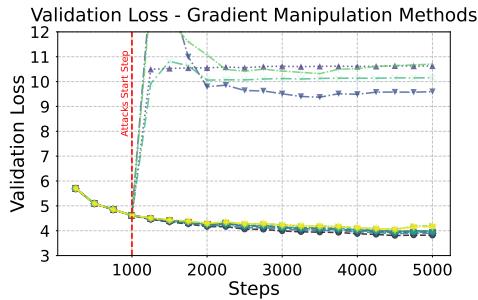
2488

2489

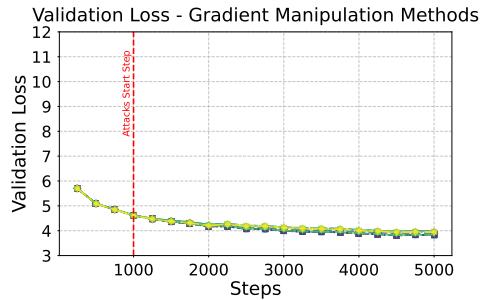
2490

2491

Vanilla (No Attack) Random Value Random Sign (10%) Additive Bias Random Invisible Noise (95%)
 Constant (Zeros) Scaling (-1) Random Sign (30%) Invisible Noise (90%) Invisible Noise (99%)
 Constant (Ones) Random Sign (1%) Delay (100 Steps)

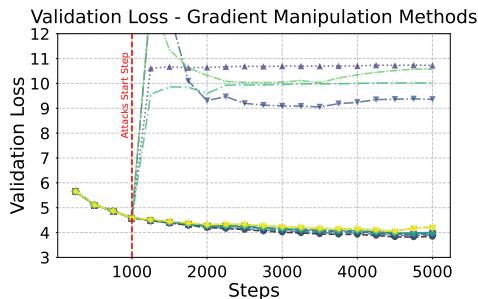


Without verification

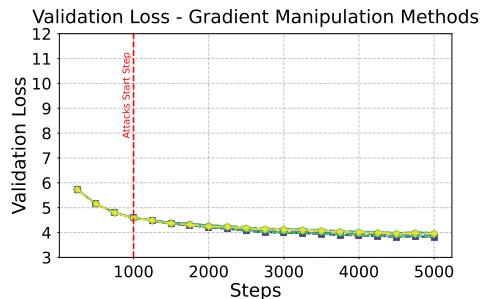


With verification

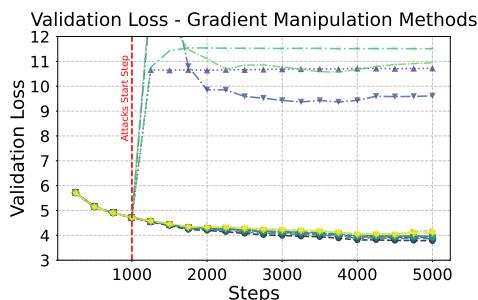
(a) C4 Dataset



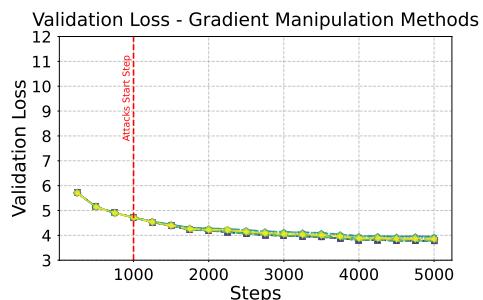
Without verification



(b) FineWeb Dataset



Without verification



(c) OpenWebText Dataset

Figure 10: **Validation loss** comparing our verification mechanism against baseline vanilla training under **gradient manipulation attacks**. We evaluate Llama-3-0.6B across three datasets (C4, FineWeb, and OpenWebText). The dotted red line marks the transition from warm-up to the attack phase, where Byzantine nodes begin submitting adversarial gradients. Our verification approach maintains stable validation performance while without it, the baseline shows significant degradation post-attack.

2528

2529

2530

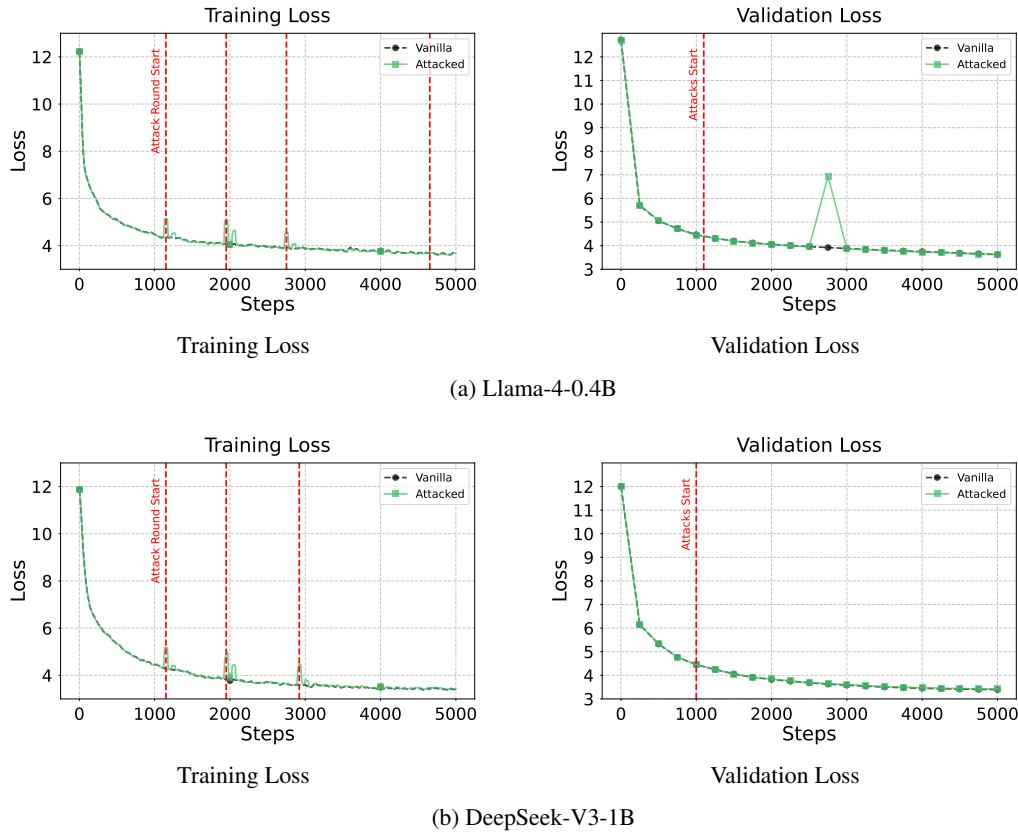
2531

2532

2533

2534

2535



2565 Figure 11: Training and validation loss evolution for MoE-based models trained on FineWeb-EDU dataset over
2566 5k iterations. The mixed activation attack scenario assumes 37.5% Byzantine nodes at each pipeline stage,
2567 with 30% randomly selected nodes performing a randomly chosen attack at a randomly sampled iteration. Each node
2568 has a different activation manipulation method (as outlined in Sec. 2.1).

2569
2570
2571
2572
2573
2574
2575
2576
2577
2578
2579
2580
2581
2582
2583
2584
2585
2586
2587
2588
2589
2590
2591

Training and Validation Loss Evolution for MoE-based Models. In Sec. 5, we discussed how SENTINEL extends to MoE-based models such as Llama-4 and DeepSeek-V3. Here, we show the training and validation loss evolution throughout training. As seen from Fig. 11, training loss starts to display higher values after a round of attacks start. However, the validation loss keeps going down similar to the vanilla baseline and match its performance.

Evolution of Adaptive Deviation Bounds. Our approach employs an adaptive IQR-based thresholding mechanism, outline in Alg. 5, that dynamically adjusts acceptable deviation bounds for each monitored metric. To demonstrate the effectiveness of this approach, we present the evolution of these adaptive thresholds alongside the corresponding deviations recorded for each worker across two representative layers of our Llama-3-0.6B model. Fig. 12a shows results under gradient delay attacks, while Fig. 12b illustrates behavior during activation random sign attacks. The results demonstrate that our adaptive bounds effectively encapsulate the normal operational behavior of honest workers during benign training phases. Critically, when a malicious worker initiates an attack, the adaptive bounds enable verifier nodes to immediately detect and flag the anomalous behavior, providing robust protection against adversarial interference in the distributed training process.

Large-scale Experimental Results. In Sec. 5, we discussed scaling our experiments to two additional settings: (1) a 16×16 mesh topology with 256 workers, and (2) a 1.2B parameter model on an 8×8 mesh with 64 workers. Here we present the complete results for both configurations across all attack types and malicious ratios tested. Tab. 15 shows detailed performance metrics including validation loss, precision, recall, and F1-score for each experimental condition. Our comprehensive results corroborate the main findings regarding the effectiveness of our proposed verification approach in detecting malicious actors that attempt to disrupt distributed training.

