

OASIS: ONLINE SAMPLE SELECTION FOR CONTINUAL INSTRUCTION TUNING

Minjae Lee^{1,*} Minhyuk Seo^{2,*} Tingyu Qu² Tinne Tuytelaars² Jonghyun Choi^{1,†}

¹ Seoul National University ² KU Leuven
 {mj1020, jonghyunchoi}@snu.ac.kr
 {minhyuk.seo, tingyu.qu, tinne.tuytelaars}@kuleuven.be

ABSTRACT

In continual instruction tuning (CIT) scenarios, where new instruction tuning data continuously arrive in an online streaming manner, training delays from large-scale data significantly hinder real-time adaptation. Data selection can mitigate this overhead, but existing strategies often rely on pre-trained reference models, which are impractical in CIT setups since future data are unknown. Recent reference model-free online sample selection methods address this, but typically select a fixed number of samples per batch (*e.g.*, top- k), making them vulnerable to distribution shifts where informativeness varies across batches. To address these limitations, we propose OASIS, an adaptive online sample selection approach for CIT that (1) selects informative samples by estimating each sample’s informativeness relative to all previously seen data, beyond batch-level constraints, and (2) minimizes informative redundancy of selected samples through iterative selection score updates. Experiments on various large foundation models show that OASIS, using only 25% of the data, achieves comparable performance to full-data training and outperforms the state-of-the-art sampling methods.

1 INTRODUCTION

A key factor in the success of recent large foundation models (LFMs), including LLMs and multi-modal LLMs (MLLMs), is training on large-scale instruction tuning data (Bai et al., 2023b; Chen et al., 2023). While emerging high-quality datasets enable LFMs to better adapt to user preferences and contexts (Maharana et al., 2025; Lau et al., 2024; Zhang, 2024), scaling such datasets also amplifies risks of overfitting and delays in training time (Ma et al., 2023; Liu et al., 2024b).

Unlike conventional training paradigms, where LFMs are optimized on static datasets prior to deployment, real-world applications often require adaptation to continuously arriving data streams (Seo et al., 2024b), *e.g.*, for learning new concepts. This motivates continual instruction tuning (CIT) (Chen et al., 2024; Maharana et al., 2025), where models adapt to shifting data distributions. However, training with the massive influx of continuously arriving data in CIT often causes forgetting of prior knowledge (Zhai et al., 2024; Zhu et al., 2024b), overfitting (He et al., 2025; Ghosh et al., 2024), and significant training delays, which hinder the model’s ability to learn from new data promptly (*i.e.*, real-time adaptation) (Koh et al., 2022; Caccia et al., 2022). For example, as shown in Fig. 1, full-data training delays the adaptation to newer tasks (T_3 and T_4) as the model spends excessive time training on earlier tasks (T_1 and T_2).

Data selection methods (Sorscher et al., 2022; Lee et al., 2024; Abbas et al., 2024) mitigate training delays by selecting informative subsets, but most rely on a reference model trained on the entire data in advance (Mindermann et al., 2022; Shin et al., 2023). Training a reference model is infeasible in CIT, where new data arrive continuously and future data distributions remain unknown. While reference model-free methods (Qin et al., 2024; Hong et al., 2024) eliminate such reliance, their strategy of selecting a fixed number of samples (*e.g.*, top- k) from each streaming batch fails to capture inter-batch and intra-batch informativeness.

*indicates equal contribution. † indicates corresponding authors. JC is with ECE, IPAI, ASRI at SNU.




Figure 1: **Real-time adaptation under equal training time.** Bar width indicates training data volume in the CIT data stream. While 'Full' trains on all data, TIVE (Liu et al., 2024b), Adapt- ∞ (Maharana et al., 2025), and OASIS use 25% selected data. Under equal training time, 'Full' degrades on newly arrived tasks (e.g., T_3 , T_4), since sequential training on all data provides sufficient time for earlier tasks (e.g., T_1 , T_2) but insufficient time for new ones. TIVE and Adapt- ∞ achieve only marginal speedup despite using 25% data, as the backward-pass selection overhead limits real-time adaptation. OASIS uses inference-only selection with minimal overhead, enabling efficient sample selection and strong adaptation to new tasks.

Specifically, since informative samples are unevenly distributed over time (Seo et al., 2025), e.g., some batches contain many novel or forgotten samples while others contain few, fixed-size selection per batch often includes uninformative samples while missing critical ones, thereby overlooking inter-batch informativeness variation. Moreover, in continual data streams, similar instances often recur periodically (Koh et al., 2023), e.g., Christmas data in winter, swimsuit data in summer. As a result, many samples within a batch receive similar selection scores (Hong et al., 2024), thus selecting the top- k without considering intra-batch similarity leads to redundancy in the chosen subset.

To address these limitations, we propose **Online Adaptive Sample selection via Informative Statistics (OASIS)**, which (i) selects informative samples by estimating each sample’s informativeness relative to all previously encountered data, beyond batch-level constraints, and (ii) reduces the redundancy by considering sample-wise similarity within each batch. Specifically, OASIS maintains online estimates of the global mean and variance of informativeness as batches arrive, and uses these statistics to assess the relative informativeness of each sample. Moreover, to avoid selecting similar samples within a batch (Hong et al., 2024), OASIS iteratively updates the samples’ selection scores: once a sample is chosen, the scores of remaining candidate samples in the batch are adjusted according to their shared information with the selected sample. Note that OASIS requires only a forward pass for sample selection, roughly half the cost of a backward pass (Huo et al., 2018), enabling efficient selection and reducing training delays, thereby facilitating real-time adaptation to newly encountered tasks, as shown in Fig. 1.

We empirically validate OASIS by comparing it with recent sample selection baselines across various LLMs (e.g., Llama-3.1-8B (Grattafiori et al., 2024) and Qwen3-8B (Yang et al., 2025)) and MLLMs (e.g., LLaVA-1.5-7B (Liu et al., 2024a) and Qwen-VL-2.5-7B (Bai et al., 2025)) on multiple CIT benchmarks. In particular, OASIS incurs only a 1.51% performance degradation compared to full-data training, while training with only 25% of samples on MICVIT benchmark.

We summarize our contributions as follows:

- We propose an adaptive sample selection strategy that selects batch samples based on their informativeness relative to the entire dataset, moving beyond batch-wise selection.
- We propose a redundancy reduction strategy that leverages sample-wise similarity to minimize redundancy among selected samples.
- By combining these two strategies, our proposed OASIS significantly outperforms baselines in CIT through extensive evaluations.

2 RELATED WORK

Continual Instruction Tuning. Existing instruction tuning methods often focus on fixed tasks (Zhu et al., 2024a; Bai et al., 2025), overlooking continuously emerging instruction tuning data (Maharana et al., 2025; Guo et al., 2025). To adapt LLMs to dynamically changing data distributions, *continual instruction tuning* (CIT) has been proposed, which aims to learn new tasks while preserving knowledge from previously encountered instruction tuning data (He et al., 2023). To reflect

real-world distribution shifts, various CIT benchmarks have been proposed in both text-only (*e.g.* Long Sequence (Razdaibiedina et al., 2023), TRACE (Wang et al., 2023)) and multi-modal (*e.g.* COAST (Cao et al., 2024), UCIT (Guo et al., 2025)) domains, along with corresponding strategies (*e.g.* Fwd-Prompt (Zheng et al., 2024), SEFE (Chen et al., 2025)). However, they train LFM on all data, leading to overfitting (Rice et al., 2020; Zhai et al., 2024), high computational costs (Wang et al., 2024; Panos et al., 2025), and training delays (Caccia et al., 2022; Ghunaim et al., 2023). Although aL-SAR (Seo et al., 2025) improves the training efficiency of LFM through dynamic layer freezing based on batch informativeness, it still relies on the entire dataset, limiting its overall efficiency.

Data Selection. Motivated by the observation that not all data contribute equally to learning, recent work explores selecting informative subsets to match full-data performance with lower training cost (Lee et al., 2024; Abbas et al., 2024; Qin et al., 2024). Bayesian (Deng et al., 2023b) and RHO-LOSS (Mindermann et al., 2022) enhance training efficiency by training only with a selected subset of data, but both require a reference model pretrained on the full dataset for data selection. TIVE (Liu et al., 2024b) uses smaller reference data, but requires full-layer gradient computation, incurring high computational overhead that offsets any computation savings from data selection. This reliance on reference models limits their use in CIT, where sequentially arriving data makes full-data pretraining infeasible. To address this, reference model-free online sample selection methods like GradNorm (Katharopoulos & Fleuret, 2018), InfoBatch (Qin et al., 2024), and DivBS (Hong et al., 2024) have been proposed. However, they select a fixed number of samples per batch based on difficulty or dissimilarity, limiting adaptability to shifting data distributions, where some batches may contain more informative samples (*e.g.*, newly encountered or forgotten) than others. SelfSup (Sorscher et al., 2022), COINCIDE (Lee et al., 2024), and DBP (Abbas et al., 2024) select samples by K -means clustering with a sensitive hyperparameter K , which is difficult to tune under non-i.i.d. streams due to representation shifts (Książek et al., 2025). Recently, Adapt- ∞ (Maharana et al., 2025) tackles data selection under shifting data distributions, but it assumes known task boundaries and relies on costly intermediate gradients, limiting its real-world applicability.

3 APPROACH

We first present the problem statement for online sample selection in CIT in Sec.3.1. We then introduce our method, **Online Adaptive Sample selection via Informative Statistics (OASIS)**, in Sec. 3.2, which consists of two components: **Online Relative Informativeness Selection (ORIS)** in Sec.3.3 and **Similarity-aware Information Redundancy Elimination (SIREN)** in Sec. 3.4.

3.1 PROBLEM STATEMENT OF ONLINE SAMPLE SELECTION IN CIT

CIT trains an LFM on a data stream \mathcal{D} comprising a sequence of T tasks, *i.e.*, $\mathcal{D} = \mathcal{D}_1, \dots, \mathcal{D}_T$, where each $\mathcal{D}_i = \{(x_1^i, y_1^i), (x_2^i, y_2^i), \dots\}$ denotes the dataset for task i . Note that explicit task boundaries may be absent, and the data distribution across tasks can even be identical. Reflecting real-world scenarios where data are collected continuously over time (Koh et al., 2022), online CIT assumes that data arrive as an online stream of samples, denoted as $(x_1^i, y_1^i), (x_2^i, y_2^i), \dots$, whereas *offline* CIT provides each task chunk \mathcal{D}_i at once. At timestep t , a batch \mathcal{B}_t with batch size N_B is drawn from \mathcal{D} . A subset $\mathcal{B}_t^* \subset \mathcal{B}_t$ containing N_s samples ($N_s < N_B$) is then selected according to a predefined selection ratio, and only this selected subset \mathcal{B}_t^* is used for training. Given model parameters θ of an LFM f and loss function \mathcal{L} , the objective is:

$$\min_{\theta} \mathbb{E}_{\mathcal{B}_t \sim \mathcal{D}} [\mathcal{L}(f_{\theta}(\mathcal{B}_{t,x}^*), \mathcal{B}_{t,y}^*)]. \quad (1)$$

3.2 ONLINE ADAPTIVE SAMPLE SELECTION VIA INFORMATIVE STATISTICS

Existing online sample selection methods typically select a fixed number of samples per batch, *e.g.*, top- N_S selection scores, which struggles under distribution shifts (*e.g.*, CIT), due to: (i) varying numbers of informative samples per batch over time, and (ii) redundancy among co-occurring similar samples within a batch (Hong et al., 2024).

To address these limitations, we propose two strategies: (i) **Online Relative Informativeness Selection (ORIS)**, which selects samples based on their relative informativeness across the entire data distribution, rather than intra-batch ranking; and (ii) **Similarity-aware Information Redundancy Elimination**




Figure 2: **Overview of our proposed OASIS.** For each online batch \mathcal{B}_t : (1) OASIS first scores the informativeness I for all sample in batch $\mathcal{B}_t = \{d_1^{(t)}, d_2^{(t)}, \dots\}$ (Eq. 2); (2) It then iteratively reduces redundancy by adjusting I of other samples to \tilde{I} based on their gradient similarity $S_{i,j}$ to the most informative sample (here, $d_2^{(t)}$). (Sec. 3.4); (3) OASIS computes relative informativeness \hat{I} by normalizing the updated informativeness \tilde{I} using EMA μ_t and EMV σ_t (Eq. 4); (4) Finally, OASIS computes selection probability P_S and selects samples exceeding a uniformly drawn threshold r , resulting in a selected subset $\mathcal{B}_t^* \subset \mathcal{B}_t$ (Eq. 6). Model f_θ is then trained using only \mathcal{B}_t^* .

(SIREN), which mitigates redundancy by updating, at each selection step, the selection scores of remaining samples within a batch according to their similarity with already selected ones.

Specifically, for each training batch, ORIS computes sample-wise informativeness using Fisher Information (FI) and estimates relative informativeness within the overall FI distribution. During this calculation, SIREN adjusts these scores using sample-wise gradient similarity to account for the impact of training one sample on others. Integrating ORIS and SIREN, we call our method Online Adaptive Sample selection via Informative Statistics (OASIS). We provide an overview of OASIS in Fig. 2, and pseudocode in Sec. A.14.

