LEARNING TO ACQUIRE VIA LOOK-AHEAD LOSS ESTIMATION FOR ACTIVE LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

A common strategy in (deep) active learning is to acquire samples that the current model finds most challenging, using signals such as uncertainty or loss. This introduces a myopic bias in acquisition, as it relies on the assumption that samples identified as challenging at acquisition time remain challenging throughout subsequent optimization within the same acquisition round. Across image, text, and audio classification tasks, we find substantial within-round temporal misalignment: agreement between current and future per-sample loss rankings degrades rapidly as training epochs increase, so samples that appear difficult under the current model may become less challenging after subsequent training epochs. In this work, we present AHEAD (Active learning with Horizon-aware Estimation of Anticipated Difficulty), a framework that learns sample difficulty from estimated future loss rather than observed current loss. Concretely, AHEAD predicts each sample's relative loss T epochs ahead, providing a future-oriented supervision signal that avoids relying on outdated loss observations and improves robustness to rapidly evolving training dynamics. The predicted future difficulty serves as a consistent target for training a representation-based importance function that reweights labeled samples during training. Because its score depends only on sample representations, it can be directly applied to unlabeled data as a label-free acquisition function. The result is a single mechanism for emphasizing samples that remain difficult during subsequent optimization rather than those that are momentarily challenging. Across six datasets spanning three modalities, AHEAD consistently outperforms strong uncertainty-, diversity-, and learning-based active learning baselines.

1 INTRODUCTION

Active learning (AL) reduces annotation costs by selecting unlabeled samples whose labels are expected to provide the greatest benefit to the learner. Given a pretrained learner f_θ and a large unlabeled pool $\mathcal{U}$, at each acquisition round the learner scores the samples in $\mathcal{U}$ with an acquisition function and queries labels for the top-scoring batch. This technique is employed in a wide range of machine learning applications where labels are expensive, including medical imaging (Budd et al., 2021), speech recognition (Riccardi & Hakkani-Tur, 2005), and scientific discovery (Lookman et al., 2019; Ding et al., 2025; 2026; Malik et al., 2026). The acquisition function is the central design choice, and prior work offers three main families: uncertainty-based methods score samples by the model's predictive uncertainty (Gal et al., 2017; Beluch et al., 2018), diversity-based methods score batches by their coverage of the data distribution (Sener & Savarese, 2018; Ash et al., 2020), and learning-based methods learn acquisition criteria themselves, including loss prediction (Yoo & Kweon, 2019), bilevel optimization against held-out validation objectives (Killamsetty et al., 2021), and learned acquisition policies (Konyushkova et al., 2017; Ding et al., 2024; Wang et al., 2025).

For all three families, acquisition decisions are made based on the learner's current state. However, in deep active learning, the learner is further trained on the expanded labeled dataset after each batch is acquired, during which the difficulty of individual samples can change substantially. Thus, a sample that appears difficult or informative at acquisition time may become easy after further optimization, while samples that remain difficult may not initially be prioritized. Previous studies have shown that samples may alternate between being learned and forgotten (Toneva et al., 2019), and that example difficulty evolves throughout optimization (Baldock et al., 2021). Indeed, we observe

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

substantial temporal misalignment during optimization: the ranking of samples by current loss, a proxy for training difficulty, differs from their ranking after T training epochs across image, text, and audio classification tasks (Section 3.2). In this regime, acquisition based on current difficulty can be misaligned with future optimization needs: a sample with high current loss may become less challenging after subsequent optimization, while samples that remain difficult after further optimization may not be prioritized by the current acquisition criterion.

In this work, we study AL under this temporal misalignment and propose to score samples by their predicted future difficulty after subsequent optimization rather than their observed current difficulty. Concretely, we train a predictor to estimate each sample's relative loss T epochs ahead, and use this look-ahead signal as the supervision target for a representation-based importance function. The importance function reweights learner updates during training and, since it takes only sample representations as input, can also score unlabeled samples during acquisition. Our method differs from loss-prediction methods (Yoo & Kweon, 2019), which estimate current loss, whereas we predict future loss to align the importance signal with a later optimization state. It also differs from learning-based acquisition methods such as AutoAL (Wang et al., 2025), which optimize acquisition functions based on their effect on learner performance. Instead, we formulate importance learning as a future-difficulty prediction task and supervise the function using loss estimates after T epochs of optimization, shifting the scoring quantity from current difficulty to anticipated future difficulty.