2592

2593

2594

2595

2596

2597

2598

2599

2600

2601

2602

2603

2604

2605

2606

2607

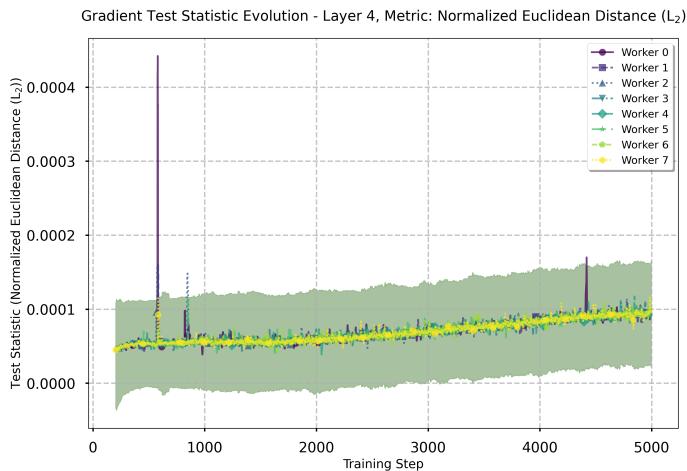
2608

2609

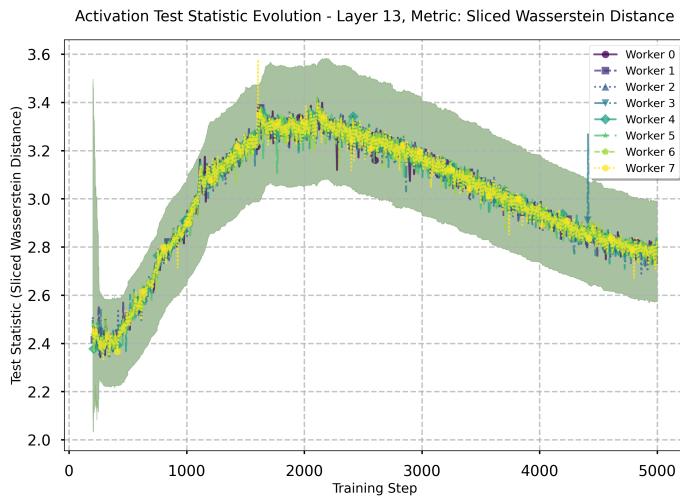
2610

2611

2612



(a) Normalized L_2 distance evolution of gradients at layer 4 under gradient delay attack. Worker 0 initiates attack at iteration 4411 and is immediately flagged.



(b) Sliced Wasserstein distance evolution of activations at layer 13 under random sign attack (1%). Workers 7 and 3 initiate attacks at iterations 1600 and 4411, respectively, and are immediately flagged.

Figure 12: Evolution of adaptive deviation bounds and worker statistics. The proposed thresholding mechanism adapts to natural distribution shifts while detecting Byzantine behavior.

Additionally, we demonstrated in Tab. 4 that we can scale the model size easily and SENTINEL still successfully work for these settings. Fig. 13 shows the training and validation loss behavior throughout training. It demonstrates that our approach mitigates the malicious workers that aim to lead the model into divergence and allows the training to continue with a similar behavior to the vanilla training.

Longer Training & Alternative Architectures. Two natural questions arise from our approach: whether our verification method remains stable under longer training regimes, and whether it generalizes to alternative transformer architectures beyond our initial experiments. To address these concerns, we conduct extended evaluations on two different models: NanoGPT-0.25B and Llama-3-0.6B, training each for 30,000 iterations. This training duration corresponds to approximately 2B tokens for NanoGPT-0.25B ($7 \times$ the parameter count) and 3B tokens for Llama-3-0.6B ($5 \times$ the parameter count), following established scaling laws in LLM training (Hoffmann et al., 2022).

Table 15: Attack detection performance for large-scale Llama-3 training on C4 dataset. Metrics shown include precision, recall, F1 score (all as percentages), average detection speed (in iterations), and validation loss. For all experiments, we assume a 37.5% Byzantine workers per stage (thus, for 16×16 mesh we have 6 malicious vs. 10 honest workers per stage, while for 8×8 mesh their ratio is 3:5.)

SETUP	MODE	ATTACK	SENTINEL (OURS)				
			PR. (%) \uparrow	RE. (%) \uparrow	F1 (%) \uparrow	DET. SPEED \downarrow	VAL. LOSS \downarrow
0.6B ON 16×16 MESH	ACTIVATION	Random Value	100.0	100.0	100.0	7.96	3.900
		Delay (100-steps)	85.7	100.0	92.3	14.39	3.945
		Bias Addition	100.0	25.6	40.8	65.15	3.981
		Invisible Noise (99%)	100.0	100.0	100.0	7.96	3.898
	GRADIENT	Random Value	100.0	100.0	100.0	124.08	3.895
		Delay (100-steps)	100.0	100.0	100.0	9.05	3.890
		Bias Addition	100.0	100.0	100.0	1.69	3.894
		Invisible Noise (99%)	98.7	93.6	96.0	14.27	3.915
1.2B ON 8×8 MESH	-	None (Vanilla)	100.0	100.0	100.0	N/A	3.723
	ACTIVATION	Random Value	100.0	100.0	100.0	4.33	3.723
		Delay (100-steps)	37.5	100.0	54.5	67.0	3.774
		Bias Addition	0.0	0.0	0.0	N/A	3.738
	GRADIENT	Random Value	100.0	100.0	100.0	1.0	3.726
		Delay (100-steps)	100.0	100.0	100.0	1.0	3.722
		Bias Addition	0.0	0.0	0.0	N/A	3.805
		Invisible Noise (99%)	100.0	100.0	100.0	9.2	3.725

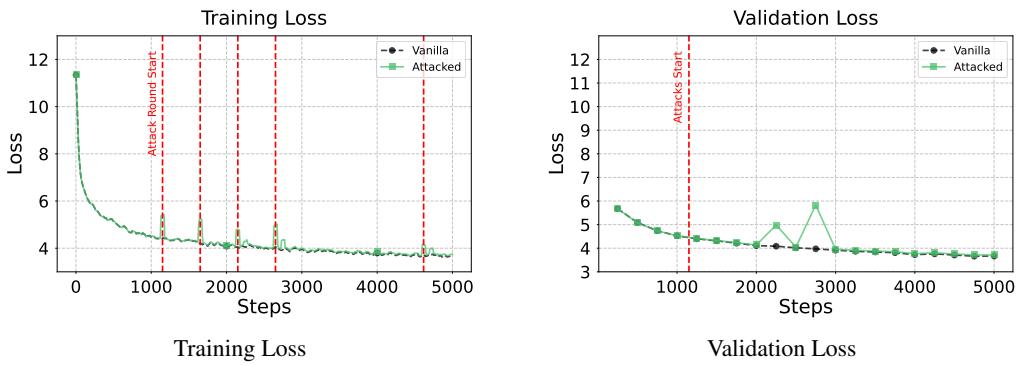


Figure 13: Training and validation loss evolution for Llama-3-4B trained on FineWeb dataset over 5k iterations. The mixed activation attack scenario assumes 37.5% Byzantine nodes at each pipeline stage, with 20% randomly selected nodes performing a randomly chosen attack at a randomly sampled iteration. Each node has a different activation manipulation method (as outlined in Sec. 2.1).

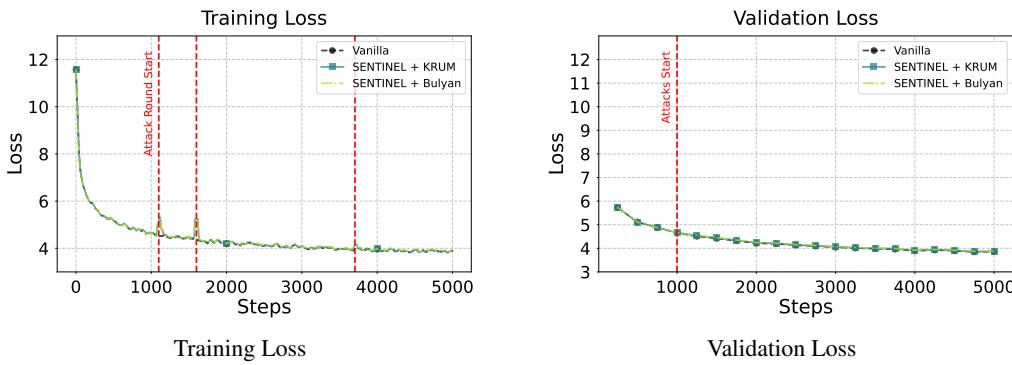
We evaluate our method’s robustness by simulating a challenging adversarial environment where 50% of nodes are malicious at each transformer stage. At each attack round, one randomly selected malicious node performs a randomly chosen attack from Tab. 12 under a “no collusion” assumption. Tab. 16 summarizes our detection performance across this extended training period for both architectures.