3.3 ONLINE RELATIVE INFORMATIVENESS SELECTION

Informativeness I . To select informative samples, we first define the sample-wise informativeness I based on Fisher Information (FI), which measures a model’s potential information gain from each example (Deng et al., 2023a). Computing full FI is costly, since it requires forward and backward passes across all layers, which would negate the efficiency gains of training on only a selected subset. We therefore adopt an efficient approximation: (i) computing gradients only at the last layer θ_L , since gradients of preceding layers are proportional to it by the chain rule (Koh et al., 2023), and (ii) using a first-order diagonal FI approximation that avoids Hessian computation (Kirkpatrick et al., 2017; Soen & Sun, 2021). Formally, we define the informativeness of the i th sample $d_i^{(t)} = (x_i^{(t)}, y_i^{(t)})$ in batch \mathcal{B}_t as:

$$I_i^{(t)} = \text{tr} \left[\left(\nabla_{\theta_L} L(d_i^{(t)}) \right) \cdot \left(\nabla_{\theta_L} L(d_i^{(t)}) \right)^\top \right], \quad (2)$$

where $L(d_i^{(t)}) = \ell(f_\theta(x_i^{(t)}), y_i^{(t)})$, ℓ is the loss function, and $\text{tr}(\cdot)$ denotes the trace operator. While existing methods use gradients from middle (Maharana et al., 2025) or all layers (Liu et al., 2024b) for selection, incurring costs close to full-data training, ORIS requires only last-layer gradients, yielding near inference-only cost.

Relative Informativeness \hat{I} . However, we cannot compare I across batches, as θ in Eq. 2 changes over time due to model updates at each batch iteration. In other words, cross-batch comparison of I would require re-forwarding all previously encountered samples with the latest θ , which is computationally prohibitive and often infeasible in online CIT scenarios where previous samples may be discarded (Maharana et al., 2025). To avoid these issues, instead of tracking all samples’ I , we maintain a running mean of I over seen samples and evaluate each new sample by its deviation from this mean, yielding a relative informativeness over all previously observed data.

To be specific, inspired by the exponential decay model of forgetting (Shin & Lee, 2020; Mahto et al., 2020; Chien et al., 2021), we maintain exponential moving average (EMA) and variance (EMV) of I , which exponentially decay the influence of previous statistics. Formally, at each timestep t , given

batch \mathcal{B}_t of size $N_{\mathcal{B}}$ and decay factor β , we update EMA μ_t and EMV σ_t of I as:

$$\begin{aligned}\mu_t &= \beta \bar{I}^{(t)} + (1 - \beta)\mu_{t-1}, \\ \sigma_t &= \beta (\bar{I}^{(t)} - \mu_{t-1})^2 + (1 - \beta)\sigma_{t-1}.\end{aligned}\tag{3}$$

where $\bar{I}^{(t)} = \frac{1}{N_{\mathcal{B}}} \sum_{i=1}^{N_{\mathcal{B}}} I_i^{(t)}$ denotes average $I^{(t)}$ of samples in \mathcal{B}_t .

Using these statistics, we estimate each sample’s relative informativeness. By Theorem 3.1, normalization with the EMA μ_{t-1} and EMV σ_{t-1} (updated up to step $t - 1$) yields a standard normal distribution, enabling deviation measurement of each sample via Z -score normalization (see Sec. A.1 for proof). Formally, we define relative informativeness of the i_{th} sample with informativeness $I_i^{(t)}$ in batch \mathcal{B}_t at time step t as:

$$\hat{I}_i^{(t)} = \frac{I_i^{(t)} - \mu_{t-1}}{\sigma_{t-1}}.\tag{4}$$

Note that we normalize with μ_{t-1} and σ_{t-1} rather than μ_t and σ_t , since the latter already incorporate $I_i^{(t)}$ and thus induce dependence between numerator and denominator (*i.e.*, self-normalization bias (Jing et al., 2003; de la Peña et al., 2009)).

Theorem 3.1. *Let $I_i^{(t)}$ be defined by Eq. 2 for samples in \mathcal{B}_t , and let $(\mu_{t-1}, \sigma_{t-1}^2)$ be the EMA and EMV from Eq. 3 computed from batches prior to t . Assume local stationarity and local weak dependence of $\{I_i^{(t)}\}$ with uniformly bounded second moments. Then*

$$\hat{I}_i^{(t)} = \frac{I_i^{(t)} - \mu_{t-1}}{\sigma_{t-1}} \approx \mathcal{N}(0, 1).\tag{5}$$

Probabilistic Selection. We then determine whether to select each sample based on $\hat{I}^{(t)}$. Inspired by coresets methods (Bang et al., 2021; Seo et al., 2024a), which select both informative and easy-to-learn samples, we adopt probabilistic selection based on \hat{I} rather than deterministically selecting those above a fixed threshold I_T (set by the target selection ratio). Specifically, we apply Sigmoid to $\hat{I} - I_T$, and perform Bernoulli sampling, allowing samples below the threshold to be included with nonzero probability. Formally, with $r \sim \mathcal{U}(0, 1)$, where $\mathcal{U}(0, 1)$ denotes the uniform distribution over the interval $[0, 1]$, we define the selected samples set \mathcal{B}_t^* from batch \mathcal{B}_t as:

$$\mathcal{B}_t^* = \{d_i^{(t)} \in \mathcal{B}_t \mid \text{sigmoid}(\hat{I}_i^{(t)} - I_T) > r\}.\tag{6}$$

Note that existing online selection methods require knowledge of the total sample count to fix per-batch quotas. In contrast, ORIS uses only the selection ratio to determine I_T , enabling operation on endless real-world data streams. We provide details on determining I_T in Sec. A.6.

3.4 SIMILARITY-AWARE INFORMATION REDUNDANCY ELIMINATION

In real-world data streams, similar task samples often arrive consecutively. Selecting samples within \mathcal{B}_t using ORIS without considering intra-batch similarity can thus lead to redundancy. To mitigate this, we adjust I based on sample-wise similarity. Ideally, one would iteratively select the most informative unchosen sample, train the model, and then recompute I for the remaining candidates using the updated model, thereby accounting for the reduced informativeness of samples similar to those already trained on (Hekimoglu et al., 2023). However, this procedure incurs prohibitive computational cost due to repeated I computation, rendering batch-wise training impractical.

Instead of re-forwarding to recompute I , we approximate its change induced by selecting other samples. Inspired by Du et al. (2018); Seo et al. (2025), which show that gradient alignment correlates with informational overlap, when a sample is selected, we multiply its I by its gradient similarity with each remaining sample and subtract the result from their I values, thereby reducing the informativeness of samples that overlap with the selected one. Formally, given a batch $\mathcal{B}_t = \{d_i^{(t)}\}_{i=1}^{N_{\mathcal{B}}}$, we first calculate I for each sample and sort them in descending order, *i.e.*, $I_1^{(t)} > \dots > I_{N_{\mathcal{B}}}^{(t)}$. We then add the most informative sample (*i.e.*, d_1) to the set H , the set of samples that is assumed to

have been trained on. For each remaining sample $d_i^{(t)} \in \mathcal{B}_t \setminus H$, we define updated informativeness $\tilde{I}_i^{(t)}$ as:

$$\tilde{I}_i^{(t)} = I_i^{(t)} - \sum_{h \in H} \cos(g_i, g_h) \cdot I_h, \quad (7)$$

where $g_i = \nabla \ell_{\theta_L}(f(x_i), y_i)$ refers to the last-layer gradient of sample $d_i = (x_i, y_i)$. We iteratively add $d_1, \dots, d_{N_{\mathcal{B}}}$ to H and repeat the process iteratively, as shown in Algorithm 1.

However, when $|H| > 1$, overlapping similarities between samples in H cause duplicate subtractions. We resolve this using the inclusion-exclusion principle to capture higher-order redundancy as:

$$\begin{aligned} \tilde{I}_i^{(t)} = & I_i^{(t)} - \sum_{h \in H} \cos(g_i, g_h) \cdot I_h \\ & + \sum_{\substack{U \subseteq H \\ |U| \geq 2}} (-1)^{|U|} \cos(g_i, \bar{g}_U) \cdot \bar{I}_U, \end{aligned} \quad (8)$$

where U denotes all possible non-empty subsets of H with $|U| \geq 2$, $\bar{g}_U = \frac{1}{|U|} \sum_{u \in U} g_u$ and $\bar{I}_U = \frac{1}{|U|} \sum_{u \in U} I_u$ denote the average gradient and average I over subset $U \subseteq H$, respectively.

We use $\tilde{I}^{(t)}$ to calculate relative informativeness $\hat{I}^{(t)}$ instead of directly using $I^{(t)}$, as shown in Fig. 2. By applying SIREN before computing $\hat{I}^{(t)}$ in ORIS, we effectively reduce redundancy without requiring any re-forwarding process.

4 EXPERIMENTS

4.1 EXPERIMENTAL SETUP

In all experiments, we report the mean and standard deviation of the results from three different seeds. Moreover, we conduct Welch’s t-test with a significance level of 0.05. The highest performance is highlighted in bold, and results not significantly different from the best are underlined.

Models. For MLLMs, we use LLaVA-1.5 (Liu et al., 2024a) and Qwen-VL-2.5 (Bai et al., 2025) as our MLLMs. In the main paper, we focus on LLaVA-1.5-7B and Qwen-VL-2.5-7B, while we provide experiments with other model sizes of LLaVAs (1B, 3B, and 13B) and Qwen-VLs (0.5B). During training, we update only the LoRA adapters (Hu et al., 2022), keeping the LLM frozen for training efficiency (Ye et al., 2023). For LLMs, we use Llama-3.1-8B (Grattafiori et al., 2024) and Qwen3-8B (Yang et al., 2025).

Baselines. We compare OASIS with recent state-of-the-art sample selection methods, including GradNorm (Katharopoulos & Fleuret, 2018), Self-Sup (Sorscher et al., 2022), COINCIDE (Lee et al., 2024), DBP (Abbas et al., 2024), InfoBatch (Qin et al., 2024), DivBS (Hong et al., 2024), TIVE (Liu et al., 2024b), Adapt- ∞ (Maharana et al., 2025), and Random selection, which, despite its simplicity, often surpasses SOTA selection methods as noted in Gupta et al. (2023); Zheng et al. (2023); Maharana et al. (2025). While baselines select a fixed number of samples per batch, OASIS dynamically selects samples probabilistically. For fair comparisons, we ensure OASIS uses a comparable or fewer total samples (see Sec. A.13. See Sec. A.4 for the baselines’ details).

Metrics. We report A_{last} , the accuracy measured at the end of training, and A_{avg} , the average accuracy measured at each task boundary.

Benchmarks. We evaluate on a range of CIT benchmarks, including text-only benchmarks (Long Sequence (Razdaibiedina et al., 2023), TRACE (Wang et al., 2023)) and multi-modal benchmarks (COAST (Cao et al., 2024), Adapt (Maharana et al., 2025)). However, COAST assumes balanced task sizes (*i.e.*, 20k samples per task), which contradicts real-world data imbalance, while Adapt includes datasets containing COCO (Lin et al., 2014) images that overlap with LLaVA’s instruction-tuning data. To address these limitations, we introduce **Multi-image Imbalanced Continual Visual Instruction Tuning (MICVIT)**, a new benchmark that removes such overlaps while preserving naturally imbalanced task distributions. To better reflect real-world complexity, MICVIT primarily adopts multi-image datasets.

Implementation Details. We fine-tune models using learning rates from their original papers for one epoch: $2e-5$ for both LLaVA-1.5-7B and Qwen-VL-2.5-7B and $3e-4$ for both LLaMA-3.1-8B and Qwen3-8B. We use the Adam optimizer with no decay and a Cosine LR scheduler. For batch size N_B , we use 16 for all experiments. For LoRA finetuning, we adopt LoRA with rank 128 at all linear layers in the LLM backbone. We set the number of batch iterations per sample encounter to COAST, Adapt, and MICVIT as 0.125, 0.125, and 0.0625, respectively. We assume an infinite memory setup, which assumes that all encountered samples can be stored in an episodic memory. For the batch retrieval method, we adopt memory-only training (Koh et al., 2022; Seo et al., 2024b), where training batches are only retrieved from the episodic memory at each iteration, enhancing robustness to distribution shifts (Koh et al., 2023; Seo et al., 2025). We set the EMA ratio β to 0.9 for all datasets. We conduct experiments on NVIDIA RTX A6000 GPUs. Every experiment takes no more than 2 days.

More details on benchmarks can be found in A.3, respectively.

4.2 QUANTITATIVE ANALYSIS

Comparison across Various Sample Selection Ratios. We compare OASIS with the baselines by training LLaVA-1.5-7B on MICVIT under various selection ratios (*i.e.*, 6.25%, 12.5%, and 25.0%). As shown in Tab. 2, OASIS significantly outperforms baselines, with the largest gap at the lowest sample selection ratio (*i.e.*, 6.25%). Since each sample has a greater impact at smaller coreset sizes (Zheng et al., 2023; Jafari et al., 2024), this highlights OASIS’s strength in identifying informative data. Moreover, at the 25.0% ratio, OASIS nearly matches full-data training, showing only a 1–2% drop in A_{last} . Finally, gains in both A_{last} and A_{avg} indicate that OASIS not only boosts final accuracy but also accelerates convergence by prioritizing informative samples.

Methods using K-means clustering (*i.e.*, Self-Sup, COINCIDE, DBP, TIVE, and Adapt- ∞) often perform worse than random selection due to imbalanced clustering under imbalanced data distributions, as detailed in Sec. A.12.

Comparison of Computational Cost. We compare the computational cost of OASIS and baselines, as shown in Fig. 3. OASIS achieves higher performance while requiring fewer FLOPs. See Sec. A.11 for a detailed computational analysis.