Contributions. We make four contributions:

1. We formulate sample importance as an *adversarially* learned weighting function and identify a temporal mismatch when learning this function from current loss: the learned weights shape subsequent learner updates, while their supervision is based on losses measured at the current learner state, before those updates (Section 3.1).
2. We address this mismatch with look-ahead supervision, using a pairwise-ranking predictor to estimate each sample's relative loss T epochs ahead and using the predicted future loss as the target for the importance function (Section 3.2).
3. We show that the resulting importance function serves not only to reweight training updates but also as a label-free acquisition function, unifying training-time importance weighting and active acquisition within a single mechanism (Section 3.3).
4. We evaluate the framework on six classification datasets spanning image, text, and audio modalities, showing improvements over representative uncertainty-, diversity-, and learning-based AL methods and isolating the contribution of look-ahead supervision through ablations (Section 4).

2 Related Work

Our work builds on several related lines of work in active learning, data valuation, and training dynamics. We highlight the closest connections here.

Model-based acquisition criteria in AL. A large body of work designs acquisition functions from the current learner state: uncertainty-based methods query samples with high predictive uncertainty (Gal et al., 2017; Beluch et al., 2018), diversity-based methods select batches that cover the representation space (Sener & Savarese, 2017; Sinha et al., 2019), and hybrid criteria combine both signals via gradient embeddings (Ash et al., 2020). At scale, recent methods extend acquisition to large batches and web-scale unlabeled pools (Citovsky et al., 2021; Coleman et al., 2022), while stochastic acquisition strategies introduce randomness into batch selection to reduce dependence on a single ranking (Kirsch et al., 2023). However, these criteria evaluate sample importance using the learner state available at the time of selection, without explicitly modeling how sample difficulty evolves during subsequent optimization. Our approach instead scores samples using predicted future difficulty T epochs ahead rather than difficulty observed at selection time.

Learning-based acquisition and data valuation. A second line of work is to learns the acquisition criterion itself, including error-reduction prediction (Konyushkova et al., 2017), loss prediction (Yoo & Kweon, 2019), bilevel subset selection (Killamsetty et al., 2021), learned ranking

policies (Ding et al., 2024), and differentiable acquisition optimization (Wang et al., 2025). Regarding our motivation, Mindermann et al. (2022) accounts for training difficulty by discounting the irreducible loss relative to a reference model. However, these learning-based methods optimize acquisition criteria using task-level objectives rather than explicitly predicting how the difficulty of individual samples evolves under future optimization. A related perspective comes from data valuation methods, which assess samples by their contribution to the final model performance rather than their current utility. Datamodel approaches estimate such effects through training-subset regressions (Ilyas et al., 2022) and extensions to other settings (Dass et al., 2026). However, these approaches require offline computation over labeled subsets and are substantially more expensive than online acquisition. In contrast, we supervise the acquisition criterion with each sample's predicted loss T epochs ahead. Since the learned score depends only on sample representations, the same mechanism reweights labeled samples during training and ranks the unlabeled pool during acquisition, without requiring labels for acquisition scoring or an additional held-out validation set.

Sample reweighting during training. Complementary to acquisition, reweighting methods adjust the contribution of samples within the labeled set during training. Fixed heuristics derive sample weights from current training signals: focal loss upweights hard samples (Lin et al., 2017), importance sampling weights samples based on loss or gradient norm (Katharopoulos & Fleuret, 2018), and online batch selection prioritizes high-loss samples observed during training (Loshchilov & Hutter, 2015). Learned schemes meta-optimize the weighting function against a held-out validation set (Ren et al., 2018; Shu et al., 2019). However, these methods generally define sample importance using signals derived from the current learner state, without explicitly modeling how sample difficulty evolves during future optimization. By comparison, our weighting function is supervised using a future-loss signal, aligning its supervision with the learner state T epochs ahead.

Training dynamics and example difficulty. A separate line of work studies how per-sample training difficulty evolves during optimization: samples can flip between learned and forgotten (Toneva et al., 2019), examples are learned at different depths and stages of training (Baldock et al., 2021), and early-training scores can identify examples that remain influential later in training (Paul et al., 2021). These works collectively establish that sample difficulty evolves over training and that early training signals carry information about later importance. Their primary goal, however, is to characterize training dynamics rather than to use future difficulty explicitly as an optimization target. We build on this observation by predicting future difficulty from the current learner state and using that prediction both to supervise training-time reweighting and to score samples for active acquisition.