Our results demonstrate consistent stability across both model architectures and all three datasets, achieving high F1-scores ($> 81\%$) for attack detection. Importantly, we observe that the median detection speed across all successfully detected attack types is 5.0 iterations, which aligns precisely with our acceptable number of violations threshold. This indicates that we can detect the majority of attacks with significant training impact within our predefined tolerance window. Even when some attacks remain undetected, they exhibit negligible impact on training convergence, as illustrated in Figs. 15 and 16. The validation loss under our verification method closely tracks the vanilla baseline throughout the entire training duration for both models, corroborating our theoretical analysis from Theorem 1 and demonstrating the method’s architectural flexibility and long-term stability.

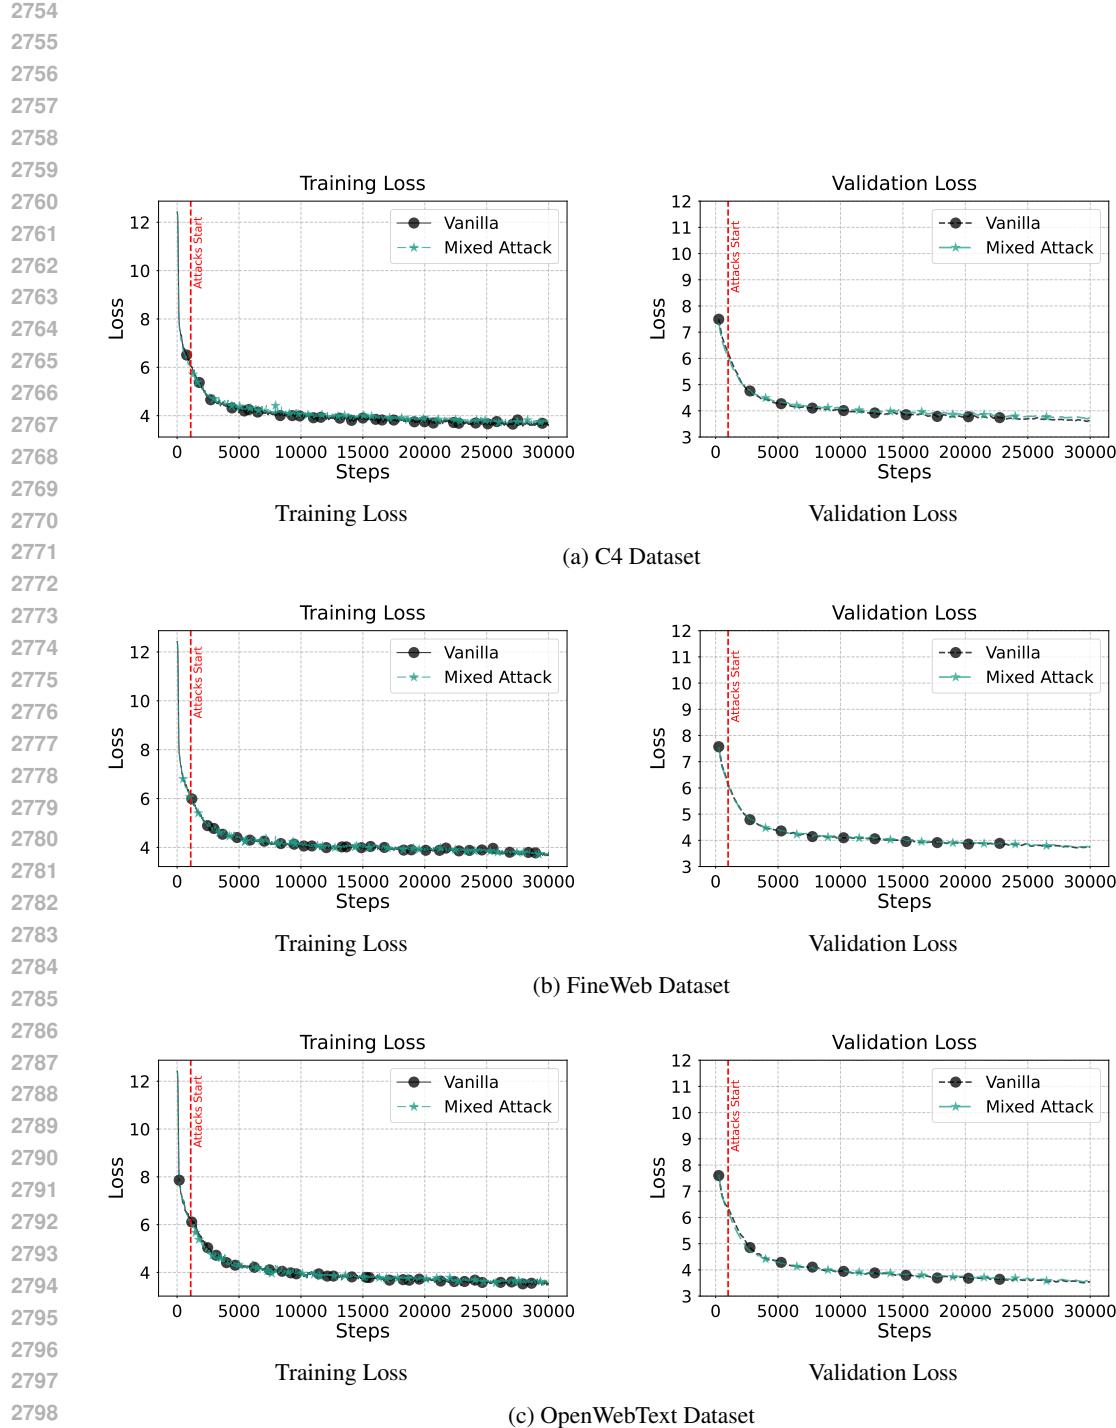
2700 Table 16: Detection performance for training NanoGPT-0.25B and Llama-3-0.6B against mixed attacks for 30k
 2701 iterations.

2703 MODEL	2704 DATASET	2705 SENTINEL (OURS)					2706 VANILLA	
		2707 PR. (%) \uparrow	2708 RE. (%) \uparrow	2709 F1 (%) \uparrow	2710 MED. SPEED \downarrow	2711 AVG. SPEED \downarrow	2712 VAL. LOSS \downarrow	2713 VAL. LOSS \downarrow
2707 NANOGPT-0.25B	C4	91.2	86.1	88.6	5.0	44.07	3.747	3.650
	FW	91.2	86.1	88.6	5.0	62.13	3.752	3.731
	OW	73.3	91.7	81.5	5.0	320.43	3.571	3.531
2712 LLAMA-3-0.6B	C4	76.0	86.4	80.9	5.0	263.51	3.357	3.347
	FW	88.6	88.6	88.6	5.0	108.73	3.465	3.459
	OW	78.0	88.6	83.0	5.0	44.06	3.316	3.313