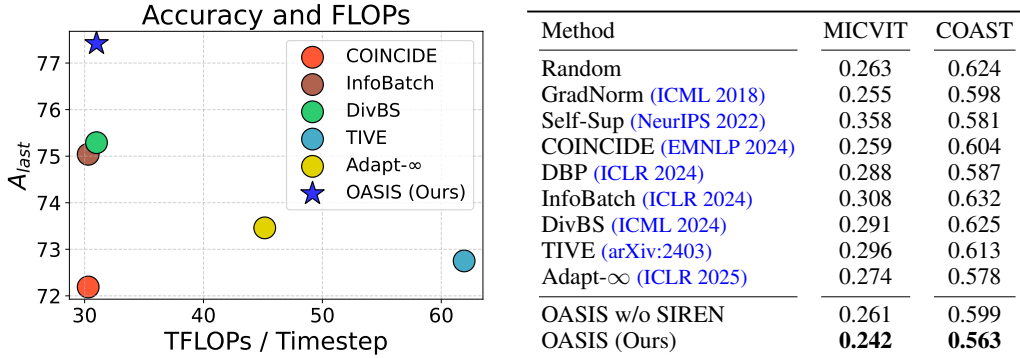


Figure 3: **Accuracy and FLOPs with 25% selection ratio on MICVIT.** The top-left corner illustrates effective and efficient sample selection.

Comparison across Various Benchmarks. We also compare OASIS with baselines using a 6.25% selection ratio on existing CIT benchmarks, namely COAST, Adapt, Long Sequence, and TRACE (Tab. 3). As shown in the table, OASIS consistently outperforms the baselines across multiple CIT benchmarks, demonstrating its robustness and generalizability. We provide additional results with various selection ratios in Sec. A.2

Comparison with Qwen-VL. Beyond LLaVA-1.5-7B, we further test OASIS on Qwen-VL-2.5-7B. As shown in Tab. 4, OASIS consistently outperforms baselines, demonstrating its model-agnostic generalizability. See Sec. A.5 for results on different benchmarks across different selection ratios on Qwen-VL.

Method	Selection Ratio (%)					
	6.25		12.5		25.0	
	$A_{avg} \uparrow$	$A_{last} \uparrow$	$A_{avg} \uparrow$	$A_{last} \uparrow$	$A_{avg} \uparrow$	$A_{last} \uparrow$
Full-Data Training	71.80±0.44	79.66±0.43	71.80±0.44	79.66±0.43	71.80±0.44	79.66±0.43
Random	61.15±0.34	67.29±0.61	64.50±0.22	71.33±0.47	65.48±0.73	73.84±0.45
GradNorm (ICML 2018)	61.23±0.83	66.33±0.73	63.68±1.39	70.51±0.10	65.81±1.37	72.03±0.94
Self-Sup (NeurIPS 2022)	58.05±0.77	64.61±1.16	62.84±0.73	69.48±0.62	64.39±1.26	71.81±0.13
COINCIDE (EMNLP 2024)	59.28±1.45	64.83±0.45	62.33±0.92	69.50±0.39	64.92±0.81	72.25±1.17
DBP (ICLR 2024)	57.14±0.51	63.24±0.82	60.64±1.06	67.85±1.19	62.14±1.42	70.13±1.39
InfoBatch (ICLR 2024)	60.82±0.75	68.88±1.09	64.87±0.68	73.63±0.28	66.41±1.36	74.70±0.52
DivBS (ICML 2024)	61.07±0.58	69.06±0.54	65.31±0.41	<u>73.76±0.85</u>	64.22±0.71	<u>75.38±1.25</u>
TIVE (arXiv:2403)	58.75±1.20	64.19±1.40	61.27±0.59	68.38±0.41	64.29±1.02	72.76±0.83
Adapt-∞ (ICLR 2025)	59.13±0.32	64.37±0.85	61.53±0.84	69.82±0.99	65.50±0.52	73.57±1.22
OASIS (Ours)	64.39±0.58	71.76±0.72	66.89±0.68	75.60±0.26	68.84±0.37	77.95±0.93

Table 2: **Quantitative comparison between sample selection methods on MICVIT benchmark.** LLaVA-1.5-7B is used as the MLLM. Bold indicates the highest performance; underlined results are within the 0.05 t-test significance level. 'Full-Data Training' uses all data without selection.

Method	MLLM Benchmark				LLM Benchmark			
	COAST		Adapt		Long Sequence		TRACE	
	$A_{avg} \uparrow$	$A_{last} \uparrow$	$A_{avg} \uparrow$	$A_{last} \uparrow$	$A_{avg} \uparrow$	$A_{last} \uparrow$	$A_{avg} \uparrow$	$A_{last} \uparrow$
Full-Data Training	31.56±1.42	39.06±0.55	55.89±0.20	54.26±0.23	85.83±0.48	82.83±0.47	50.84±0.70	57.73±0.54
Random	23.57±0.17	30.80±0.30	47.64±0.21	40.26±0.71	70.90±0.13	68.02±0.18	43.44±0.55	52.27±0.03
GradNorm	21.91±0.47	28.08±0.49	44.92±0.06	36.78±0.52	70.70±0.05	68.36±0.75	41.39±0.86	50.29±0.69
Self-Sup	22.16±0.27	30.09±0.51	43.85±0.75	34.49±1.37	67.88±0.28	64.24±0.43	37.82±0.39	47.23±0.44
COINCIDE	22.56±0.90	29.89±1.42	44.56±0.84	35.83±0.48	69.02±0.59	66.82±0.69	39.14±0.85	49.42±0.82
DBP	20.28±0.81	28.67±0.32	43.25±0.45	34.75±0.54	66.74±0.82	64.95±0.06	35.93±0.26	45.86±0.98
InfoBatch	22.93±0.73	29.14±0.56	47.65±0.30	40.60±0.96	<u>74.38±0.37</u>	69.84±0.24	41.26±0.27	50.19±0.13
DivBS	23.41±0.14	31.72±0.18	47.25±0.74	40.38±0.62	<u>74.37±0.75</u>	<u>70.65±0.57</u>	44.25±0.08	<u>52.23±0.76</u>
TIVE	21.15±0.09	28.12±0.10	44.72±0.58	35.28±1.37	68.15±0.41	65.79±0.90	36.57±0.84	47.04±0.02
Adapt-∞	22.42±0.16	29.39±0.72	45.06±0.34	35.51±1.29	71.48±0.97	68.45±0.31	42.38±0.14	50.48±0.44
OASIS (Ours)	25.67±0.35	34.23±0.38	49.98±0.27	43.94±0.31	75.26±0.35	71.91±0.55	45.56±0.35	53.48±0.17

Table 3: **Quantitative comparison across MLLM and LLM benchmarks at a 6.25% selection ratio.** We employ LLaVA-1.5-7B (MLLM) and LLaMA-3.1-8B (LLM). Bold indicates the highest performance; underlined results are within the 0.05 t-test significance level. 'Full-Data Training' uses all data without selection.

Comparison of Selected Samples' Diversity. We quantify diversity using kernel density, computed as the mean pairwise similarity under a Gaussian kernel, following Lee et al. (2024). As shown in Tab. 1, OASIS selects the most diverse set among all methods. We attribute this to the probabilistic sampling in ORIS, which promotes diversity, and the redundancy reduction in SIREN.

Method	Selection Ratio (%)					
	6.25		12.5		25.0	
	$A_{avg} \uparrow$	$A_{last} \uparrow$	$A_{avg} \uparrow$	$A_{last} \uparrow$	$A_{avg} \uparrow$	$A_{last} \uparrow$
Full-Data Training	72.31±0.42	78.18±0.75	72.31±0.42	78.18±0.75	72.31±0.42	78.18±0.75
Random	59.23±0.85	65.94±0.38	62.18±0.19	69.03±0.10	67.44±0.12	74.12±0.42
GradNorm (ICML 2018)	57.49±0.92	63.38±0.64	63.90±0.22	69.34±0.47	67.91±0.55	73.39±0.30
Self-Sup (NeurIPS 2022)	54.29±0.56	61.42±0.31	61.37±0.99	68.32±0.38	64.34±0.49	70.71±0.56
COINCIDE (EMNLP 2024)	57.34±0.77	62.90±0.48	62.38±0.73	69.23±0.44	65.12±0.38	71.04±0.72
DBP (ICLR 2024)	54.53±0.43	60.38±0.56	60.63±0.30	67.81±0.52	62.42±0.42	72.24±0.61
InfoBatch (ICLR 2024)	58.14±0.28	65.38±0.87	64.04±0.42	68.26±1.31	65.69±0.30	73.51±0.45
DivBS (ICML 2024)	60.28±0.98	67.48±0.29	65.26±1.34	70.48±0.38	66.18±0.74	75.57±0.42
TIVE (arXiv:2403)	55.84±0.49	63.37±0.83	62.11±0.45	67.42±0.73	64.64±0.59	72.16±0.81
Adapt-∞ (ICLR 2025)	56.32±0.97	62.48±0.38	62.35±0.49	67.82±1.24	65.95±0.61	72.18±0.16
OASIS (Ours)	61.64±0.43	69.43±0.72	66.54±0.88	72.35±0.28	70.23±0.27	76.41±0.41

Table 4: **Quantitative comparison with Qwen-VL-2.5-7B on MICVIT benchmark.** We use Qwen-VL-2.5-7B as the MLLM. Bold indicates the highest performance; underlined results are within the 0.05 t-test significance level. 'Full-Data Training' uses all data without selection.




Figure 4: **Comparison of fast adaptation performance.** After CIT of LLaVA-1.5-7B on subsets (25% of the full data), selected using each sample selection baseline from MICVIT, we fine-tune the model for 100 epochs on each downstream task (*i.e.*, VISION, COMICS, MagicBrush, and DreamSim).

Method	MICVIT		COAST		Adapt	
	$A_{avg} \uparrow$	$A_{last} \uparrow$	$A_{avg} \uparrow$	$A_{last} \uparrow$	$A_{avg} \uparrow$	$A_{last} \uparrow$
Vanilla	62.54±0.55	67.16±0.39	21.35±0.38	29.39±0.26	46.36±0.53	38.57±0.38
(+) ORIS	65.55±0.41	72.45±0.30	25.80±0.32	35.28±0.52	50.52±0.42	44.65±0.10
(+) ORIS & SIREN (Ours)	67.58±0.46	74.51±0.48	27.83±0.50	37.29±0.13	51.33±1.16	46.22±0.57

Table 5: **Ablations for proposed components of OASIS.** We train LLaVA-1.5-7B using 6.25% of the dataset samples with 0.125, 0.25, and 0.25 online iterations in MICVIT, COAST, and Adapt, respectively. Vanilla: selecting a fixed number of samples with the highest FI per batch.

Method	Model Scale					
	LLaVA-1.5-1B		LLaVA-1.5-3B		LLaVA-1.5-13B	
	$A_{avg} \uparrow$	$A_{last} \uparrow$	$A_{avg} \uparrow$	$A_{last} \uparrow$	$A_{avg} \uparrow$	$A_{last} \uparrow$
Random	55.04±0.12	59.75±0.30	57.37±0.12	60.39±0.68	67.49±0.47	73.60±0.49
GradNorm (ICML 2018)	53.44±0.05	57.39±0.22	56.19±0.70	60.73±0.08	65.32±0.53	71.50±0.98
Self-Sup (NeurIPS 2022)	52.01±0.33	55.96±0.27	54.59±0.53	56.76±0.50	64.71±0.71	71.62±0.57
COINCIDE (EMNLP 2024)	53.03±0.51	56.94±0.16	56.04±0.40	59.55±0.39	65.47±0.85	72.25±0.86
DBP (ICLR 2024)	51.02±0.23	55.62±0.34	53.32±0.68	55.89±0.39	62.27±0.02	68.37±0.02
InfoBatch (ICLR 2024)	55.48±0.08	60.29±1.09	58.92±1.31	60.39±0.64	66.78±0.45	72.04±0.93
DivBS (ICML 2024)	57.24±0.12	<u>60.22±0.42</u>	<u>58.02±0.72</u>	62.15±0.52	67.28±0.27	73.63±0.18
TIVE (arXiv:2403)	51.34±0.11	56.26±0.08	54.82±0.12	58.90±0.37	64.74±0.02	71.50±0.02
Adapt-∞ (ICLR 2025)	54.09±0.15	57.09±0.21	55.29±0.45	56.81±0.69	63.69±0.68	70.46±0.81
OASIS (Ours)	57.86±0.10	62.17±0.25	60.46±0.26	64.01±0.40	68.04±0.21	74.60±0.22

Table 6: **Quantitative comparison between online sample selection methods across model scales on MICVIT under selection ratio 6.25%.** Bold indicates the highest performance; underlined results are within the 0.05 t-test significance level.

Comparison of Fast Adaptation Performance. We evaluate fast adaptation on unseen downstream tasks following CIT, where CIT serves as upstream continual pre-training. This setup assesses the generalizability of models trained on subsets selected by different online sample selection methods. Downstream tasks include COMICS (Iyyer et al., 2017), VISION (Bai et al., 2023a), DreamSim (Fu et al., 2023), and MagicBrush (Zhang et al., 2023), while MICVIT serve as upstream datasets. Results in Fig. 4 show that subsets selected by OASIS yield consistently superior performance across all downstream tasks, demonstrating superior generalizability. Note that zero-shot performance (*i.e.*, iteration 0) can vary depending on the similarity between the downstream tasks and the subset of data selected by each baseline, resulting in inconsistent ordering between baselines. However, as fine-tuning on downstream tasks progresses, OASIS exhibits consistently faster adaptation, demonstrating the superior generalizability of models trained on subsets selected by OASIS.

Comparison on Various Model Scales In addition to LLaVA-1.5-7B and Qwen-VL-2.5-7B, we compare baselines with various sizes of LLaVAs (*i.e.*, LLaVA-1.5-1B, 3B, and 13B) in Tab. 6. As shown in the table, OASIS consistently outperforms the baselines across model sizes by a significant margin, demonstrating its general applicability. We also provide model scaling results across different benchmarks and different model architecture (Qwen-VL) in Sec. A.7.