3 METHODOLOGY

We develop our method in two steps, each addressing a limitation of existing approaches. We first learn a sample-wise importance function through adversarial reweighting (Section 3.1); unlike existing reweighting schemes that rely on current-state signals, we identify and resolve the temporal mismatch in the supervision signal used to learn the importance function (Section 3.2). We then show that the resulting importance score, which depends only on sample representations, directly yields a label-free acquisition function (Section 3.3), connecting training-time reweighting and active acquisition within a single framework. Future difficulty is not intended to measure intrinsic sample value directly. Instead, we use future loss as an optimization-aligned proxy, motivated by the observation that samples remaining difficult after subsequent updates are precisely those whose supervision continues to affect the learner.

3.1 LEARNABLE IMPORTANCE VIA ADVERSARIAL REWEIGHTING

Let f_θ denote a classifier with feature extractor h_θ. For a labeled example (x_i, y_i), we write its representation and task loss at training time t as

$$z_i^t = h_{\theta^t}(x_i), \qquad \ell_i^t = \mathcal{L}(f_{\theta^t}(x_i), y_i)\,. \tag{1}$$

Why learn the weighting function? A natural approach to sample reweighting is to derive weights from the current loss: focal loss emphasizes hard samples (Lin et al., 2017), importance sampling weights samples by loss or gradient norm (Katharopoulos & Fleuret, 2018), and online

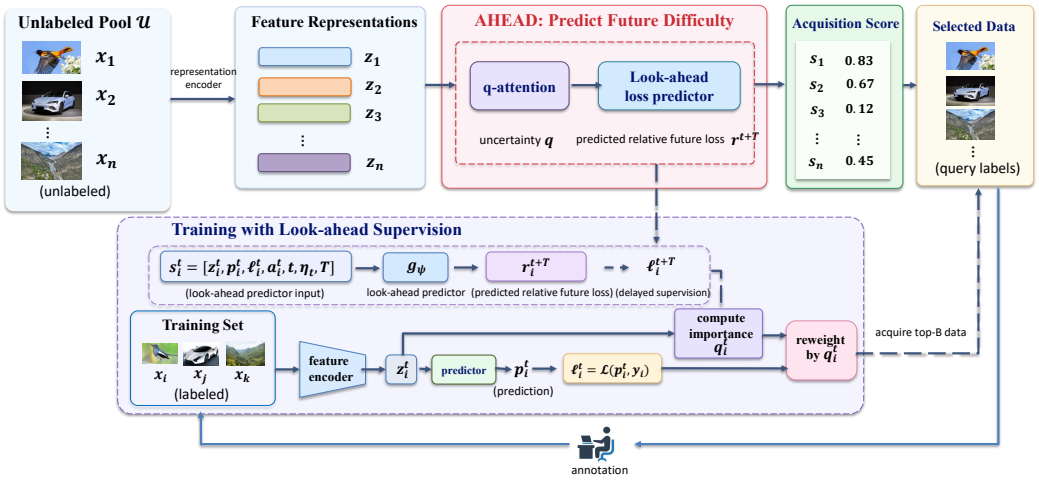

Figure 1: Overall Architecture of AHEAD. The look-ahead loss predictor estimates each sample's relative future difficulty from its current training state. The predicted future loss supervises the importance function, which reweights labeled samples during training and directly provides a label-free acquisition score for the unlabeled pool.

batch selection prioritizes recently high-loss samples (Loshchilov & Hutter, 2016). However, fixed loss transformations cannot adapt their weighting strategy across different training stages and do not account for how the weighting decision influences subsequent optimization. A second line of work thus *learns* the weighting function, typically by meta-optimizing weights against a clean validation set (Ren et al., 2018; Shu et al., 2019). In our setting, note that no such clean reference set is available: the labeled pool is small by construction, and holding out part of it for meta-validation would directly consume annotation budget. For these reasons, we instead learn the importance function against the learner itself, following the adversarial optimization perspective adopted in prior acquisition learning methods (Wang et al., 2025).