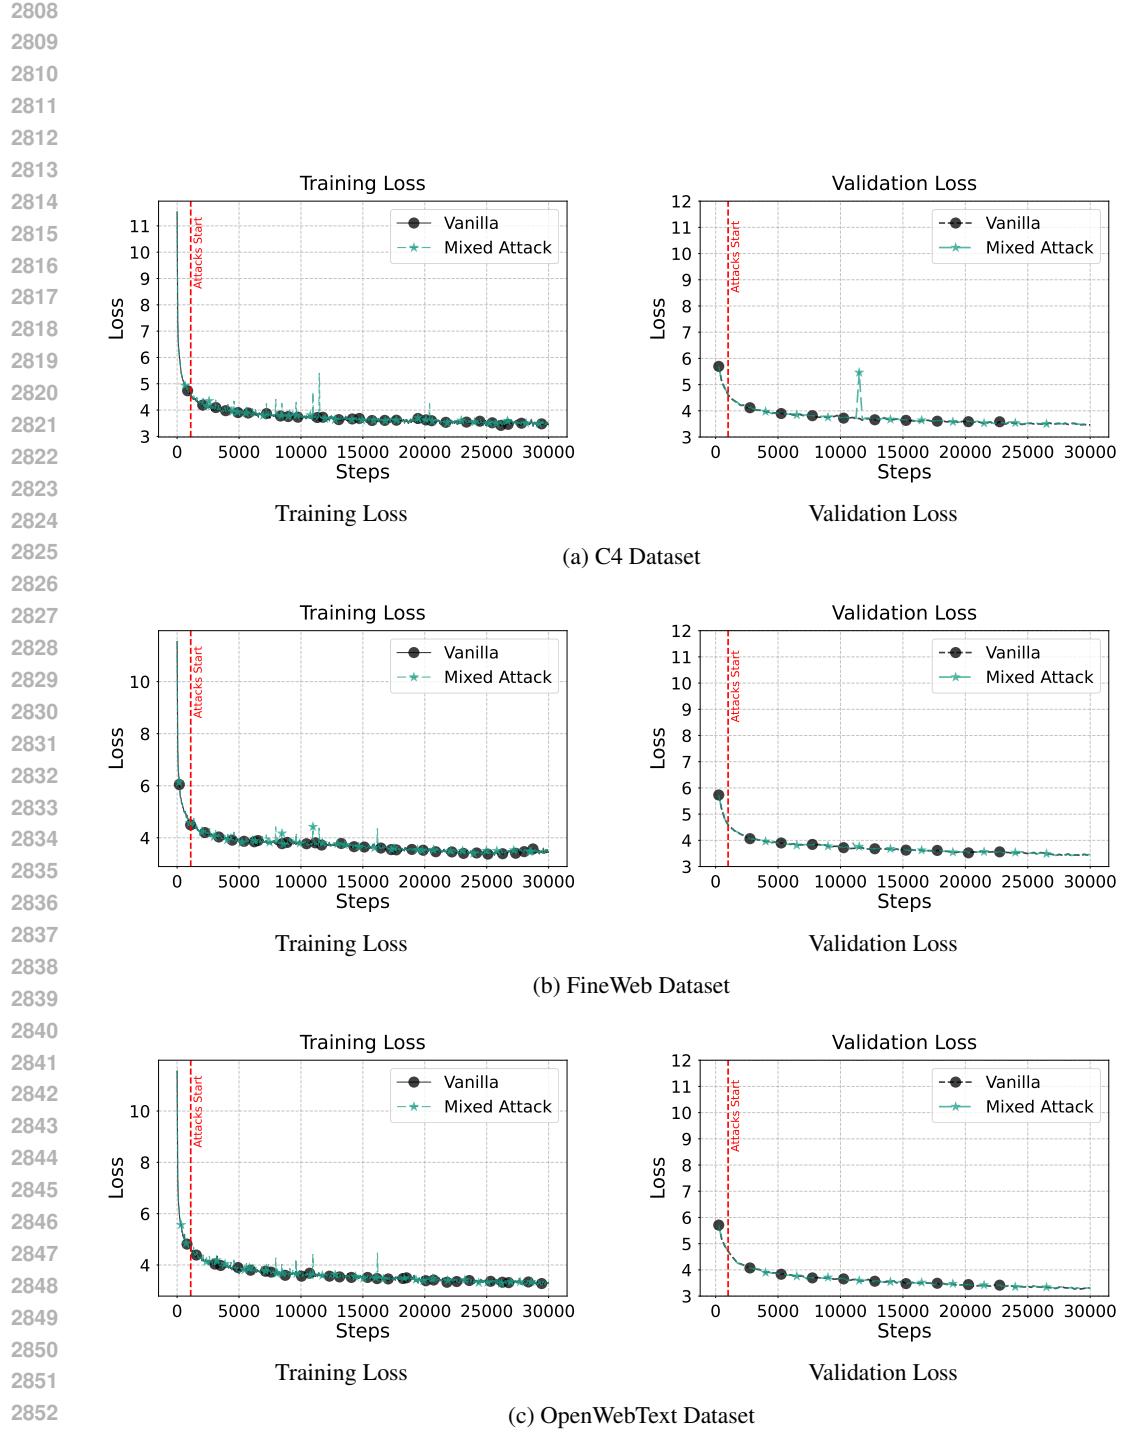
2720
 2721 **Complementary Nature of DP Defense Methods to SENTINEL.** Throughout the paper, we
 2722 explained in detail how the defense methods in the data parallel domain proposed by prior
 2723 Byzantine-tolerant literature (e.g., please see (Mhamdi et al., 2018; Gorbunov et al., 2022; Malinovsky et al.,
 2724 2024)) are complementary to SENTINEL that secures the pipeline parallel dimension by verifying the
 2725 activations and activation gradients transmitted during forward and backward pass. To demonstrate
 2726 this complementary nature, we combine our method with two well-known robust aggregation tech-
 2727 niques used to defend against parameter gradient attacks. In particular, we use Krum (Blanchard et al.,
 2728 2017) and Bulyan (Mhamdi et al., 2018) instead of vanilla gradient averaging. Our goal is to show that
 2729 these methods would not impact the operation of SENTINEL in a negative way as they are operating
 2730 in an orthogonal dimension. Fig. 14 depicts the training and validation loss for a Llama-3-0.6B model
 2731 trained against mixed attacks. As seen, **presence of robust aggregation methods does not interfere**
 2732 **with how SENTINEL works**, and our method can ensure a convergence rate that follows the vanilla,
 2733 non-attacked baseline.



2750 Figure 14: Training and validation loss evolution for Llama-3-0.6B trained on FineWeb dataset over 5k iterations
 2751 **in the presence of robust aggregation methods during gradient averaging.** The mixed attack scenario
 2752 assumes 37.5% Byzantine nodes at each pipeline stage, with 33% randomly selected nodes performing a
 2753 randomly chosen attack at a randomly sampled iteration. Each node has a different manipulation method (as
 2754 outlined in Sec. 2.1).



2800 Figure 15: Training and validation loss evolution for **NanoGPT-0.25B** model on C4, FineWeb, and OpenWebText
2801 datasets over 30k iterations. The mixed attack scenario assumes 50% Byzantine nodes at each pipeline stage,
2802 with one randomly selected node performing a randomly chosen attack at a randomly sampled iteration. Each
2803 node has a different mode (activation vs. gradient) and manipulation method (as outlined in Sec. 2.1).
2804
2805
2806
2807



2853
2854 Figure 16: Training and validation loss evolution for **Llama-3-0.6B** model on C4, FineWeb, and OpenWebText
2855 datasets over 30k iterations. The mixed attack scenario assumes 50% Byzantine nodes at each pipeline stage,
2856 with one randomly selected node performing a randomly chosen attack at a randomly sampled iteration. Each
2857 node has a different mode (activation vs. gradient) and manipulation method (as outlined in Sec. 2.1).
2858
2859
2860
2861

2862 F.3 ADDITIONAL ABLATION STUDIES
28632864 In this section, we present additional ablation studies on the impact of various components used in
2865 SENTINEL.2866
2867 **Impact of EMA in Verification.** To demonstrate how the temporal training dynamics captured
2868 within the EMA affects the detection of malicious workers, we compare our approach against a
2869 naïve version that simply compares submitted signals with the average. For this experiment, we start
2870 training a Llama-3-0.6B model while malicious workers employ activation delay attacks. To cancel
2871 out the EMA, we set $\beta = 0$ which essentially means that we would compare against the instantaneous
2872 activation average. Our results are shown in Tab. 17. As seen, comparison with the instantaneous mean
2873 is not enough to protect against malicious workers. This corroborates the importance of temporal
2874 patterns in detecting pipeline parallel attacks. **Additionally, we measure sensitivity of SENTINEL to**
2875 **the EMA decay rate by modifying $\beta_h = \beta_g$ to 0.6 and 0.99. As shown in Tab. 17, the sensitivity to**
2876 **the decay rates is not significant.**

2877 Table 17: Ablation study on the impact of EMA in SENTINEL.

REFERENCE POINT	METRICS			
	PR. (%) ↑	RE. (%) ↑	F1 (%) ↑	VAL. LOSS ↓
AVERAGE ($\beta_h = \beta_g = 0$)	23.08	100.0	37.5	6.248
SENTINEL ($\beta_h = 0.9$, $\beta_g = 0.8$)	100.0	100.0	100.0	3.826
SENTINEL ($\beta_h = \beta_g = 0.60$)	94.7	100.0	97.3	3.894
SENTINEL ($\beta_h = \beta_g = 0.99$)	90.0	100.0	94.7	3.875

2885
2886 **Impact of Distance Metrics.** The sensitivity of different attacks to various distance metrics varies
2887 between activation and activation gradient attacks, which is why we require multiple distance metrics.
2888 As discussed in Sec. 3.1, other optimal distance metric choices could provide a unified solution (e.g.,
2889 neural network classifiers), which we defer to future work. Here, to evaluate the impact of various
2890 distance metrics, we conducted an ablation study using the mixed attack setting from Tab. 2 for
2891 training a Llama-3-0.6B on the C4 dataset, employing only one distance metric at a time for detection
2892 against mixture attacks. Our results are shown in Tab. 18. As seen, combining all metrics yields
2893 optimal performance against the mixture of all activation and activation gradient attacks.
2894

2895 Table 18: Ablation study on distance metrics against mixed attacks.

DISTANCE METRIC	METRICS			
	PR. (%) ↑	RE. (%) ↑	F1 (%) ↑	VAL. LOSS ↓
SFR	42.9	83.3	56.6	8.882
SWD	40.0	94.4	56.2	6.332
NORMALIZED L_2	75.0	75.0	75.0	10.274
ABSOLUTE DEVIATION L_1	71.1	88.9	79.0	3.883
ALL (SENTINEL)	83.7	92.3	87.8	3.831

2903
2904 **Impact of Random Seeds.** Due to limited computational resources and substantial experimental
2905 costs, we have not reported error bars throughout the paper. To demonstrate the statistical integrity
2906 of our approach, we computed error bars for two of the most challenging activation attacks: delay
2907 (100-steps) and invisible noise (99%). We randomly selected 5 seeds and repeated the experiments
2908 from Tab. 1, randomizing both network initialization and malicious worker selection. The results
2909 provided in Tab. 19 indicate the statistical significance of our findings.
29102911 G SENTINEL IN THE WILD: VERIFICATION FOR DECENTRALIZED LLM
2912 TRAINING USING SWARM PARALLELISM
29132914 In this section, we detail the adaptation of SENTINEL to SWARM parallelism (Ryabinin et al.,
2915 2023) and showcase its capabilities for decentralized training of LLMs. We first provide a high-
level overview of SWARM’s operational dynamics. Then, we demonstrate the compatibility of our

2916 Table 19: Statistical significance analysis with error bars for the performance of SENTINEL against activation
 2917 delay and invisible noise attack.