MICVIT				COAST			
Order	Method	$A_{avg} \uparrow$	$A_{last} \uparrow$	Order	Method	$A_{avg} \uparrow$	$A_{last} \uparrow$
ONCHoPDH	Random	61.15±0.34	67.29±0.61	cdim	Random	23.57±0.17	30.80±0.30
	InfoBatch	60.82±0.75	68.88±1.09		InfoBatch	22.93±0.73	29.14±0.56
	DivBS	61.07±0.58	69.06±0.54		DivBS	23.41±0.14	31.72±0.18
	OASIS (Ours)	64.39±0.58	71.76±0.72		OASIS (Ours)	25.67±0.35	34.23±0.38
HDPHoCNO	Random	57.85±0.63	66.94±0.74	imcd	Random	21.45±0.39	29.25±0.45
	InfoBatch	58.04±0.47	67.69±0.80		InfoBatch	22.39±0.53	30.04±0.72
	DivBS	58.39±0.28	69.32±0.53		DivBS	23.84±0.26	31.10±0.56
	OASIS (Ours)	61.87±0.82	70.98±0.64		OASIS (Ours)	25.01±0.41	33.96±0.53
PNHOHoCD	Random	60.31±0.88	68.18±0.62	dmci	Random	22.17±0.42	30.64±0.33
	InfoBatch	60.25±0.35	67.52±0.16		InfoBatch	20.31±0.25	29.53±1.28
	DivBS	61.70±0.74	69.16±0.43		DivBS	22.92±0.33	32.31±0.74
	OASIS (Ours)	63.12±0.58	71.31±0.91		OASIS (Ours)	24.36±0.59	33.28±0.26

Table 7: **Ablation on task order on MICVIT and COAST benchmarks.** In MICVIT, ONCHoPDH, HDPHoCNO, and PNHOHoCD denote task orders of Bongard-OpenWorld \rightarrow NLVR2 \rightarrow Co-Instruct-DB, Bongard-HOI \rightarrow PatternCom \rightarrow DVQA \rightarrow HQ Edit, HQ Edit \rightarrow DVQA \rightarrow PatternCom \rightarrow Bongard-HOI \rightarrow NLVR2 \rightarrow Bongard-OpenWorld, and PatternCom \rightarrow NLVR2 \rightarrow HQ Edit \rightarrow Bongard-Openworld \rightarrow Bongard-HOI \rightarrow Co-Instruct-DB \rightarrow DVQA, respectively. In COAST, cdim, imcd, and dmci denote task orders of ChartQA \rightarrow DocVQA \rightarrow IconQA \rightarrow MedicalQA, IconQA \rightarrow MedicalQA \rightarrow ChartQA \rightarrow DocVQA, and DocVQA \rightarrow MedicalQA \rightarrow ChartQA \rightarrow IconQA, respectively.

Comparison on Different Task Order In Tab. 7, we examine the effect of task order using three different task sequences for the top-3 performing baselines and OASIS on COAST and MICVIT. For each setting, 6.25% of the training data is selected using OASIS or the top three baseline methods. The results demonstrate that regardless of the change in task order, OASIS outperforms other selection methods.

4.3 ABLATION STUDY

Ablation Study on Different Components. We ablate OASIS to investigate the benefit of each proposed component, and summarize the results in Tab. 5. As shown, ORIS significantly improves performance by identifying informative samples, while SIREN further enhances performance by reducing redundancy between selected samples. Consistent improvements across benchmarks validate the effectiveness of our proposed components.

Additional Analyses. We further evaluate OASIS across different model scales (*e.g.*, 0.5B, 1B, 3B, and 13B) (Sec. A.7) and analyze informativeness metric evaluations (Sec. A.8), effect of EMA ratio β (Sec. A.9), accuracy over time (Sec. A.10), and comparison of the number of selected samples (Sec. A.13).

5 CONCLUSION

We address the challenge of achieving high performance while training on a small subset of data in the online CIT setup. Prior works select a fixed number of samples per batch, which struggles in the CIT setup due to varying informativeness and frequent redundancy from co-occurring similar samples. To address this, we propose OASIS, comprising two components: ORIS, which selects informative samples by estimating each sample’s informativeness relative to all previously encountered data, and SIREN, which reduces redundancy via sample-wise similarity. Extensive experiments on diverse CIT benchmarks demonstrate that our proposed OASIS consistently selects informative samples, outperforms existing baselines, and achieves performance comparable to full-data training.

Our method efficiently selects informative samples in CIT, requiring only a forward pass without any backward computation. A promising future direction is to remove this forward-pass requirement entirely, enabling near-instantaneous adaptation and substantially faster processing for real-time CIT applications.

REFERENCES

- Amro Abbas, Kushal Tirumala, Dániel Simig, Surya Ganguli, and Ari S Morcos. Semdedup: Data-efficient learning at web-scale through semantic deduplication. *arXiv preprint arXiv:2303.09540*, 2023.
- Amro Abbas, Evgenia Rusak, Kushal Tirumala, Wieland Brendel, Kamalika Chaudhuri, and Ari S Morcos. Effective pruning of web-scale datasets based on complexity of concept clusters. *arXiv preprint arXiv:2401.04578*, 2024.
- Haoping Bai, Shancong Mou, Tatiana Likhomanenko, Ramazan Gokberk Cinbis, Oncel Tuzel, Ping Huang, Jiulong Shan, Jianjun Shi, and Meng Cao. Vision datasets: A benchmark for vision-based industrial inspection. *arXiv preprint arXiv:2306.07890*, 2023a.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023b.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025.
- Jihwan Bang, Heesu Kim, YoungJoon Yoo, Jung-Woo Ha, and Jonghyun Choi. Rainbow memory: Continual learning with a memory of diverse samples. In *CVPR*, 2021.
- Lucas Caccia, Jing Xu, Myle Ott, Marcaurelio Ranzato, and Ludovic Denoyer. On anytime learning at macroscale. In *CoLLAs*. PMLR, 2022.
- Meng Cao, Yuyang Liu, Yingfei Liu, Tiancai Wang, Jiahua Dong, Henghui Ding, Xiangyu Zhang, Ian Reid, and Xiaodan Liang. Continual llava: Continual instruction tuning in large vision-language models. *arXiv preprint arXiv:2411.02564*, 2024.
- Cheng Chen, Junchen Zhu, Xu Luo, Hengtao Shen, Jingkuan Song, and Lianli Gao. Coin: A benchmark of continual instruction tuning for multimodal large language models. In *NeurIPS*, 2024.
- Jinpeng Chen, Runmin Cong, Yuzhi Zhao, Hongzheng Yang, Guangneng Hu, Horace Ho Shing Ip, and Sam Kwong. Sefe: Superficial and essential forgetting eliminator for multimodal continual instruction tuning. *arXiv preprint arXiv:2505.02486*, 2025.
- Keqin Chen, Zhao Zhang, Weili Zeng, Richong Zhang, Feng Zhu, and Rui Zhao. Shikra: Unleashing multimodal llm’s referential dialogue magic. *arXiv preprint arXiv:2306.15195*, 2023.
- Louis HY Chen, Larry Goldstein, and Qi-Man Shao. *Normal approximation by Stein’s method*. Springer Science & Business Media, 2010.
- Hsiang-Yun Sherry Chien, Javier S Turek, Nicole Beckage, Vy A Vo, Christopher J Honey, and Ted L Willke. Slower is better: revisiting the forgetting mechanism in lstm for slower information decay. *arXiv preprint arXiv:2105.05944*, 2021.
- Cody Coleman, Christopher Yeh, Stephen Mussmann, Baharan Mirzasoleiman, Peter Bailis, Percy Liang, Jure Leskovec, and Matei Zaharia. Selection via proxy: Efficient data selection for deep learning. In *ICLR*, 2020.
- Anirban DasGupta. *Asymptotic theory of statistics and probability*, volume 180. Springer, 2008.
- Sam Dauncey, Christopher C. Holmes, Christopher Williams, and Fabian Falck. Approximations to the fisher information metric of deep generative models for out-of-distribution detection. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856.
- Victor H. de la Peña, Tze Leung Lai, and Qi-Man Shao. *Self-Normalized Processes: Limit Theory and Statistical Applications*. Springer, 2009.
- Danruo Deng, Guangyong Chen, Yang Yu, Furui Liu, and Pheng-Ann Heng. Uncertainty estimation by fisher information-based evidential deep learning. In *ICML*. PMLR, 2023a.

- Zhijie Deng, Peng Cui, and Jun Zhu. Towards accelerated model training via bayesian data selection. In *NeurIPS*, 2023b.
- Yunshu Du, Wojciech M Czarnecki, Siddhant M Jayakumar, Mehrdad Farajtabar, Razvan Pascanu, and Balaji Lakshminarayanan. Adapting auxiliary losses using gradient similarity. *arXiv preprint arXiv:1812.02224*, 2018.
- Stephanie Fu, Netanel Tamir, Shobhita Sundaram, Lucy Chai, Richard Zhang, Tali Dekel, and Phillip Isola. Dreamsim: Learning new dimensions of human visual similarity using synthetic data. In *NeurIPS*, volume 36, pp. 50742–50768, 2023.
- Sreyan Ghosh, Chandra Kiran Reddy Evuru, Sonal Kumar, Deepali Aneja, Zeyu Jin, Ramani Duraiswami, Dinesh Manocha, et al. A closer look at the limitations of instruction tuning. *arXiv preprint arXiv:2402.05119*, 2024.
- Yasir Ghunaim, Adel Bibi, Kumail Alhamoud, Motasem Alfarra, Hasan Abed Al Kader Hammoud, Ameysa Prabhu, Philip HS Torr, and Bernard Ghanem. Real-time evaluation in online continual learning: A new hope. In *CVPR*, 2023.
- Annette Rios Gonzales, Nicolas Spring, Tannon Kew, Marek Kostrzewa, Andreas Säuberli, Mathias Müller, and Sarah Ebling. A new dataset and efficient baselines for document-level text simplification in german. In *ACL Workshop*, pp. 152–161, 2021.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Haiyang Guo, Fanhu Zeng, Ziwei Xiang, Fei Zhu, Da-Han Wang, Xu-Yao Zhang, and Cheng-Lin Liu. Hide-llava: Hierarchical decoupling for continual instruction tuning of multimodal large language model. *arXiv preprint arXiv:2503.12941*, 2025.
- Animesh Gupta, Irtiza Hasan, Dilip K Prasad, and Deepak K Gupta. Data-efficient training of cnns and transformers with coresets: A stability perspective. *arXiv preprint arXiv:2303.02095*, 2023.
- Jinghan He, Haiyun Guo, Ming Tang, and Jinqiao Wang. Continual instruction tuning for large multimodal models. *arXiv preprint arXiv:2311.16206*, 2023.
- Xuehai He, Yichen Zhang, Luntian Mou, Eric Xing, and Pengtao Xie. Pathvqa: 30000+ questions for medical visual question answering. *arXiv preprint arXiv:2003.10286*, 2020.
- Yongquan He, Wenyuan Zhang, Xuancheng Huang, Peng Zhang, Lingxun Meng, Xiang Zhou, Ke Zeng, and Xunliang Cai. Don't half-listen: Capturing key-part information in continual instruction tuning. In *ACL*, 2025.
- Aral Hekimoglu, Adrian Brucker, Alper Kagan Kayali, Michael Schmidt, and Alvaro Marcos-Ramiro. Active learning for object detection with non-redundant informative sampling. *arXiv preprint arXiv:2307.08414*, 2023.
- Feng Hong, Yueming Lyu, Jiangchao Yao, Ya Zhang, Ivor W Tsang, and Yanfeng Wang. Diversified batch selection for training acceleration. *arXiv preprint arXiv:2406.04872*, 2024.
- Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. In *ICLR*, 2022.
- Yebowen Hu, Tim Ganter, Hanieh Deilamsalehy, Franck Dernoncourt, Hassan Foroosh, and Fei Liu. Meetingbank: A benchmark dataset for meeting summarization. In *ACL*, 2023.
- Mude Hui, Siwei Yang, Bingchen Zhao, Yichun Shi, Heng Wang, Peng Wang, Cihang Xie, and Yuyin Zhou. HQ-edit: A high-quality dataset for instruction-based image editing. In *ICLR*, 2025.
- Zhouyuan Huo, Bin Gu, and Heng Huang. Training neural networks using features replay. In *NeurIPS*, 2018.