Parameterization. We parameterize importance with a scoring network a_ϕ acting on sample representations, and normalize scores within each mini-batch $\mathcal{B}$:

$$a_i^t = a_\phi(z_i^t), \qquad q_i^t = \frac{\exp(a_i^t/\tau_q)}{\sum_{j\in\mathcal{B}} \exp(a_j^t/\tau_q)}, \tag{2}$$

so that $q^t \in \Delta^{|\mathcal{B}|-1}$ defines a distribution of training importance over the batch. Directly minimizing the weighted task loss with respect to ϕ would drive the importance function to suppress high-loss samples. We instead optimize the learner and importance function in opposite directions:

$$\min_\theta \max_\phi \; \mathcal{J}(\theta, \phi) = \sum_{i\in\mathcal{B}} q_i \ell_i + \tau_q H(q), \qquad H(q) = -\sum_{i\in\mathcal{B}} q_i \log q_i. \tag{3}$$

The entropy term prevents the weighting distribution from collapsing prematurely onto a small number of samples. The following proposition characterizes the inner maximization when the weights are free variables.

Proposition 3.1 (Entropy-regularized sample weighting). *For a fixed difficulty vector $d \in \mathbb{R}^{|\mathcal{B}|}$ and temperature $\tau_q > 0$, consider the free-weight optimization problem over the probability simplex,*

$$\max_{q\in\Delta^{|\mathcal{B}|-1}} q^\top d + \tau_q H(q). \tag{4}$$

Its unique solution is

$$q_i^\star(d) = \frac{\exp(d_i/\tau_q)}{\sum_j \exp(d_j/\tau_q)}. \tag{5}$$

Consequently, $d_i > d_j$ implies $q_i^\star > q_j^\star$ in the ordering sense, i.e., $d_i > d_j \Rightarrow q_i^\star > q_j^\star$.

Proposition 3.1 shows that the maximization problem admits a closed-form solution for a given difficulty vector. Therefore, the sample weights can be obtained either analytically through this solution or learned through a parametric importance function. In Section 4.2.2, we compare these two alternatives when the difficulty signal is provided by the proposed look-ahead loss prediction. The proof of proposition is provided in Appendix A.

3.2 Look-Ahead Importance Learning

Current loss is misaligned with future optimization. Which signal should supervise the importance function? Existing reweighting methods typically rely on information available at the current learner state, such as the current loss ℓ^t. Yet at learner state θ^t, the weights q^t are used to update the learner to θ^{t+1}, so the loss that reflects the consequence of the weighting decision, ℓ^{t+1}, is not yet available when q^t is determined:

$$\ell^t \longrightarrow q^t \longrightarrow \theta^{t+1} \longrightarrow \ell^{t+1}. \tag{6}$$

This temporal mismatch motivates us to replace the immediate supervision signal with future difficulty predicted T epochs ahead. The current loss therefore provides a supervision target that may not reflect future optimization difficulty: samples flip between learned and forgotten throughout training (Toneva et al., 2019), and per-sample difficulty evolves with the learner (Baldock et al., 2021). Empirically, we observe that sample difficulty rankings change substantially during optimization, motivating the need to predict future difficulty rather than rely on the current loss.

Related works address this mismatch from different perspectives. Yoo & Kweon (2019) trains a module to predict per-sample loss, but focuses on the current loss rather than future optimization difficulty. Mindermann et al. (2022) identifies samples that are not yet learned but learnable using a reference model; however, its criterion is still defined relative to a current model state and requires an auxiliary model. Killamsetty et al. (2021) selects subsets by optimizing validation loss after a one-step virtual update, using a held-out labeled set for evaluation. In contrast, we address this temporal misalignment by predicting each sample's loss T epochs ahead from its current state and supervising the importance function with this look-ahead signal. The feasibility of predicting future importance is supported by Paul et al. (2021), who show that early-training signals can identify samples that remain important later in training. We extend this idea by explicitly learning to predict future sample difficulty and using this prediction as the supervisory signal for importance learning.

Stabilized target generation. Both the learner and the future-loss predictor evolve during training, making a directly generated look-ahead target unstable. We therefore generate supervision targets using an EMA copy $\bar{\psi}$ of the predictor:

$$\bar{\psi} \leftarrow \rho\bar{\psi} + (1-\rho)\psi. \tag{7}$$

At the beginning of each look-ahead block, we use the EMA predictor to generate a standardized target $\tilde{r}_i^{\text{target}}$ that estimates each sample's relative difficulty T epochs ahead. This target is held fixed throughout the subsequent T-epoch block, while the online predictor continues to receive delayed supervision from realized future losses as they become available. Thus, only the target seen by the importance function is frozen. This separates the predictor's update timescale from the importance function's supervision timescale. Further stabilization details appear in Appendix C.