ATTACK	METRICS			
	PR. (%) \uparrow	RE. (%) \uparrow	F1 (%) \uparrow	VAL. LOSS \downarrow
DELAY (100-steps)	99.2 \pm 1.6	100.0 \pm 0.0	99.6 \pm 0.8	3.843 \pm 0.008
INVISIBLE NOISE (99%)	100.0 \pm 0.0	100.0 \pm 0.0	100.0 \pm 0.0	3.832 \pm 0.004

2924 verification mechanism with communication-efficient compression techniques employed in distributed
 2925 SWARM training. We describe how SENTINEL integrates with SWARM’s existing infrastructure by
 2926 leveraging its trainer node architecture. Additionally, we analyze the critical role of verification in
 2927 the presence of SWARM’s stochastic wiring mechanism, which introduces additional failure modes
 2928 beyond traditional PP. Finally, we present comprehensive experimental settings and detailed results
 2929 from our SWARM experiments.

2932 G.1 SWARM PARALLELISM OVERVIEW

2934 As briefly discussed in Sec. 5.1, SWARM parallelism can be viewed as a stochastic DP/PP training
 2935 approach. At a high level, each worker gets assigned a random layer/stage of the model to process (PP-
 2936 axis) while other workers process different micro-batches for the same stage (DP-axis). Unlike traditional
 2937 fixed meshes used in distributed training frameworks such as `torch.distributed` (Paszke et al., 2019; Li et al., 2020), SWARM parallelism operates as pools of workers available at each stage
 2938 that process data for subsequent stages.

2940 Coordination of training in such a stochastic environment is handled by so-called trainer nodes. Each
 2941 trainer node is responsible for processing a single micro-batch of data end-to-end. In particular, each
 2942 trainer node receives a disjoint micro-batch of data and sends it to a worker at stage 0. Once the
 2943 worker processes the data, the trainer receives the result and forwards it to the next stage. The trainer
 2944 continues this process until the micro-batch has passed through all layers in the forward pass. Then,
 2945 the trainer begins the backward pass by traversing the stages in reverse order. With this architecture,
 2946 workers do not need to save the intermediate activations for backward since the trainer maintains all
 2947 necessary information and can send it to different workers during the backward pass. The workers
 2948 then accumulate parameter gradients locally and update their parameters after an all-reduce operation
 2949 with all existing workers of the same stage.

2950 Trainer nodes use a Distributed Hash Table (DHT) to route their micro-batches. Since devices at
 2951 each stage process data at different speeds, trainers maintain worker load and throughput information
 2952 in the DHT so that all other trainers know when to send signals and which worker at each stage
 2953 processes data most efficiently. This mechanism is at the heart of the *stochastic wiring* that occurs
 2954 within SWARM and distinguishes it from fixed mesh DP/PP approaches (Ryabinin et al., 2023). As
 2955 demonstrated, a SWARM system can accommodate many trainer nodes since each processes a single
 2956 micro-batch, and typically multiple devices at each stage are capable of processing data batches. For
 2957 a visual representation of SWARM, please refer to Fig. 17a. A pseudo-code of how the trainer nodes
 2958 work in SWARM is also provided in Alg. 6.

2959 G.2 SUBSPACE COMPRESSION

2961 Regular SWARM parallelism enables decentralized training, but suffers from slow speeds due to
 2962 bandwidth requirements (Ryabinin et al., 2023). To relax this requirement and enable faster training,
 2963 recent work by Ramasinghe et al. (2025) has shown that training with speeds as low as 60 Mbps is
 2964 plausible using lossless compression. In particular, Ramasinghe et al. (2025) observes that most pre-
 2965 trained transformers already exhibit low-rank properties in their feedforward and attention projection
 2966 layers. They utilize this observation to construct a lossless compression mechanism using a shared
 2967 subspace that enables faster transfer of activations and gradients passed between layers in PP.

2968 In simple terms, given a signal $\mathbf{h}^{(s)} = f_s(\mathbf{h}^{(s-1)})$ that has been processed by a worker at stage s ,
 2969 instead of communicating $\mathbf{h}^{(s)} \in \mathbb{R}^{b \times n \times d}$, they compress it using a subspace compression matrix

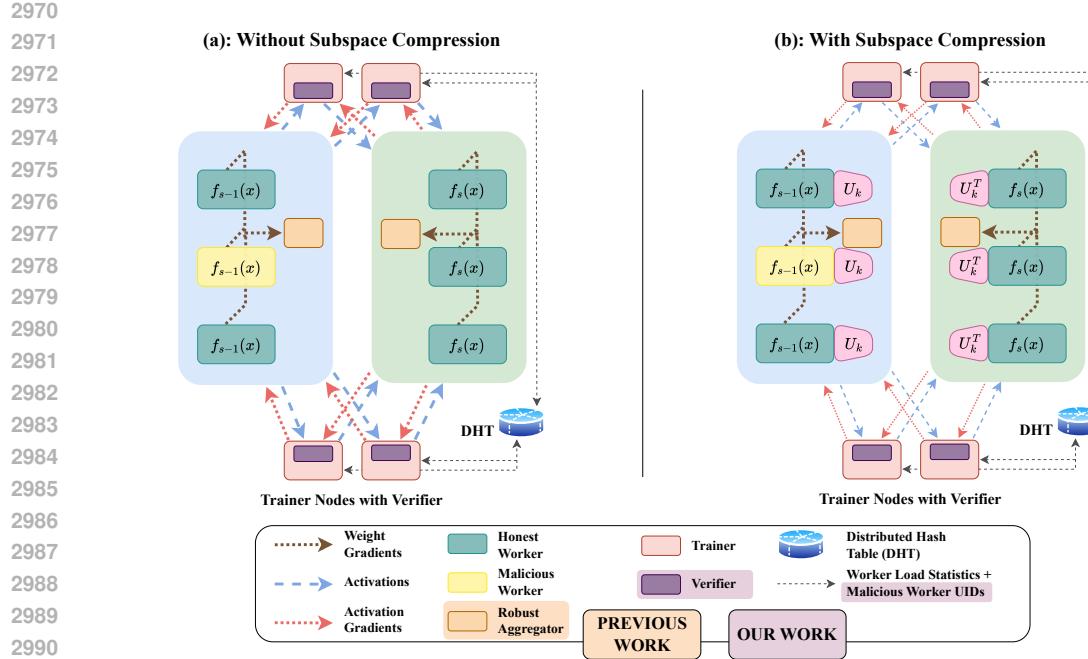


Figure 17: (a) SWARM (Ryabinin et al., 2023) parallelism utilizes pipeline and data parallelism to train neural networks. Compared to traditional DP/PP methods, SWARM utilizes trainer nodes to stochastically route micro-batches of input data through the model layers. Trainer nodes leverage a Distributed Hash Table (DHT) to communicate worker load/throughput to enable a better device utilization within the SWARM. (b) Subspace compression (Ramasinghe et al., 2025) complements SWARM by adding lossless compression at the boundary of worker devices for faster communication of activations and activation gradients from/to trainer nodes. SENTINEL can be adapted to both settings and enable worker verification. It sits naturally within the trainer nodes that are already coordinating signal transmission among workers.

U_k that projects it to a k -dimensional subspace where $k \ll d$ (Ramasinghe et al., 2025):

$$h_{\text{compressed}}^{(s)} = (h^{(s)} - \text{PE} - \mathbf{T}_{\text{fixed}}[t_{1:n}, :])U_k. \quad (68)$$

Here, PE denotes the positional embeddings and $\mathbf{T}_{\text{fixed}}[t_{1:n}, :]$ are the fixed token embeddings. At layer $s + 1$, the receiver decompresses this using (Ramasinghe et al., 2025):

$$h_{\text{recovered}}^{(s)} = h_{\text{compressed}}^{(s)}U_k^T + \text{PE} + \mathbf{T}_{\text{fixed}}[t_{1:n}, :] = h^{(s)} \quad (69)$$

This compression scheme, tailored for training transformer-based LLMs, can enable training across four geographical regions with bandwidth as low as 60 Mbps. Therefore, we will incorporate it as part of our realistic SWARM integration. For an overview of SWARM with subspace compression, please refer to Fig. 17b.