- Mohit Iyyer, Varun Manjunatha, Anupam Guha, Yogarshi Vyas, Jordan Boyd-Graber, Hal Daume, and Larry S Davis. The amazing mysteries of the gutter: Drawing inferences between panels in comic book narratives. In *CVPR*, pp. 7186–7195, 2017.
- Mohammad Jafari, Yimeng Zhang, Yihua Zhang, and Sijia Liu. The power of few: Accelerating and enhancing data reweighting with coresets selection. In *ICASSP*, pp. 7100–7104. IEEE, 2024.
- Dongfu Jiang, Xuan He, Huaye Zeng, Cong Wei, Max Ku, Qian Liu, and Wenhui Chen. Mantis: Interleaved multi-image instruction tuning. *Transactions on Machine Learning Research*, 2024.
- Huaizu Jiang, Xiaojian Ma, Weili Nie, Zhiding Yu, Yuke Zhu, and Anima Anandkumar. Bongard-hoi: Benchmarking few-shot visual reasoning for human-object interactions. In *CVPR*, pp. 19056–19065, 2022.
- Bing-Yi Jing, Qi-Man Shao, and Qi Wang. Self-normalized cramer-type large deviations for independent random variables. *Annals of Probability*, 31(4):2167–2215, 2003.
- Kushal Kafle, Brian Price, Scott Cohen, and Christopher Kanan. Dvqa: Understanding data visualizations via question answering. In *CVPR*, pp. 5648–5656, 2018.
- Angelos Katharopoulos and François Fleuret. Not all samples are created equal: Deep learning with importance sampling. In *ICML*, pp. 2525–2534. PMLR, 2018.
- Been Kim, Rajiv Khanna, and Oluwasanmi O Koyejo. Examples are not enough, learn to criticize! criticism for interpretability. In *Advances in neural information processing systems*, 2016.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.
- Hyunseo Koh, Dahyun Kim, Jung-Woo Ha, and Jonghyun Choi. Online continual learning on class incremental blurry task configuration with anytime inference. In *ICLR*, 2022.
- Hyunseo Koh, Minhyuk Seo, Jihwan Bang, Hwanjun Song, Deokki Hong, Seulki Park, Jung-Woo Ha, and Jonghyun Choi. Online boundary-free continual learning by scheduled data prior. In *ICLR*, 2023.
- Kamil Książek, Hubert Jastrzębski, Bartosz Trojan, Krzysztof Pniaczek, Michał Karp, and Jacek Tabor. Fenec: Enhancing continual learning via feature clustering with neighbor-or logit-based classification. *arXiv preprint arXiv:2503.14301*, 2025.
- Allison Lau, Younwoo Choi, Vahid Balazadeh, Keertana Chidambaram, Vasilis Syrgkanis, and Rahul G Krishnan. Personalized adaptation via in-context preference learning. *arXiv preprint arXiv:2410.14001*, 2024.
- Jaewoo Lee, Boyang Li, and Sung Ju Hwang. Concept-skill transferability-based data selection for large vision-language models. In *EMNLP*, 2024.
- Lei Li, Yuwei Yin, Shicheng Li, Liang Chen, Peiyi Wang, Shuhuai Ren, Mukai Li, Yazheng Yang, Jingjing Xu, Xu Sun, et al. M³ it: A large-scale dataset towards multi-modal multilingual instruction tuning. *arXiv preprint arXiv:2306.04387*, 2023a.
- Ming Li, Yong Zhang, Zhitao Li, Jiuhai Chen, Lichang Chen, Ning Cheng, Jianzong Wang, Tianyi Zhou, and Jing Xiao. From quantity to quality: Boosting llm performance with self-guided data selection for instruction tuning. *arXiv preprint arXiv:2308.12032*, 2023b.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. In *CVPR*, 2024a.
- Zikang Liu, Kun Zhou, Wayne Xin Zhao, Dawei Gao, Yaliang Li, and Ji-Rong Wen. Less is more: High-value data selection for visual instruction tuning. *arXiv preprint arXiv:2403.09559*, 2024b.

- Pan Lu, Liang Qiu, Jiaqi Chen, Tony Xia, Yizhou Zhao, Wei Zhang, Zhou Yu, Xiaodan Liang, and Song-Chun Zhu. Iconqa: A new benchmark for abstract diagram understanding and visual language reasoning. In *NeurIPS*, 2021a.
- Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. Learn to explain: Multimodal reasoning via thought chains for science question answering. *Advances in Neural Information Processing Systems*, 35:2507–2521, 2022.
- Shuai Lu, Daya Guo, Shuo Ren, Junjie Huang, Alexey Svyatkovskiy, Ambrosio Blanco, Colin Clement, Dawn Drain, Daxin Jiang, Duyu Tang, et al. Codexglue: A machine learning benchmark dataset for code understanding and generation. In *NeurIPS*, 2021b.
- Chengcheng Ma, Yang Liu, Jiankang Deng, Lingxi Xie, Weiming Dong, and Changsheng Xu. Understanding and mitigating overfitting in prompt tuning for vision-language models. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(9):4616–4629, 2023.
- Andrew Maas, Raymond E Daly, Peter T Pham, Dan Huang, Andrew Y Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies*, pp. 142–150, 2011.
- Adyasha Maharana, Jaehong Yoon, Tianlong Chen, and Mohit Bansal. Adapt-\$\infty\$: Scalable continual multimodal instruction tuning via dynamic data selection. In *ICLR*, 2025.
- Shivangi Mahto, Vy Ai Vo, Javier S Turek, and Alexander Huth. Multi-timescale representation learning in lstm language models. In *ICLR*, 2020.
- Ahmed Masry, Xuan Long Do, Jia Qing Tan, Shafiq Joty, and Enamul Hoque. Chartqa: A benchmark for question answering about charts with visual and logical reasoning. In *ACL*, pp. 2263–2279, 2022.
- Minesh Mathew, Dimosthenis Karatzas, and CV Jawahar. Docvqa: A dataset for vqa on document images. In *WACV*, 2021.
- Sören Mindermann, Jan M Brauner, Muhammed T Razzak, Mrinank Sharma, Andreas Kirsch, Winnie Xu, Benedikt Höltgen, Aidan N Gomez, Adrien Morisot, Sebastian Farquhar, et al. Prioritized training on points that are learnable, worth learning, and not yet learnt. In *ICML*. PMLR, 2022.
- Swaroop Mishra, Arindam Mitra, Neeraj Varshney, Bhavdeep Sachdeva, Peter Clark, Chitta Baral, and Ashwin Kalyan. Numglue: A suite of fundamental yet challenging mathematical reasoning tasks. In *ACL*, 2022.
- Aristeidis Panos, Rahaf Aljundi, Daniel Olmeda Reino, and Richard E Turner. Efficient few-shot continual learning in vision-language models. *arXiv preprint arXiv:2502.04098*, 2025.
- Mansheej Paul, Surya Ganguli, and Gintare Karolina Dziugaite. Deep learning on a data diet: Finding important examples early in training. In *NeurIPS*, 2021.
- Bill Psomas, Ioannis Kakogeorgiou, Nikos Efthymiadis, Giorgos Toliás, Ondřej Chum, Yannis Avrithis, and Konstantinos Karantzas. Composed image retrieval for remote sensing. In *IGARSS*, pp. 8526–8534. IEEE, 2024.
- Ziheng Qin, Kai Wang, Zangwei Zheng, Jianyang Gu, Xiangyu Peng, Zhaopan Xu, Daquan Zhou, Lei Shang, Baigui Sun, Xuansong Xie, et al. Infobatch: Lossless training speed up by unbiased dynamic data pruning. In *ICLR*, 2024.
- Anastasia Razdaibiedina, Yuning Mao, Rui Hou, Madian Khabsa, Mike Lewis, and Amjad Almahairi. Progressive prompts: Continual learning for language models. In *ICLR*, 2023.
- Leslie Rice, Eric Wong, and Zico Kolter. Overfitting in adversarially robust deep learning. In *ICML*, 2020.
- David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience replay for continual learning. In *NeurIPS*, 2019.

- Minhyuk Seo, Seongwon Cho, Minjae Lee, Diganta Misra, Hyeonbeom Choi, Seon Joo Kim, and Jonghyun Choi. Just say the name: Online continual learning with category names only via data generation. *arXiv preprint arXiv:2403.10853*, 2024a.
- Minhyuk Seo, Hyunseo Koh, Wonje Jeung, Minjae Lee, San Kim, Hankook Lee, Sungjun Cho, Sungik Choi, Hyunwoo Kim, and Jonghyun Choi. Learning equi-angular representations for online continual learning. In *CVPR*, pp. 23933–23942, 2024b.
- Minhyuk Seo, Hyunseo Koh, and Jonghyun Choi. Budgeted online continual learning by adaptive layer freezing and frequency-based sampling. In *ICLR*, 2025.
- Agam Shah, Suvan Paturi, and Sudheer Chava. Trillion dollar words: A new financial dataset, task & market analysis. In *ACL*, 2023.
- Hyo-Sang Shin and Hae-In Lee. A new exponential forgetting algorithm for recursive least-squares parameter estimation. *arXiv preprint arXiv:2004.03910*, 2020.
- Seungjae Shin, Heesun Bae, Donghyeok Shin, Weonyoung Joo, and Il-Chul Moon. Loss-curvature matching for dataset selection and condensation. In *AISTATS*. PMLR, 2023.
- Alexander Soen and Ke Sun. On the variance of the fisher information for deep learning. In *NeurIPS*, 2021.
- Ben Sorscher, Robert Geirhos, Shashank Shekhar, Surya Ganguli, and Ari Morcos. Beyond neural scaling laws: beating power law scaling via data pruning. In *NeurIPS*, 2022.
- Alane Suhr, Mike Lewis, James Yeh, and Yoav Artzi. A corpus of natural language for visual reasoning. In *ACL*, 2017.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*, 2018.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. Superglue: A stickier benchmark for general-purpose language understanding systems. *Advances in neural information processing systems*, 32, 2019.
- Xiao Wang, Yuansen Zhang, Tianze Chen, Songyang Gao, Senjie Jin, Xianjun Yang, Zhiheng Xi, Rui Zheng, Yicheng Zou, Tao Gui, et al. Trace: A comprehensive benchmark for continual learning in large language models. *arXiv preprint arXiv:2310.06762*, 2023.
- Ziqi Wang, Chang Che, Qi Wang, Yangyang Li, Zenglin Shi, and Meng Wang. Separable mixture of low-rank adaptation for continual visual instruction tuning. *arXiv preprint arXiv:2411.13949*, 2024.
- Haoning Wu, Hanwei Zhu, Zicheng Zhang, Erli Zhang, Chaofeng Chen, Liang Liao, Chunyi Li, Annan Wang, Wenxiu Sun, Qiong Yan, et al. Towards open-ended visual quality comparison. In *ECCV*, pp. 360–377. Springer, 2024a.
- Rujie Wu, Xiaojian Ma, Zhenliang Zhang, Wei Wang, Qing Li, Song-Chun Zhu, and Yizhou Wang. Bongard-openworld: Few-shot reasoning for free-form visual concepts in the real world. In *ICLR*, 2024b.
- Zhiyang Xu, Ying Shen, and Lifu Huang. Multiinstruct: Improving multi-modal zero-shot learning via instruction tuning. In Anna Rogers, Jordan L. Boyd-Graber, and Naoaki Okazaki (eds.), *ACL*, pp. 11445–11465. Association for Computational Linguistics, 2023.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.
- Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye, Ming Yan, Yiyang Zhou, Junyang Wang, Anwen Hu, Pengcheng Shi, Yaya Shi, et al. mplug-owl: Modularization empowers large language models with multimodality. *arXiv preprint arXiv:2304.14178*, 2023.

- Zhenfei Yin, Jiong Wang, Jianjian Cao, Zhelun Shi, Dingning Liu, Mukai Li, Xiaoshui Huang, Zhiyong Wang, Lu Sheng, Lei Bai, et al. Lamm: Language-assisted multi-modal instruction-tuning dataset, framework, and benchmark. In *NeurIPS*, 2023.
- Yuexiang Zhai, Shengbang Tong, Xiao Li, Mu Cai, Qing Qu, Yong Jae Lee, and Yi Ma. Investigating the catastrophic forgetting in multimodal large language models. In *CPAL*, 2024.
- Jiarui Zhang. Guided profile generation improves personalization with llms. In *ACL*, 2024.
- Kai Zhang, Lingbo Mo, Wenhui Chen, Huan Sun, and Yu Su. Magicbrush: A manually annotated dataset for instruction-guided image editing. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *NeurIPS*. Curran Associates, Inc., 2023.
- Chenye Zhao, Yingjie Li, and Cornelia Caragea. C-stance: A large dataset for chinese zero-shot stance detection. In *ACL*, pp. 13369–13385, 2023.
- Haizhong Zheng, Rui Liu, Fan Lai, and Atul Prakash. Coverage-centric coreset selection for high pruning rates. In *ICLR*, 2023.
- Junhao Zheng, Qianli Ma, Zhen Liu, Binqun Wu, and Huawen Feng. Beyond anti-forgetting: Multi-modal continual instruction tuning with positive forward transfer. *arXiv preprint arXiv:2401.09181*, 2024.
- Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023.
- Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. MiniGPT-4: Enhancing vision-language understanding with advanced large language models. In *ICLR*, 2024a. URL <https://openreview.net/forum?id=1tZbq88f27>.
- Didi Zhu, Zhongyi Sun, Zexi Li, Tao Shen, Ke Yan, Shouhong Ding, Kun Kuang, and Chao Wu. Model tailor: Mitigating catastrophic forgetting in multi-modal large language models. In *ICML*, 2024b.

A APPENDIX

A.1 PROOF OF THEOREM 3.1

In continual instruction tuning (CIT), the data distribution evolves gradually over time (*e.g.*, more Christmas-related samples in winter and more air-conditioner-related samples in summer) (Koh et al., 2023). Although this constitutes a distribution shift, within short horizons (*e.g.*, a week), the distribution of training batches can be regarded as approximately stable under memory-only training (Koh et al., 2022; Seo et al., 2024a), even in the presence of sharp task boundaries. Specifically, memory-only training constructs batches by only retrieving from an episodic memory containing previously encountered samples, thereby inducing smoother changes of training distribution, unlike ER (Rolnick et al., 2019), which mixes continuously encountered streaming data with memory samples and may cause abrupt shifts. As a result, our choice of memory-only training for CIT motivates the local stationarity assumption (Assumption 1) and local weak dependence assumption (Assumption 2) introduced below.

Setting. At time t , we observe a batch $\{(x_i^{(t)}, y_i^{(t)})\}_{i=1}^{N_B}$ and calculate sample-wise FI $\{I_i^{(t)}\}_{i=1}^{N_B}$ using Fisher Information Matrix (Eq. 2), while updating EMA μ_t and EMV σ_t of FI (Eq. 3) using FI averaged by batch ($\bar{I}^{(t)} = \frac{1}{N_B} \sum_{i=1}^{N_B} I_i^{(t)}$). We decompose $\bar{I}^{(t)}$ as

$$\bar{I}^{(t)} = m_t + \varepsilon_t,$$

where $m_t := \mathbb{E}[\bar{I}^{(t)}]$ denotes the underlying (population) FI at time t , and $\varepsilon_t := \bar{I}^{(t)} - m_t$ represents the centered sampling noise.