3.2.1 Relative Future-Loss Prediction

Direct regression of future loss values can be challenging in practice (Yoo & Kweon, 2019; Ding et al., 2024). Fortunately, our reweighting objective only requires the *relative* difficulty of samples, not calibrated loss values. We therefore predict relative future loss using a pairwise ranking objective, similar to RankNet (Burges et al., 2005).

For sample i at epoch t, we assemble a state vector and predict a relative score:

$$s_i^t = \left[z_i^t,\ p_i^t,\ \ell_i^t,\ a_i^t,\ \bar{t},\ \eta_t,\ T\right], \qquad r_i^{t+T} = g_\psi(s_i^t), \tag{8}$$

where p_i^t is the predictive probability vector, η_t is the learning rate, and $\bar{t}$ denotes normalized training time and prediction horizon. The score r_i^{t+T} is not required to match the numerical value of ℓ_i^{t+T}.

Table 1: Performance comparison between the backbone model and AHEAD across six datasets. AHEAD incorporates the proposed look-ahead loss prediction and reweighting mechanism. Accuracy is reported for image and text datasets, while Equal Error Rate (EER) is reported for ASVspoof2019. Mean and standard deviation over multiple runs are shown.

	CIFAR10	**CIFAR100**	**CINIC10**	**SVHN**	**AG News**	**ASVSpoof2019**
backbone	87.78±0.50	53.47±0.89	73.79±0.16	94.77±0.46	92.13±0.23	2.78±0.27
AHEAD	92.23±0.49	76.16±0.49	79.32±0.14	95.24±0.46	94.73±0.23	2.32±0.81

The current importance score a_i^t is included as a state feature, allowing the predictor to account for the learner's current weighting behavior.

Given the realized future losses, we normalize them within the corresponding snapshot within each look-ahead block and define the pairwise target and predicted logit as

$$u_i^{t+T} = \frac{\ell_i^{t+T} - \mu_\ell^{t+T}}{\sigma_\ell^{t+T} + \epsilon}, \qquad P_{ij}^{t+T} = \sigma\left(\frac{u_i^{t+T} - u_j^{t+T}}{\tau_y}\right), \qquad \widehat{\Delta}_{ij}^{t+T} = \frac{r_i^{t+T} - r_j^{t+T}}{\tau_p}. \tag{9}$$

The predictor is trained with

$$\mathcal{L}_{\text{pred}} = \frac{1}{|\mathcal{P}|} \sum_{(i,j)\in\mathcal{P}} \text{BCEWithLogits}\left(\widehat{\Delta}_{ij}^{t+T}, P_{ij}^{t+T}\right). \tag{10}$$

The soft target retains information about the relative loss gap while avoiding direct regression on future loss values. To avoid future information leakage, states recorded at epoch t are supervised only after their corresponding losses at $t+T$ become available.

3.2.2 Predictive Reweighting

Given the look-ahead prediction, we replace the current difficulty signal with the predicted future difficulty when updating the importance function. Let $\tilde{r}_i^{t+T}$ denote the standardized future score. The importance and learner objectives are

$$\mathcal{L}_q = -\sum_{i\in\mathcal{B}} q_i^t \tilde{r}_i^{\text{target}} - \tau_q H(q^t), \qquad \mathcal{L}_\theta = \sum_{i\in\mathcal{B}} \left(\lambda_q \text{sg}(q_i^t) + \lambda_{\text{base}}\right)\ell_i^t, \tag{11}$$

where $\text{sg}(\cdot)$ denotes stop-gradient, and λ_q and λ_{base} control the contributions of importance-weighted and uniform training, respectively. The two optimization steps alternate: a_ϕ is updated against the predicted future difficulty, and the resulting weights are detached and used to update the learner. The full training protocol can be seen in Appendix B.