G.2.1 SENTINEL COMPATIBILITY WITH SUBSPACE COMPRESSION

To demonstrate the compatibility of our proposed method with subspace compression, we integrate this compression algorithm into our fixed mesh `TorchTitan` implementation. We assume that malicious workers would manipulate the compressed signals since these are what is being sent to subsequent workers.

In particular, we use the first k dimensions of the transmitted signal $h_{\text{compressed}}^{(s)} \in \mathbb{R}^{b \times n \times k+1}$, as the last dimension corresponds to the fixed token embeddings which may vary significantly between different batches. Specifically, SENTINEL uses $h_{\text{compressed}}^{(s)}[:, :, :, -1]$ for activation verification using EMA (and similarly $g_{\text{compressed}}^{(s)}[:, :, :, -1]$ for gradients). The rest of the algorithm remains unchanged.

3024 **Algorithm 6** SWARM Trainer Node (Ryabinin et al., 2023)

```

3025
3026 class TrainerNode:
3027     def __init__(self, dht, stage_uids):
3028         self.dht = dht
3029         self.stages = [RemoteExpert(uid, dht) for uid in stage_uids]
3030
3031     def forward(self, input_batch):
3032         """Process microbatch through all pipeline stages"""
3033         hidden = input_batch
3034         for stage in self.stages:
3035             hidden = stage.process(hidden)
3036         return hidden
3037
3038 class RemoteExpert:
3039     def __init__(self, stage_uid, dht):
3040         self.stage_uid = stage_uid
3041         self.dht = dht
3042
3043     def process(self, data):
3044         """Route data to available worker at this stage"""
3045         worker = self.dht.find_available_worker(self.stage_uid)
3046         result = worker.forward(data)
3047         return result

```

3048 Using these assumptions, we train our 16-layer Llama-3-0.6B on the FineWeb-EDU dataset using the
3049 settings given in Tab. 10. The only difference is that for these experiments we use a local batch-size
3050 8 with 8 gradient accumulation steps to reach a target batch-size of 512 per optimizer step. For
3051 subspace compression, we use a compression factor of 25 for 96% compression. As demonstrated
3052 by our results in Tab. 20, SENTINEL can easily be adapted to this setting and demonstrates efficient
3053 detection capability against various kinds of attacks. We also plot the validation loss for all these
3054 attacks in Fig. 18, demonstrating uninterrupted training in all cases.

3055 Table 20: Attack detection performance across different attack modes for training Llama-3-0.6B with subspace
3056 compression (Ramasinhe et al., 2025). Metrics shown include precision, recall, F1 score (all as percentages),
3057 and validation loss at 5000 steps. In all scenarios, each stage has 3:5 malicious to honest ratio.

3058 MODE	3059 ATTACK	3060 DETECTION PERFORMANCE			3061 TRAINING
		3062 PR. (%) \uparrow	3063 RE. (%) \uparrow	3064 F1 (%) \uparrow	
3065 ACTIVATION	3066 Constant (Zeros)	100.00	100.00	100.00	4.2490
	3067 Constant (Ones)	100.00	100.00	100.00	4.2486
	3068 Random Value	100.00	100.00	100.00	4.2999
	3069 Bias Addition (Constant)	100.00	75.00	85.71	4.2484
	3070 Bias Addition (Random)	0.00	0.00	0.00	4.3248
	3071 Delay (100-steps)	100.00	100.00	100.00	4.2677
3072 GRADIENT	3073 Constant (Zeros)	100.00	100.00	100.00	4.2470
	3074 Constant (Ones)	100.00	100.00	100.00	4.2639
	3075 Random Value	100.00	100.00	100.00	4.2726
	3076 Bias Addition (Constant)	80.00	100.00	88.89	4.2751
	3077 Bias Addition (Random)	100.00	100.00	100.00	4.2608
	3078 Delay (100-steps)	100.00	100.00	100.00	4.2567
3079 MIXED MODE		96.97	88.89	92.75	4.2779

3080 G.3 SENTINEL INTEGRATION WITH SWARM

3081 Now that we have laid out the background on SWARM and how subspace compression acts as a
3082 complementary factor that enables decentralized training under restricted bandwidths, let us detail how
3083 SENTINEL can be integrated into this realistic, production-ready decentralized training ecosystem.

3078

3079

3080

3081

3082

3083

3084

3085

3086

3087

3088

3089

3090

3091

3092

3093

3094

3095

3096

3097

3098

3099

3100

3101

3102

3103

3104

3105

3106

3107

3108

3109

3110

3111

3112

3113

3114

3115

3116

3117

3118

3119

3120

3121

3122

3123

3124

3125

3126

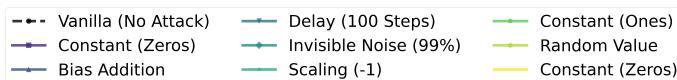
3127

3128

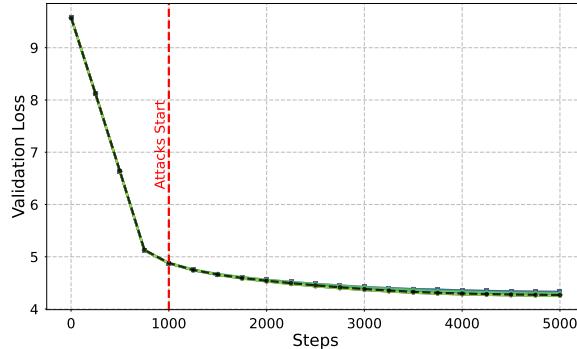
3129

3130

3131

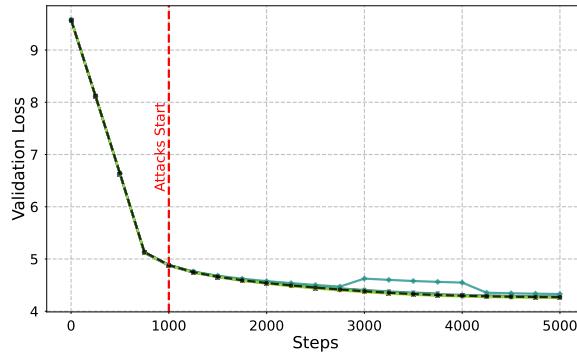


Validation Loss - Activation Manipulation Methods



(a) Activation Attacks

Validation Loss - Gradient Manipulation Methods



(b) Gradient Attacks

Figure 18: Validation loss evolution for training Llama-3-0.6B models that employ subspace compression (Ramasinha et al., 2025) under various activation/gradient manipulation attacks. Despite the compression in the activation/gradients, SENTINEL can successfully prevent divergence from vanilla training loss.

As discussed, trainer nodes are in a natural position to take on the verifier role. As mentioned in App. A Q3, trainer nodes only perform CPU-based operations and are readily available at a fraction of worker node costs. Thus, it is not far-fetched to operate such nodes for training coordination. Additionally, data integrity heavily depends on the trustworthiness of trainer nodes, and if this role were delegated to untrusted parties, other major issues regarding backdoor and privacy attacks would arise. Therefore, it is both economically feasible and logically crucial to have trusted trainer nodes.

To add SENTINEL verification to trainer nodes, we employ a mechanism for trainers to maintain an EMA of the signals they distribute across different stages. In particular, each trainer stores the EMA of all layer outputs since it operates on them end-to-end. Specifically, trainer i maintains:

$$\{\mathbf{m}_{t,i}^{(s)} \mid 1 \leq s \leq p\} \quad (70)$$

where

$$\mathbf{m}_{t,i}^{(s)}(\mathbf{h}) = \beta_h \mathbf{m}_{t-1,i}^{(s)}(\mathbf{h}) + (1 - \beta_h) \mathbf{h}_{t,i}^{(s,r)}, \quad (71)$$

is the momentum at step t and stage s . A similar set is also kept for gradient EMAs, but we omit it for brevity. Using these EMA states, each trainer can run SENTINEL verification as signals are being processed by the workers.