Assumption 1 (Local Stationarity). *Although the data distribution P_t and model parameters θ_t evolve over time in continual learning, we assume that such changes are smooth. In particular, within a short horizon comparable to the EMA effective window $L_{\text{eff}} \approx \frac{1+\beta}{1-\beta}$, the mean $m_t = \mathbb{E}[\bar{I}^{(t)}]$ and variance $s_t^2 = \text{Var}(\bar{I}^{(t)})$ remain nearly constant. Formally, for temporal lag $h \in \mathbb{Z}$ with $|h| \leq L_{\text{eff}}$,*

$$\sup_{|h| \leq L_{\text{eff}}} |m_{t+h} - m_t| = o(\sqrt{1-\beta}), \quad (9)$$

$$\sup_{|h| \leq L_{\text{eff}}} |s_{t+h}^2 - s_t^2| = o(1). \quad (10)$$

Assumption 2 (Local Weak Dependence). *The noise $\{\varepsilon_t\}$ is assumed to be locally weakly dependent. That is, while sequential batches are correlated due to temporal proximity, correlations decay sufficiently fast as the temporal lag h grows. Formally, for lag $h \rightarrow \infty$,*

$$\text{Cov}(\varepsilon_t, \varepsilon_{t+h}) \rightarrow 0.$$

Lemma A.1 ($\hat{I}^{(t)}$ has finite $(2 + \delta)$ -moments). *We estimate the Fisher information (FI) as the trace of the Fisher information matrix (FIM) in Eq. 2. Since the diagonal entries of the FIM are typically large in magnitude, the FI of a sample $(x_i^{(t)}, y_i^{(t)}) \in \mathcal{B}_t$, denoted $I_i^{(t)}$, can be approximated by a chi-square distribution (Dauncey et al., 2024),*

$$I_i^{(t)} \stackrel{d}{\approx} \sigma_t^2 \chi_{k_t}^2,$$

where σ_t^2 is a scale parameter chosen such that the mean and variance match those of the Fisher information trace, i.e.,

$$\mathbb{E}[I_i^{(t)}] \approx \sigma_t^2 k_t, \quad \text{Var}(I_i^{(t)}) \approx 2\sigma_t^4 k_t.$$

Because both the sum of chi-square variables remains chi-square distributed, the batch-averaged FI $\bar{I}^{(t)}$ also follow chi-square approximations:

$$\bar{I}^{(t)} \stackrel{d}{\approx} \frac{\sigma_t^2}{B} \chi_{k_t \mathcal{B}_t}^2.$$

Since chi-square distributions have finite moments of all orders, it follows that for every $\delta > 0$,

$$\mathbb{E}\left[|\bar{I}^{(t)}|^{2+\delta}\right] < \infty, \quad \mathbb{E}\left[|I_i^{(t)}|^{2+\delta}\right] < \infty.$$

Lemma A.2 (Normal approximation rate for FI trace via chi-square). *Let $X \sim \chi_k^2$ and consider the standardized variable $\tilde{X} = (X - k)/\sqrt{2k}$. Then classical Berry-Esseen-type results (DasGupta (2008); Chen et al. (2010)) guarantee that*

$$\sup_{x \in \mathbb{R}} \left| \Pr(\tilde{X} \leq x) - \Phi(x) \right| = O\left(\frac{1}{\sqrt{k}}\right).$$

Applied to $\bar{I}^{(t)}$ with $\bar{I}^{(t)} \approx \frac{\sigma_t^2}{B} \chi_{k_t \mathcal{B}_t}^2$, we obtain

$$\begin{aligned} \sup_x \left| \Pr\left(\frac{\bar{I}^{(t)} - \mathbb{E}[\bar{I}^{(t)}]}{\sqrt{\text{Var}(\bar{I}^{(t)})}} \leq x\right) - \Phi(x) \right| \\ = O\left(\frac{1}{\sqrt{k_t \mathcal{B}_t}}\right). \end{aligned} \quad (11)$$

Lemma A.3 (EMA/EMV are consistent approximations of the true mean and variance). *Let $m_t = \mathbb{E}[\bar{I}^{(t)}]$ be the true mean of the batch-averaged FI and $s_t^2 = \text{Var}(I_i^{(t)})$ the true sample-wise variance at time t . Let μ_{t-1} and v_{t-1} be the EMA and EMV computed from past batches, and define $\hat{\sigma}_{t-1}^2 := B \cdot v_{t-1}$. Then, under Assumptions 1 and 2,*

$$|\mu_{t-1} - m_t| = o(s_t), \quad \frac{\hat{\sigma}_{t-1}^2}{s_t^2} \xrightarrow{p} 1.$$

In words: the EMA mean μ_{t-1} is essentially the same as the true mean m_t , and the EMV-based variance estimator $\hat{\sigma}_{t-1}^2$ consistently estimates the true variance s_t^2 .

Proof Lemma A.3. Mean part. The EMA is a weighted moving average:

$$\mu_{t-1} = (1 - \beta) \sum_{k \geq 1} \beta^{k-1} \bar{I}_{t-k}.$$

Local stationarity (Assumption 1) says that the true mean m_t does not change much over the EMA window $L_{\text{eff}} \approx \frac{1+\beta}{1-\beta}$. Hence each difference $|m_{t-k} - m_t|$ is small, and the exponential weights $(1 - \beta)\beta^{k-1}$ quickly decay. As a result, the EMA bias

$$|\mu_{t-1} - m_t| = (1 - \beta) \sum_{k \geq 1} \beta^{k-1} |m_{t-k} - m_t|$$

is negligible compared to the natural variability s_t ; formally, it is $o(s_t)$.

Variance part. The EMV is defined as

$$v_{t-1} = (1 - \lambda) \sum_{k \geq 1} \lambda^{k-1} (\bar{I}_{t-k} - \mu_{t-k})^2.$$

This is an exponentially weighted average of past squared deviations. Local weak dependence (Assumption 2) ensures that correlations between far-apart batches vanish, so v_{t-1} consistently estimates $\text{Var}(\bar{I}_t)$ (this is the usual consistency of exponential-kernel HAC estimators). Finally, since $\text{Var}(\bar{I}_t) = s_t^2/B$ under i.i.d. samples in the batch, multiplying v_{t-1} by B yields a consistent estimator of the sample-wise variance s_t^2 .

Thus, both the EMA mean and the EMV variance converge to their true population counterparts. \square

Proof of Theorem 3.1. We prove that the one-step-lag standardized statistic $Z_{t,i} = \frac{I_i^{(t)} - \mu_{t-1}}{\hat{\sigma}_{t-1}}$ is well-approximated by $\mathcal{N}(0, 1)$ under Assumptions 1–2 and the chi-square structure of FI traces.

Step 1: Within-batch normal approximation. By Lemma A.1, the sample-wise FI admits a (scaled) chi-square representation: $I_i^{(t)} \stackrel{d}{\approx} \sigma_t^2 \chi_{k_t}^2$ for some effective degrees-of-freedom k_t . Therefore, by Lemma A.2,

$$\sup_{x \in \mathbb{R}} \left| \Pr\left(\frac{I_i^{(t)} - m_t}{s_t} \leq x\right) - \Phi(x) \right| = O(k_t^{-1/2}), \quad (12)$$

where $m_t = \mathbb{E}[\bar{I}_t] = \mathbb{E}[I_i^{(t)}]$ and $s_t^2 = \text{Var}(I_i^{(t)})$.

Step 2: EMA/EMV ratio-consistency. The ideal normalization requires (m_t, s_t) , but in practice we use the EMA mean μ_{t-1} and EMV variance $\hat{\sigma}_{t-1}^2$ from past batches. By Lemma A.3, local stationarity ensures that $|\mu_{t-1} - m_t| = o(s_t)$, and local weak dependence guarantees $\hat{\sigma}_{t-1}^2/s_t^2 \xrightarrow{p} 1$. Hence $(\mu_{t-1}, \hat{\sigma}_{t-1})$ consistently approximate (m_t, s_t) , justifying their use as practical normalizers.

Step 3: Slutsky/self-normalization. We can decompose $Z_{t,i}$ as

$$Z_{t,i} = \frac{I_i^{(t)} - m_t}{s_t} \cdot \frac{s_t}{\hat{\sigma}_{t-1}} - \frac{\mu_{t-1} - m_t}{\hat{\sigma}_{t-1}}. \quad (13)$$

By equation 12, the standardized term $(I_i^{(t)} - m_t)/s_t$ is $O(k_t^{-1/2})$, close to $\mathcal{N}(0, 1)$ in Kolmogorov distance. Lemma A.3 further ensures that $s_t/\hat{\sigma}_{t-1} \xrightarrow{p} 1$ and $(\mu_{t-1} - m_t)/\hat{\sigma}_{t-1} = o_p(1)$. Hence, both adjustment factors vanish asymptotically, and by Slutsky's theorem, the distribution of $Z_{t,i}$ retains the normal approximation:

$$\sup_{x \in \mathbb{R}} \left| \Pr(Z_{t,i} \leq x) - \Phi(x) \right| \leq C_1 k_t^{-1/2} + o(1).$$

Step 4: One-step lag and leakage. We used $(\mu_{t-1}, \hat{\sigma}_{t-1})$ built from *past* batches only, so $I_i^{(t)}$ is not used inside its own normalizer; hence no self-normalization bias occurs. (If one includes t in the EMA/EMV, the induced bias is $O(1 - \beta)$.)

Combining the steps establishes the stated claim $Z_{t,i} \approx \mathcal{N}(0, 1)$ and the quantitative bound. \square

Method	Selection Ratio (%)							
	12.5				25.0			
	COAST		Adapt		COAST		Adapt	
$A_{avg} \uparrow$	$A_{last} \uparrow$	$A_{avg} \uparrow$	$A_{last} \uparrow$	$A_{avg} \uparrow$	$A_{last} \uparrow$	$A_{avg} \uparrow$	$A_{last} \uparrow$	
Full-Data Training	31.56±1.42	39.06±0.55	55.89±0.20	54.26±0.23	31.56±1.42	39.06±0.55	55.89±0.20	54.26±0.23
Random	24.58±0.17	32.45±0.92	49.04±0.49	44.19±0.62	25.82±0.10	34.28±0.28	50.68±0.30	47.58±0.34
GradNorm (ICML 2018)	23.35±0.38	29.44±0.28	48.47±0.63	45.36±0.62	25.67±0.21	32.85±0.36	51.29±0.54	48.96±0.81
Self-Sup (NeurIPS 2022)	23.72±0.27	31.37±0.75	47.04±0.26	43.94±1.38	25.17±0.37	32.84±0.42	49.34±0.85	46.46±0.47
COINCIDE (EMNLP 2024)	24.45±0.84	31.08±0.59	48.52±0.55	44.76±0.43	26.22±0.34	34.90±0.57	51.74±0.64	47.40±0.31
DBP (ICLR 2024)	22.87±0.33	25.81±0.37	46.24±0.36	43.30±1.08	24.74±0.70	32.13±0.64	48.58±0.43	45.65±0.84
InfoBatch (ICLR 2024)	24.78±0.58	31.74±0.62	49.77±0.82	47.54±0.36	25.30±0.35	34.49±0.45	51.55±0.47	49.04±0.77
DivBS (ICML 2024)	25.13±0.12	33.30±0.87	49.96±0.48	47.17±0.51	26.91±0.02	35.01±0.16	52.36±0.05	<u>50.02±1.21</u>
TIVE (arXiv:2403)	23.41±0.23	30.17±0.43	47.75±0.73	44.38±0.99	25.18±0.78	34.08±0.60	49.69±1.24	45.97±0.88
Adapt-∞ (ICLR 2025)	24.73±0.36	31.69±0.35	48.38±0.20	45.42±0.46	26.07±0.25	33.50±0.04	50.56±0.25	47.54±0.53
OASIS (Ours)	27.13±0.70	35.42±0.49	51.73±0.33	49.02±0.85	28.72±0.63	37.55±0.28	54.58±0.14	51.87±0.49

Table 8: **Quantitative comparison on COAST and Adapt under different selection ratios (12.5% and 25.0%)**. LLaVA-1.5-7B is used as the MLLM. Bold indicates the highest performance; underlined results are within the 0.05 t-test significance level. 'Full-Data Training' uses all data without selection.

A.2 EXPERIMENTAL RESULTS ON VARIOUS BENCHMARKS

In addition to MICVIT, we compare OASIS with sample selection methods in COAST (Cao et al., 2024) and Adapt (Maharana et al., 2025) under selection ratios of 12.5% and 25.0% in Tab. 8. OASIS not only outperforms selection baselines but also achieves comparable results with a model trained on all data (*i.e.*, Full-Data Training) in diverse benchmarks.

A.3 BENCHMARK CONFIGURATION DETAILS

We use Long Sequence (Razdaibiedina et al., 2023) and TRACE (Wang et al., 2023) as our text-only benchmarks. Long Sequence consists of datasets from two existing CL benchmarks, GLUE (Wang et al., 2018) and SuperGLUE (Wang et al., 2019), along with additional IMDB movie reviews dataset (Maas et al., 2011). TRACE benchmark consists of 8 tasks, NumGLUE-ds (Mishra et al., 2022), ScienceQA (Lu et al., 2022), Py150 (Lu et al., 2021b), C-STANCE (Zhao et al., 2023), FOMC (Shah et al., 2023), 20Minuten (Gonzales et al., 2021), NumGLUE-cm (Mishra et al., 2022), and MeetingBank (Hu et al., 2023), encompassing domain-specific tasks, multilingual capabilities, code generation, and mathematical reasoning.

We use COAST (Cao et al., 2024), Adapt (Maharana et al., 2025), and MICVIT for our multi-modal benchmarks. For the COAST benchmark, we use the COAST-domain, which emulates a scenario where MLLMs needs to continuously adapt to diverse domains. Specifically, it consists of DocVQA (Mathew et al., 2021), ChartQA (Masry et al., 2022), IconQA (Lu et al., 2021a), and MedicalQA (He et al., 2020), which correspond to the document, chart, icon, and medical domains, respectively. Adapt benchmark consists of visual instruction tuning datasets, such as M3IT (Li et al., 2023a), MiniGPT4 (Zhu et al., 2023), MANTIS (Jiang et al., 2024), LaMM (Yin et al., 2023), and VisionFLAN (Xu et al., 2023). MICVIT comprises NLVR2 (Suhr et al., 2017), Bongard-OpenWorld (Wu et al., 2024b), Bongard-HOI (Jiang et al., 2022), Co-Instruct-DB (Wu et al., 2024a), DVQA (Kafle et al., 2018), HQ-Edit (Hui et al., 2025), and PatternCom (Psomas et al., 2024). We summarize the task configuration of each benchmark, including MICVIT in Tab. 9.

Dataset	# of samples / task	# of tasks	Task order
COAST (Cao et al., 2024)	20,000	4	ChartQA → DocVQA → IconQA → MedicalQA
Adapt (Maharana et al., 2025)	80,000	4	M3IT → MANTIS → LaMM → VisionFLAN
MICVIT (Ours)	50K / 30K / 100K / 46K / 17K / 25K / 79K	7	Bongard-OpenWorld → NLVR2 → Co-Instruct-DB Bongard-HOI → PatternCom → DVQA → HQ Edit DBPedia → RTE → WiC → Amazon → Yahoo → AG News → MNLI → IMDB → MultiRC → CB → Yelp → COPA → QQP → SST2 → BoolQ
Long Sequence (Razdaibiedina et al., 2023)	10,000	15	NumGLUE-ds → ScienceQA → Py150 → C-STANCE → FOMC → 20Minuten → NumGLUE-cm → MeetingBank
TRACE (Wang et al., 2023)	5,000	8	

Table 9: **Task configurations of CIT benchmarks.**

A.4 SAMPLE SELECTION BASELINES

Self-Sup (Sorscher et al., 2022). Self-Sup performs K-means clustering on the output embeddings from the final layer of the pretrained model and selects the samples closest to each cluster centroid. By doing so, the most representative sample in the underlying data distribution can be selected.

GradNorm (Katharopoulos & Fleuret, 2018). GradNorm is an importance sampling method that prioritizes the selection of highly informative samples based on gradient magnitude. It computes each sample’s gradient norm with respect to the model’s last layer parameters and assigns higher selection probabilities to samples with larger gradient norms.

COINCIDE (Lee et al., 2024). COINCIDE leverages multi-layer output representations for K-means clustering. Specifically, features are extracted from five distinct layers of the MLLM to obtain the representation of the candidate samples. Then it applies K-means clustering to these features, and computes transferability and density metrics of each cluster to determine how many samples should be selected from the respective cluster. Finally, within each cluster, the maximum mean discrepancy (MMD) (Kim et al., 2016) is computed greedily to obtain the samples that best reflect the overall cluster distribution.

DBP (Abbas et al., 2024). DBP aims to improve dataset quality by removing redundant and irrelevant samples. First, following Semdedup (Abbas et al., 2023), DBP applies K-means clustering to group similar samples with last layer outputs and eliminate semantic duplicates by retaining only the sample farthest from the cluster centroid among highly similar pairs. Next, DBP performs CLIP score filtering, removing samples with low text-image similarities. Then, to prioritize more informative data, each cluster is evaluated for its complexity using two metrics: d_{inter} , the distance from the cluster centroid to other centroids, and d_{intra} , the average cosine distance of samples to their own centroid. A complexity score $C_j = d_{\text{inter},j} \cdot d_{\text{intra},j}$ is calculated for each cluster j , and samples are selected proportionally based on this score. Within each selected cluster, samples with the highest entropy are retained, ensuring that the final dataset is both diverse and rich in information.

InfoBatch (Qin et al., 2024). InfoBatch introduces a soft pruning strategy to improve training efficiency while preserving learning dynamics. Unlike traditional hard pruning methods that permanently discard data and risk introducing bias, InfoBatch probabilistically excludes a subset of well-learned samples—identified by their low loss values—during each training epoch. Furthermore, to maintain the integrity of the training trajectory, the gradients of the remaining samples are rescaled such that the expected aggregated gradient approximates that of the full batch.

DivBS (Abbas et al., 2024). From each online batch, DivBS focuses on selecting a fixed subset of training data that is both informative and diverse by maximizing the orthogonalized representativeness. To do so, DivBS utilizes the last layer gradient. DivBS aims to explicitly reduce inter-sample redundancy, ensuring the retained samples capture complementary and non-overlapping aspects of the batch distribution.

TIVE (Liu et al., 2024b). TIVE selects samples based on two criteria: instance influence and task difficulty, both derived from gradients computed across all layers of a reference model. Instance influence quantifies how much a given sample contributes to other samples during training, while task difficulty estimates the inherent complexity of learning a given task. Guided by these two metrics, TIVE prioritizes samples that are both highly influential and associated with difficult tasks.

Adapt- ∞ (Maharana et al., 2025). Rather than utilizing layer output representations, Adapt- ∞ uses gradients from the middle layer to cluster data into pseudo-tasks, enabling the model to identify and group related skills without requiring explicit task labels. It then performs multi-way sample selection within each cluster using a pool of scoring functions—such as entropy and the proposed Image Grounding score—to retain the most informative samples. Unlike our online continual instruction tuning setting, Adapt- ∞ was introduced in an offline continual instruction tuning setting, where task boundary assumptions exist, limiting its real-world applicability.

Method	Selection Ratio (%)					
	6.25		12.5		25.0	
	$A_{avg} \uparrow$	$A_{last} \uparrow$	$A_{avg} \uparrow$	$A_{last} \uparrow$	$A_{avg} \uparrow$	$A_{last} \uparrow$
Full-Data Training	34.59±0.44	41.90±0.48	34.59±0.44	41.90±0.48	34.59±0.44	41.90±0.48
Random	24.53±0.47	31.59±0.26	26.39±0.54	34.14±0.55	27.89±0.15	37.32±0.28
GradNorm (ICML 2018)	23.85±0.39	29.69±0.84	25.96±0.57	33.28±0.82	26.47±0.19	35.13±1.14
Self-Sup (NeurIPS 2022)	22.44±0.18	30.28±0.80	24.18±1.26	32.53±0.42	25.81±0.05	33.92±0.59
COINCIDE (EMNLP 2024)	21.93±0.45	30.73±0.16	24.74±0.74	33.84±0.43	27.48±0.57	35.27±0.65
DBP (ICLR 2024)	22.58±0.29	27.92±0.27	24.01±0.48	32.43±0.58	25.35±0.52	34.79±0.61
InfoBatch (ICLR 2024)	24.60±0.36	32.88±0.54	27.48±0.39	<u>34.26±1.39</u>	28.93±0.41	38.12±0.42
DivBS (ICML 2024)	25.45±0.64	33.18±0.51	28.29±0.48	<u>36.08±0.85</u>	30.82±0.28	39.87±0.74
TIVE (arXiv:2403)	22.52±0.82	28.90±0.33	25.83±0.62	<u>32.35±0.93</u>	27.14±0.45	34.42±0.14
Adapt-∞ (ICLR 2025)	23.84±0.53	29.48±0.29	24.85±0.84	33.43±0.79	26.78±0.63	35.38±0.38
OASIS (Ours)	27.29±0.18	35.68±0.62	29.16±0.13	36.27±0.75	31.04±0.38	<u>39.38±0.84</u>

Table 10: **Quantitative comparison between online sample selection methods on COAST benchmark.** We use Qwen-VL-2.5-7B as the MLLM. Bold indicates the highest performance; underlined results are within the 0.05 t-test significance level. 'Full-Data Training' uses all data without selection.

Method	Selection Ratio (%)					
	6.25		12.5		25.0	
	$A_{avg} \uparrow$	$A_{last} \uparrow$	$A_{avg} \uparrow$	$A_{last} \uparrow$	$A_{avg} \uparrow$	$A_{last} \uparrow$
Full-Data Training	45.73±1.68	43.92±0.85	45.73±1.68	43.92±0.85	45.73±1.68	43.92±0.85
Random	34.87±0.50	27.49±0.25	36.02±0.35	32.16±0.62	38.73±0.44	34.14±0.57
GradNorm (ICML 2018)	32.53±0.36	25.41±0.72	34.70±0.28	30.25±0.43	35.38±0.40	32.26±0.42
Self-Sup (NeurIPS 2022)	30.31±0.31	22.04±0.65	31.11±0.64	27.73±0.83	33.86±0.50	31.63±1.35
COINCIDE (EMNLP 2024)	32.76±0.17	24.39±0.31	33.93±0.72	29.34±0.47	35.82±0.08	32.38±0.55
DBP (ICLR 2024)	29.82±0.52	22.26±0.36	31.16±0.27	28.04±0.14	34.24±0.61	30.05±0.72
InfoBatch (ICLR 2024)	34.95±0.20	27.19±0.18	35.46±0.83	30.77±0.54	38.25±0.72	34.10±0.50
DivBS (ICML 2024)	34.24±0.75	28.63±0.91	37.24±0.41	31.68±1.39	<u>40.32±0.53</u>	<u>36.59±1.08</u>
TIVE (arXiv:2403)	31.46±0.69	23.74±0.87	32.05±0.06	28.13±0.28	35.49±0.69	31.90±0.35
Adapt-∞ (ICLR 2025)	33.61±0.83	24.93±0.42	34.00±0.57	28.51±1.54	35.36±0.82	32.37±0.63
OASIS (Ours)	36.04±0.26	30.78±0.84	38.72±0.62	33.10±0.43	41.13±0.91	37.36±0.55

Table 11: **Quantitative comparison between online sample selection methods on Adapt benchmark.** We use Qwen-VL-2.5-7B as the MLLM. Bold indicates the highest performance; underlined results are within the 0.05 t-test significance level. 'Full-Data Training' uses all data without selection.

A.5 EXPERIMENTS WITH QWEN-VL-2.5

In addition to experiments with LLaVA-1.5-7B (Liu et al., 2024a) in Sec. 4.2, we also compare OASIS with sample selection baselines using Qwen-VL-2.5-7B (Bai et al., 2025) as the MLLM. Specifically, we compare OASIS with baselines on COAST and Adapt, and summarize the results in Tab. 10 and Tab. 11 respectively. As shown in the tables, OASIS outperforms the baselines, consistent with the results observed using LLaVA-1.5-7B.

A.6 DETAILS ON DETERMINING I_T

As described in Sec. 3.3, we probabilistically select a sample based on its relative informativeness \hat{I} , with the selection probability defined as:

$$p(\hat{I}) = \sigma(\hat{I} - I_T), \quad (14)$$

where $\sigma(\cdot)$ denotes the sigmoid function and I_T is a predefined threshold. We assume $\hat{I} \sim \mathcal{N}(0, 1)$ that the original (*i.e.*, pre-normalized) informativeness I approximately follows a normal distribution. Based on this normal distribution assumption, the expected selection rate across the distribution is:

$$f(I_T) = \mathbb{E}_{\hat{I} \sim \mathcal{N}(0,1)}[\sigma(\hat{I} - I_T)] \quad (15)$$

$$= \int_{-\infty}^{\infty} \sigma(\hat{I} - I_T) \cdot \phi(\hat{I}) d\hat{I}, \quad (16)$$

Method	Model Scale					
	LLaVA-1.5-1B		LLaVA-1.5-3B		LLaVA-1.5-13B	
	$A_{avg} \uparrow$	$A_{last} \uparrow$	$A_{avg} \uparrow$	$A_{last} \uparrow$	$A_{avg} \uparrow$	$A_{last} \uparrow$
Random	20.05±0.21	28.04±0.15	21.20±0.56	29.30±0.22	25.39±0.90	33.84±0.72
GradNorm (ICML 2018)	19.13±0.33	27.49±0.11	19.95±0.80	28.73±0.39	23.38±0.56	32.03±0.51
Self-Sup (NeurIPS 2022)	18.17±0.41	26.36±0.20	17.49±0.86	24.76±0.80	21.47±0.37	29.35±0.22
COINCIDE (EMNLP 2024)	19.46±1.17	28.27±0.09	20.11±0.16	28.74±0.90	23.77±0.79	31.74±0.27
DBP (ICLR 2024)	17.98±0.14	25.57±0.06	17.05±0.69	25.71±1.83	22.24±0.22	31.95±0.62
InfoBatch (ICLR 2024)	18.36±0.18	25.47±0.04	19.92±0.89	29.27±0.56	24.05±0.73	30.48±0.64
DivBS (ICML 2024)	21.21±0.24	28.52±0.13	21.39±0.35	31.41±0.80	26.22±0.10	35.41±0.10
TIVE (arXiv:2403)	18.04±0.25	26.35±0.30	17.47±0.74	26.11±0.62	<u>24.58±0.34</u>	<u>32.86±0.60</u>
Adapt-∞ (ICLR 2025)	19.91±0.16	26.81±0.28	18.69±0.61	28.00±0.58	22.52±0.61	31.20±0.21
OASIS (Ours)	23.03±0.18	30.37±0.12	23.08±0.17	32.84±0.13	26.52±0.52	36.81±0.65

Table 12: **Quantitative comparison between online sample selection methods across model scales on COAST benchmark under selection ratio 6.25%**. Bold indicates the highest performance; underlined results are within the 0.05 t-test significance level.