3132 **Algorithm 7** SWARM Trainer Node with SENTINEL Verification

```

3133
3134 class SENTINELTrainerNode:
3135     def __init__(self, dht, stage_uids):
3136         self.dht = dht
3137         self.stages = [RemoteExpert(uid, dht) for uid in stage_uids]
3138         self.ema_detector = EMADetector()
3139         self.banned_workers = set()
3140
3141     def forward(self, input_batch):
3142         """Process microbatch with Sentinel verification"""
3143         hidden = input_batch
3144         for i, stage in enumerate(self.stages):
3145             # Process data at current stage
3146             hidden = stage.process(hidden)
3147
3148             # Sentinel verification
3149             is_suspicious = self.ema_detector.update_and_detect(i, hidden)
3150
3151             if is_suspicious:
3152                 worker_uid = stage.get_last_worker_uid()
3153                 self._reportViolation(worker_uid)
3154
3155             return hidden
3156
3157     def _reportViolation(self, worker_uid):
3158         """Report violation to DHT for global coordination"""
3159         violations = self.dht.get(f"violations_{worker_uid}", default=0)
3160         self.dht.store(f"violations_{worker_uid}", violations + 1)
3161
3162         if violations + 1 > MAX_VIOLATIONS:
3163             self.dht.store(f"banned_{worker_uid}", True)
3164             self.banned_workers.add(worker_uid)

```

3161
3162 The challenging aspect of implementing SENTINEL in SWARM is communicating malicious behavior
3163 among trainers. Since each trainer is responsible for maintaining a separate EMA and verifying
3164 their randomly chosen workers independently, this might increase the malicious impact of bad actors.
3165 To address this issue, we utilize the DHT for lightweight communication between trainers about
3166 malicious workers. Instead of each trainer independently tracking violation counters or worker bans,
3167 they cooperate through the DHT to increment violation counters or ban workers collectively. In simple
3168 terms, we track the number of violations for each worker through their unique identifiers (UIDs) in
3169 the DHT, and if they surpass the allowed number of violations, the first trainer that observes this
3170 bans them. Similarly, we periodically check whether workers are demonstrating good behavior after
3171 transient violations and decrease their violation counts through our forgiveness strategy discussed in
3172 SENTINEL. The only slight modification that makes SENTINEL work better in SWARM is replacing
3173 tainted gradients with zero tensors, which we observed makes training more stable.

3174 A pseudo-code of SENTINEL integration with SWARM is given in Alg. 7 (backward pass is omitted
3175 for brevity).

3176 **G.3.1 IMPORTANCE OF VERIFICATION UNDER STOCHASTIC WIRING**

3177
3178 An interesting observation that we made through implementing SENTINEL in SWARM is the im-
3179 portance of highly calibrated detection thresholds in removing abundant false positives due to
3180 the interplay between “cascading effect” in PP and “stochastic wiring” in SWARM. **As discussed**
3181 **in Ryabinin et al. (2023), stochastic wiring helps the SWARM to move towards maximum device**
3182 **utilization and less idle time. However, this behavior can have an adverse effect considering adversar-**
3183 **ial actors.** Specifically, if a trainer fails to detect/flag a bad actor, it not only updates its own EMA
3184 with corrupted signals, but other trainers would also be in danger since they would eventually use
3185 that malicious worker while routing their micro-batches due to SWARM’s stochastic wiring. Thus, if
a malicious worker goes undetected, it can corrupt the EMA of all trainers and they could flag all

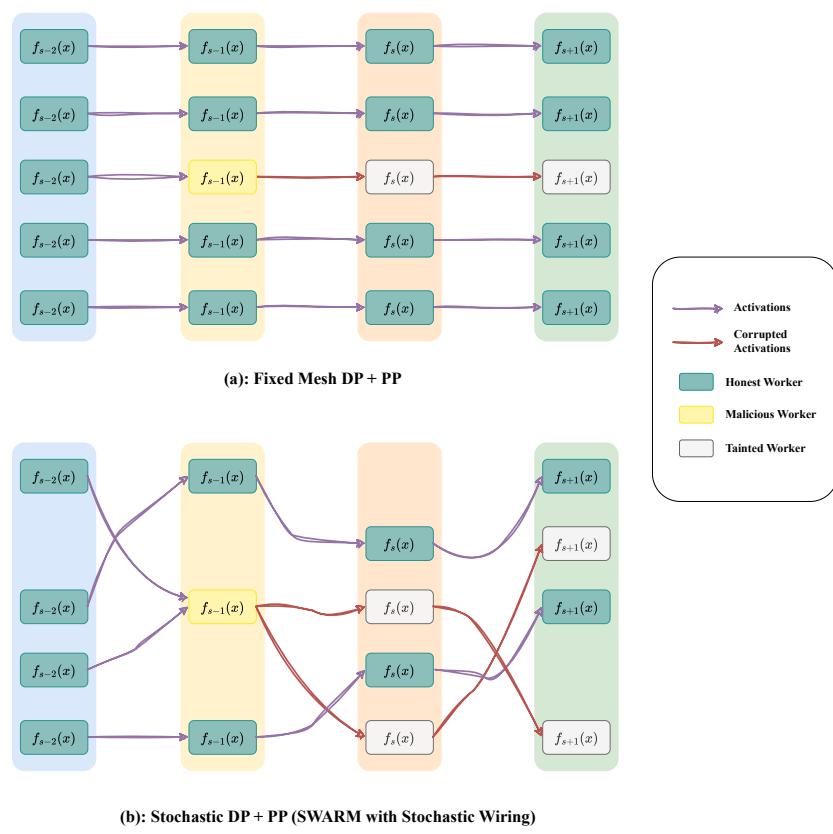


Figure 19: Visual comparison of (a) fixed mesh DP + PP vs. (b) SWARM (Ryabinin et al., 2023). Unlike fixed mesh DP + PP where each worker sends their data to a specific worker in the next stage, SWARM utilizes stochastic wiring to route the data batches between stages. This is to ensure that workers would reach their maximum utilization and have minimum idle time. This behavior, however, is a double-edge sword as by stochastically routing data batches between workers, we are effectively increasing the reachability of malicious workers that can pollute “any” worker in the subsequent stage if they go undetected. This highlights the importance of highly calibrated verification mechanisms in the real-world. Note that the signals transmitted between workers in SWARM all pass through trainer nodes, and hence, there is no peer-to-peer communication between workers. We just directly connected the workers for a better visual comparison with fixed mesh.

honest workers in that stage as malicious. In the fixed mesh (shown in Fig. 19a), in the worst case we would wrongly flag all workers from a single pipe due to the “cascading effect” that we discussed in App. D.2. In SWARM, however, we are in danger of falsely accusing all workers of the same stage due to “stochastic wiring” which when considered together with “cascading”, could mean a high false positive rate for the detection algorithm. We defer further investigation into this interesting inter-play to future work. **This phenomenon has been elaborated in Fig. 19b.**

G.4 DETAILED EXPERIMENTAL RESULTS

To evaluate our SENTINEL integration with SWARM under realistic distributed training conditions including subspace compression, we conduct comprehensive experiments that test both the detection capabilities and the system’s robustness under various attack scenarios. Our experimental setup is designed to validate the effectiveness of the proposed method under realistic conditions.

G.4.1 EXPERIMENTAL SETTINGS

Distributed Architecture. We use a Llama-3-0.6B model with 16 transformer layers, partitioned into 16 PP stages for SWARM. Each stage employs 8 parallel workers to process micro-batches concurrently, resulting in a total of 128 worker nodes. To simulate realistic distributed environments,

3240
 3241 Table 21: Detection performance of SENTINEL in a distributed SWARM with 128 workers. There are 37.5%
 3242 malicious workers that are submitting randomly chosen mixed gradient and activation attacks. Note that the low
 3243 recall is attributed to some nodes employing weak attacks that are not disruptive to training, hence they do not
 3244 get flagged as malicious but training continues without disruption. This is in line with our observation in Fig. 1
 3245 and intuition from Theorem 1 that weak attackers may survive detection but they are no harm to the training.
 3246

ATTACK MODE	METRICS		
	PR. (%) ↑	RE. (%) ↑	F1 (%) ↑
MIXED GRADIENT	100.0	75.0	85.7
MIXED ACTIVATION/GRADIENT	100.0	66.7	80.0

3250
 3251 each worker is deployed on a separate AWS instance with no direct interconnection to other nodes.
 3252 We use heterogeneous instance types based on computational requirements: `g5.2xlarge` instances
 3253 with NVIDIA A10G GPUs for the embedding and first transformer layer, `g5.4xlarge` instances
 3254 with NVIDIA A10G GPUs for the final transformer layer and projection head, and `g4dn.2xlarge`
 3255 instances with NVIDIA T4 GPUs for all intermediate transformer layers.

3256 The trainer infrastructure consists of 32 parallel trainers distributed across 4 `c6a.8xlarge` CPU
 3257 instances, with 8 trainers running simultaneously per instance. Each trainer processes a disjoint data
 3258 shard streamed from FineWeb-EDU (Penedo et al., 2024) via HuggingFace (Lhoest et al., 2021),
 3259 ensuring no data overlap during training. This configuration allows us to test the scalability of our
 3260 DHT-based violation reporting system under realistic deployment conditions.

3261 **Attack Scenarios.** We evaluate SENTINEL’s detection performance under two distinct scenarios:
 3262

1. **Gradient-Only Attacks:** A heterogeneous mixture comprising constant attacks (zeros and ones),
 3264 bias addition, random value injection, and scaling attacks.
2. **Combined Gradient and Activation Attacks:** An expanded attack suite including a mixture
 3266 of gradient and activation attacks such as constant (zeros and ones), bias addition, random
 3267 value, scale, delay, and random sign attacks. This more sophisticated threat model evaluates
 3268 SENTINEL’s comprehensive detection capabilities across the entire forward and backward pass.

3269
 3270 For both scenarios, we simulate a challenging environment with a 3:5 malicious-to-honest worker ratio
 3271 (37.5% malicious workers). In the gradient-only setting, we assume attackers operate independently
 3272 without coordination, initiating attacks at random intervals. For the combined attack scenario, we
 3273 model partial coordination where 15% of malicious workers launch synchronized attacks at random
 3274 intervals, simulating coordinated adversarial behavior while maintaining realistic assumptions about
 3275 attacker capabilities.

3276 **Training Configuration.** We use micro-batches of size 4 per worker with a total target batch-size
 3277 of 512 per optimizer step. For subspace compression, we use a compression factor of 25. All other
 3278 hyper-parameters and training settings are the same as provided in Tab. 10. Given the costs associated
 3279 with running 128 SWARM workers, we chose to train for 2500 steps only. Nevertheless, we made
 3280 sure to squeeze all the attack start times within those steps.

3281
 3282 **G.4.2 RESULTS**

3283 In this section, we present our results. From Fig. 20, we can see that in the absence of a viable
 3284 verification mechanism while training across an untrusted distributed environment, training can easily
 3285 be disrupted. In Tab. 21 we also present SENTINEL’s detection performance where we achieve greater
 3286 than 80% F1-score. Note that lower recall in mixed activation/gradient attacks is due to some nodes
 3287 employing weak attacks that are not disruptive to training, hence they do not get flagged as malicious
 3288 but training continues without disruption. This is in line with our observation in Fig. 1 and our
 3289 intuition from Theorem 1 that weak attackers may survive detection but do not harm the training.
 3290 This can be seen in Fig. 20 that training continues without divergence in both cases. These results
 3291 prove the versatility of SENTINEL in real-world applications involving distributed training.

3292
 3293 **EMA Variance across Trainers.** When employing SENTINEL in SWARM, we discussed how
 each trainer accumulate their own version of EMAs. Thus, these signals can vary from one trainer

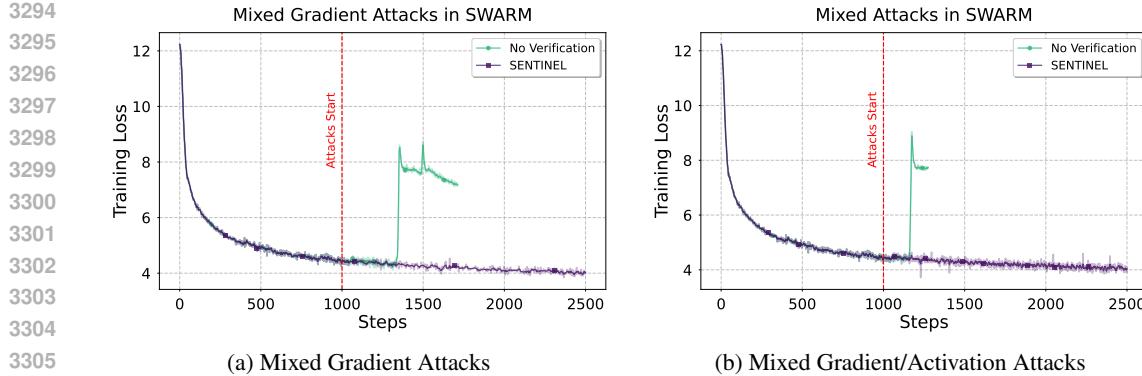


Figure 20: Loss when training Llama-3-0.6B models with subspace compression (Ramasinghe et al., 2025) in a distributed SWARM (Ryabinin et al., 2023) of 128 workers. Workers employ various activation/gradient manipulation attacks to disrupt training. While in the absence of verification training gets disrupted, SENTINEL can successfully protect training from divergence by detecting and banning malicious workers.

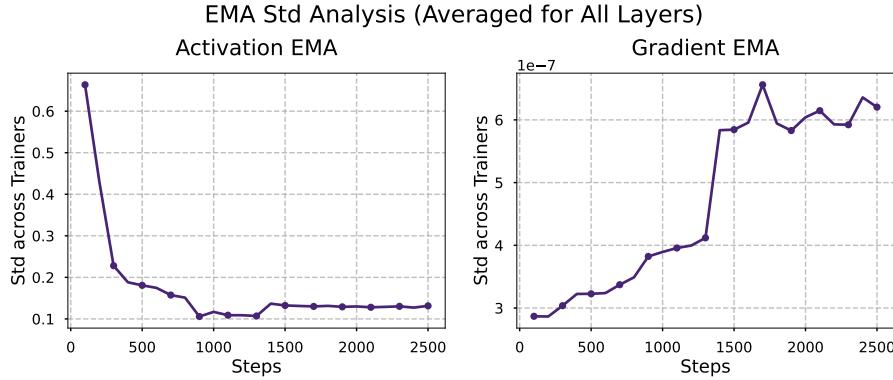


Figure 21: Standard deviation of EMA signals (activation/gradient) between the 32 trainers used in training our Llama-3-0.6B model. Since each trainer keeps the EMAs of all layers, here we report the average across all layers. As seen, activation EMA variance decreases as training progresses while gradient EMA variance increases among trainer. Despite this increase, note that the scale of gradient EMA variance is very close to zero (10^{-7}).

to another. If this variance is large, it may cause verification disruption: a worker that appears honest to one trainer could be flagged as malicious solely due to EMA variance between trainers. To demonstrate that EMAs among trainers have minimum divergence, we plot the standard deviation of the activation and activation gradient EMA among all our 32 trainers used for training in SWARM. As seen in Fig. 21, both activation and activation gradient display a controlled amount of variance across trainers which is a testament to the fact that EMAs in different trainers evolve similarly.

Test Statistic Evolution. Finally, we examine how SENTINEL is utilized by disjoint trainer nodes in SWARM. To this end, we track the test statistics for a particular worker (worker number 7 at layer 12) when it processes micro-batches from different trainers. We also track the lower and upper thresholds determined using our adaptive IQR mechanism from Alg. 5. Fig. 22 shows the evolution of these test statistics over time. From this figure, we conclude that despite no direct EMA or threshold communication between trainer nodes, all 32 trainers exhibit similar evolutionary patterns. This consistency is crucial because each trainer routes its micro-batches through different workers, requiring their detection criteria to remain aligned. Without this alignment, the system could suffer from false positives or false negatives that would disrupt training stability. The key benefit of using EMA within SENTINEL is its ability to efficiently track historical patterns, enabling this coordinated behavior across distributed trainers.

3348
 3349
 3350
 3351
 3352
 3353
 3354
 3355
 3356
 3357
 3358
 3359
 3360
 3361
 3362
 3363
 3364
 3365
 3366
 3367
 3368
 3369
 3370
 3371
 3372
 3373
 3374
 3375
 3376
 3377
 3378
 3379
 3380
 3381
 3382
 3383
 3384
 3385
 3386
 3387
 3388
 3389
 3390
 3391
 3392
 3393
 3394
 3395
 3396
 3397
 3398
 3399
 3400
 3401

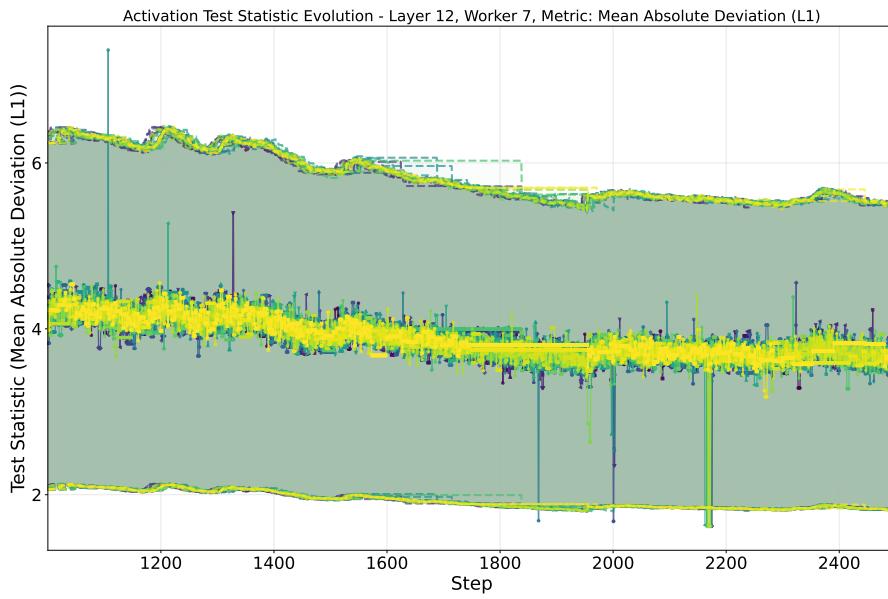


Figure 22: Evolution of adaptive deviation bounds and worker statistics for a worker processing layer 12 in our 16 layer Llama-3-0.6B. Each color represents the test statistics recorded by a different trainer. The upper and lower thresholds determined by our adaptive IQR mechanism are also shown as the top and bottom lines for each trainer. Despite no direct EMA or threshold communication between the trainers, they usually have a similar threshold. The worker starts their attack around step 2150 after which the trainers flag and ban the worker.