Method	MICVIT		COAST	
	$A_{avg} \uparrow$	$A_{last} \uparrow$	$A_{avg} \uparrow$	$A_{last} \uparrow$
Random	55.66±0.36	57.97±0.61	<u>17.64±0.50</u>	<u>23.24±1.44</u>
GradNorm	54.39±1.59	56.90±0.41	14.94±0.79	22.54±0.31
Self-Sup	53.20±0.71	55.95±0.58	15.22±0.64	19.77±0.30
COINCIDE	55.18±0.16	58.32±0.61	13.87±0.50	21.67±0.15
DBP	53.43±0.63	57.25±1.50	14.20±0.28	18.82±0.20
InfoBatch	54.90±0.46	54.30±0.54	<u>16.74±0.52</u>	<u>23.54±0.35</u>
DivBS	55.52±0.12	<u>57.22±0.17</u>	<u>17.35±0.68</u>	<u>23.57±0.15</u>
TIVE	52.15±1.06	54.12±0.59	15.33±0.22	<u>22.75±1.30</u>
Adapt-∞	54.89±0.53	55.65±0.91	<u>17.58±0.51</u>	<u>22.29±0.19</u>
OASIS (Ours)	57.66±0.17	59.84±0.64	19.02±0.71	25.08±0.59

Table 13: **Quantitative comparison between online sample selection methods on MICVIT and COAST benchmark under selection ratio 6.25% at QwenVL-0.5B**. Bold indicates the highest performance; underlined results are within the 0.05 t-test significance level.

where $\phi(z)$ denotes the standard normal probability density function.

Given target selection rate $r \in [0, 1]$ (e.g., 0.125), we find I_T such that $f(I_T) = r$. Since the integral has no closed form, we approximate it numerically using a Riemann sum and solve:

$$I_T = \arg \min_{I_T} (f(I_T) - r)^2. \quad (17)$$

Applying Eq. 17, selection ratios r of 0.0625, 0.125, and 0.25 yield corresponding I_T values of 2.06, 1.53, and 0.89, respectively. However, data distribution variations and the probabilistic selection process may cause the actual proportion of selected samples to deviate from the target ratio. To ensure fair comparison with baseline methods, we fine-tune I_T to achieve sample counts that closely approximate but remain marginally below the target selection ratio.

A.7 EXPERIMENTS ON VARIOUS MODEL SCALES

We further compare baselines with various sizes of LLaVAs (i.e., LLaVA-1.5-1B, 3B, and 13B) in Tab. 12 and Qwen-VL-2.5-0.5B in Tab. 13. As shown in the tables, OASIS consistently outperforms the baselines across architectures and model sizes by a significant margin, demonstrating its general applicability.

A.8 DIFFERENT INFORMATION METRICS

We compare OASIS using Fisher Information to other metrics for measuring the informativeness I of samples, including Entropy (Coleman et al., 2020), Perplexity (Li et al., 2023b), and EL2N (Paul et al., 2021). As shown in Tab. 14, using sample-wise FI as our Informativeness outperforms other metrics.

Method	MICVIT		COAST	
	$A_{last} \uparrow$	$A_{avg} \uparrow$	$A_{last} \uparrow$	$A_{avg} \uparrow$
Entropy	<u>62.75±0.84</u>	<u>70.08±0.99</u>	22.53±0.46	30.42±0.57
Perplexity	59.86±1.37	64.49±0.72	20.01±0.73	28.64±0.82
EL2N	60.36±1.44	67.15±0.81	23.28±0.56	31.81±0.60
FI (Ours)	64.39±0.58	71.26±0.72	25.67±0.35	34.23±0.38

Table 14: **Comparison of different information metrics.** On MICVIT and COAST benchmarks, we apply different existing information metrics in place of Fisher Information (FI) to OASIS, when selection ratio is 6.25%. Bold indicates the highest performance, and underlined results are within the 0.05 significance level of the t-test.

A.9 EFFECT OF EMA RATIO β

We further evaluate the effect of EMA ratio β on OASIS, and summarize the results in Tab. 15. As shown, extremely large or small values of β degrade performance: high β fails to capture the informativeness of past batches, while low β overemphasizes outdated informativeness. To balance this trade-off, we use a moderate value of $\beta = 0.9$ across all benchmarks and selection ratios.

EMA Ratio β	$A_{avg} \uparrow$	$A_{last} \uparrow$
0.7	22.24±0.25	31.93±0.75
0.9	24.36±0.59	33.28±0.26
0.99	22.80±0.42	32.73±0.64
0.999	23.05±0.57	32.94±0.81

Table 15: **Ablation on EMA ratio.**

A.10 ACCURACY OVER TIME

We compare the average accuracy of seen tasks at several training time points between our method and baselines. Specifically, we evaluate performance across different selection ratios and benchmarks, with results shown in Fig.5 and Fig.6, respectively. As shown in the figures, our method outperforms baselines at all time points, demonstrating that its superior performance is not limited to specific intervals but is consistently maintained throughout the entire CIT process.

A.11 COMPARISON OF COMPUTATIONAL COSTS

We compare the computational cost of OASIS with selection baselines in Tab. 16. As shown, all methods require forward passes over candidate samples. Sample selection methods (*e.g.*, Self-Sup and COINCIDE) that rely solely on features do not incur additional backward pass costs. In contrast, gradient-based methods (*e.g.*, OASIS and GradNorm) involve backward passes to compute sample-wise gradients. However, unlike TIVE and Adapt- ∞ , which introduce substantial overhead due




Figure 5: **Average accuracies over time on COAST benchmark under different selection ratios.** Performance at task t denotes the average accuracy over all seen tasks up to that point (*i.e.*, task 1 through task t). We use LLaVA-1.5-7B as the model across all selection ratios.




Figure 6: **Average accuracy over time across benchmarks.** Performance at task t denotes the average accuracy over all seen tasks up to that point (*i.e.*, task 1 through task t). We use a selection ratio of 6.25% across all benchmarks. We use LLaVA-1.5-7B as the model.

to full-layer gradient computations or intermediate layer calculations, OASIS leverages only the last-layer gradients, resulting in negligible overhead compared to the forward pass (*i.e.*, less than 3.5%).

Methods	Overhead Type	Relative \mathcal{C} to OASIS
GradNorm (ICML 2018)	Forward Pass + Last Layer Gradient Compute	1.000
Self-Sup (NeurIPS 2022)	Forward Pass	0.976
COINCIDE (EMNLP 2024)	Forward Pass	0.976
DBP (ICLR 2024)	Forward Pass	0.976
InfoBatch (ICLR 2024)	Forward Pass	0.976
DivBS (ICML 2024)	Forward Pass + Last Layer Gradient Compute	1.000
TIVE (arXiv:2403)	Forward Pass + Full Layer Gradient Compute	2.038
Adapt- ∞ (ICLR 2025)	Forward Pass + Middle Layer Gradient Compute	1.507
OASIS (Ours)	Forward Pass + Last Layer Gradient Compute	1.000

Table 16: **Comparison of computational costs.** OASIS incurs additional computational cost compared to forward-only baselines (*e.g.*, Self-Sup), but only by approximately 3.4%.

A.12 DISTRIBUTION OF DATA SELECTED BY SELECTION BASELINES

The selection baselines that consistently underperform random selection across all selection ratios, namely Self-Sup, COINCIDE, DBP, TIVE, and Adapt- ∞ , rely on K-means clustering, which often yields imbalanced clusters (*i.e.*, certain clusters contain disproportionately more data) under imbalanced distributions. Note that in continual learning, data arrives in a streaming manner, making the temporal distribution at each time point likely imbalanced, even if the overall dataset is balanced across tasks. As a result, these methods often lead to skewed data selection, over-sampling data from

certain tasks while under-sampling others. To quantify this skewness, we measure the difference in the number of selected samples between the most-sampled and least-sampled tasks for each baseline, and summarize the results in Tab. 17. As shown, these baselines exhibit significantly higher skewness than our proposed OASIS, even higher than random selection, ultimately contributing to their lower performance. In contrast, both OASIS and DivBS, the best-performing baseline, achieve much lower selection skewness, indicating a more balanced sample selection across tasks.

Method	K-means based selection	Max-Min # of selected samples across tasks
Self-Sup	O	658
COINCIDE	O	659
DBP	O	687
TIVE	O	929
Adapt- ∞	O	746
Random	X	600
DivBS	X	520
OASIS (Ours)	X	482

Table 17: **Comparison of data selection distribution.** Comparison of methods with respect to sample variability across tasks.

A.13 COMPARISON OF THE NUMBER OF SELECTED SAMPLES

While baselines select a fixed number of samples per batch, thus selecting the same number of total samples, OASIS dynamically selects samples probabilistically. For fair comparison, we ensure our approach uses comparable or fewer total samples than those selected by other baselines. We summarize the selected samples for each baseline in Tab. 18. Despite using fewer samples, our method outperforms the baselines, as demonstrated in Sec. 4.2.

Method	MICVIT			COAST			Adapt		
	6.25%	12.5%	25.0%	6.25%	12.5%	25.0%	6.25%	12.5%	25.0%
Baselines	43394	86788	173576	10000	20000	40000	40000	80000	160000
OASIS (Ours)	43158	86251	173242	9874	19772	39427	39589	79604	159154

Table 18: **Comparison of the number of selected samples.** We compare the average number of selected samples by OASIS across three different seeds with those selected by other sample selection baselines.

A.14 DETAILED ALGORITHM OF OASIS

We provide a comprehensive pseudocode of OASIS in Algorithm 1.

A.15 IMPACT STATEMENTS

In this work, we focus on advancing an online sample selection algorithm for continual instruction tuning. Our approach prioritizes training efficiency while maintaining strong performance, allowing large pre-trained MLLMs to adapt more effectively and make fairer, less biased decisions when processing continuous streams of real-world data.

A.16 DATA PRIVACY AND CONTENT SENSITIVITY

For CIT benchmarks, we use existing datasets such as COAST and Adapt, which have already filtered out privacy-sensitive content, as well as MICVIT, a combination of existing multimodal benchmarks that are also free from sensitive content.

A.17 POTENTIAL RISKS

In CIT setup, real-time data arrives continuously in a streaming manner, leading to imbalanced data distributions at each time step. This can unintentionally introduce bias throughout training.

Algorithm 1 OASIS

```

1: Input: model  $f_\theta$ , batch  $\mathcal{B}_t$ , batch size  $N_{\mathcal{B}}$ , EMA  $\mu_t$ , EMV  $\sigma_t$ , threshold  $I_T$ , number of layers  $L$ 
2: Initialize: Informativeness set  $I^{(t)} \leftarrow \emptyset$ , high informative sample set  $H \leftarrow \emptyset$ 
3: // Stage 1: Calculate Informativeness  $I^{(t)}$  I of each sample in  $\mathcal{B}_t$ 
4: for each  $(x_i^{(t)}, y_i^{(t)}) \in \mathcal{B}_t$  do
5:   Calculate gradient  $g_i \leftarrow \frac{\partial \ell(d_i^{(t)})}{\partial \theta^L}$ 
6:   Calculate Informativeness  $I_i^{(t)} \leftarrow \text{tr}(g_i \cdot g_i^\top)$ 
7:   Add to set  $I^{(t)} \leftarrow I^{(t)} \cup \{I_i^{(t)}\}$ 
8: end for
9: // Stage 2: SIREN (Similarity-aware Information Redundancy Elimination)
10: Initialize adjusted Informativeness  $\tilde{I}^{(t)} \leftarrow I^{(t)}$ 
11: while  $|H| < N_{\mathcal{B}}$  do
12:   Add most informative sample  $H \leftarrow H \cup \left\{ \arg \max_{d_i^{(t)} \in \mathcal{B}_t \setminus H} \tilde{I}_i^{(t)} \right\}$ 
13:   for each  $d_i^{(t)} \in \mathcal{B}_t \setminus H$  do
14:     for each  $d_h^{(t)} \in H$  do
15:       Calculate updated Informativeness  $\tilde{I}_i^{(t)} \leftarrow I_i^{(t)} - \cos(g_i, g_h) \cdot I_h^{(t)}$ 
16:     end for
17:     if  $|H| > 1$  then
18:       Account higher-order redundancy due to overlapping similarities between  $d_h^{(t)} \in H$ 
19:        $\tilde{I}_i^{(t)} = \tilde{I}_i^{(t)} + \sum_{\substack{U \subseteq H \\ |U| \geq 2}} (-1)^{|U|} \cos(g_i, \bar{g}_U) \cdot \bar{I}_U^{(t)}$ 
20:     end if
21:   end for
22: end while
23: // Stage 3: Calculate Relative Informativeness  $\hat{I}^{(t)}$ 
24: for each  $\tilde{I}_i^{(t)}$  in  $\tilde{I}^{(t)}$  do
25:   Calculate Relative Informativeness  $\hat{I}_i^{(t)} \leftarrow \frac{\tilde{I}_i^{(t)} - \mu_t}{\sigma_t}$ 
26: end for
27: // Stage 4: Probabilistic Sampling
28: Sample random threshold from uniform distribution  $r \sim \mathcal{U}(0, 1)$ 
29: Select samples  $\mathcal{B}_t^* \leftarrow \{(x_i^{(t)}, y_i^{(t)}) \in \mathcal{B}_t \mid \sigma(\hat{I}_i^{(t)} - I_T^{(t)}) > r\}$ 
30: Calculate average Informativeness  $\bar{I}^{(t)} = \frac{1}{N_{\mathcal{B}}} \sum_{i=1}^{N_{\mathcal{B}}} I_i^{(t)}$ 
31: Update EMA  $\mu_t = \beta \bar{I}^{(t)} + (1 - \beta) \mu_{t-1}$ 
32: Update EMV  $\sigma_t = \beta (\bar{I}^{(t)} - \mu_{t-1})^2 + (1 - \beta) \sigma_{t-1}$ 
33: Output: selected samples  $\mathcal{B}_t^*$ 

```

A.18 PARAMETERS FOR PACKAGES

For evaluation, we measure accuracy, which selects an answer among multiple candidate choices. In ORIS, we normalize the informativeness score I using Z-score normalization based on our empirical evidence from the QQ-plot.

A.19 LICENSE FOR ARTIFACTS

We utilize publicly available data, models, and codebases, as provided by the original papers for each baseline method.

A.20 USE OF AI ASSISTANCE

We use AI assistance, such as GPT-4, solely for grammatical error corrections.