3.3 From Reweighting to Active Acquisition

The importance function takes only sample representations as input, allowing it to score the unlabeled pool directly. We use its raw pre-softmax output as the acquisition score,

$$A_{\theta,\phi}(x) = a_\phi(h_\theta(x)), \tag{12}$$

and given unlabeled pool $\mathcal{U}_r$ and budget B, select

$$Q_r = \text{Top-}B_{x\in\mathcal{U}_r}\, A_{\theta_r,\phi_r}(x). \tag{13}$$

Two remarks. First, the acquisition function inherits the look-ahead supervision learned during training: a_ϕ is trained on labeled samples using predicted future difficulty as supervision, so its scores on the pool rank samples by *anticipated* rather than current difficulty, without requiring a single pool label. Second, we score with the raw output a_ϕ rather than the normalized weight q, since q depends on the other samples through mini-batch normalization and therefore cannot provide a standalone sample score.

3.4 Multi-Round Active Learning Protocol

We deploy the acquisition rule in a multi-round active learning protocol. Starting from the initial labeled set $\mathcal{L}_0$, we train a learner for K epochs before each acquisition. After querying and annotating $\mathcal{Q}_r$, we update $\mathcal{L}_{r+1} = \mathcal{L}_r \cup \mathcal{Q}_r$ and $\mathcal{U}_{r+1} = \mathcal{U}_r \setminus \mathcal{Q}_r$. At the beginning of each round, the learner and the look-ahead loss predictor are initialized and trained on the current labeled set. All look-ahead states and future-loss targets are constructed within the same round.

The look-ahead prediction objective is constructed from the training trajectory of the current round. Specifically, a state recorded at epoch t receives supervision from its corresponding loss at epoch $t + T$ once that loss becomes available. Since the labeled and unlabeled sets remain fixed during each training round, the recorded states and their future supervision remain aligned throughout the optimization process.

Each round therefore consists of look-ahead training followed by acquisition: future-loss prediction supervises the importance function, the importance function reweights learner updates, and its representation-based score selects the next batch of unlabeled samples.

4 Experimental Results

4.1 Experimental Setup

We evaluate our method on six datasets spanning three modalities: four image datasets (SVHN (Netzer et al., 2011), CINIC-10 (Darlow et al., 2018), CIFAR-10, and CIFAR-100 (Alex, 2009)), one text dataset (AG News (Zhang et al., 2015)), and one audio deepfake detection dataset (ASVspoof 2019 (Todisco et al., 2019)). For active learning evaluation, we initialize the labeled set with 500 samples for SVHN and 5,000 samples for the other datasets, acquire 200 and 2,000 samples per round respectively, and perform five acquisition rounds. We use representative backbones for each modality: ResNet-18 (He et al., 2016) trained from scratch with SGD for images, pretrained RoBERTa-base (Liu et al., 2019) for text, and randomly initialized AASIST (Jung et al., 2022) for audio. All models use cosine annealing learning rate schedules, with full optimization details provided in Appendix D.

We compare against five active learning baselines. Random selection provides a strategy-agnostic reference. CoreSet (Sener & Savarese, 2017) selects samples through k-center coverage of the representation space and evaluates the contribution of diversity-based acquisition. BADGE (Ash et al., 2020) combines uncertainty and diversity through gradient embeddings. GLISTER (Killamsetty et al., 2021) selects subsets by optimizing validation performance through bilevel optimization, providing a comparison to acquisition criteria based on future learner improvement. Learning Loss (Yoo & Kweon, 2019) learns an auxiliary module to predict the current loss of unlabeled samples and uses the predicted loss as the acquisition criterion. Finally, AutoAL (Wang et al., 2025) learns acquisition strategies through differentiable optimization and represents a recent learning-based acquisition approach.[1]

4.2 Main Results

4.2.1 Effectiveness of Look-ahead Supervision

We first evaluate whether the proposed look-ahead supervision provides an effective signal for learning sample importance across different modalities. Table 1 compares the backbone model trained with uniform weighting and our look-ahead supervised reweighting strategy across six datasets. Classification accuracy is reported for image and text datasets, while Equal Error Rate (EER) is reported for ASVspoof2019. The proposed look-ahead reweighting strategy achieves consistent improvements across all evaluated datasets, improving accuracy by 4.45, 22.69, 5.53, 0.47, and 2.60 percentage points on CIFAR10, CIFAR100, CINIC10, SVHN, and AG News, respectively. Meanwhile, it reduces EER on ASVSpoof2019 from 2.78% to 2.32%. These results demonstrate that

[1] We report AutoAL on the four image benchmarks; its differentiable acquisition procedure incurs substantially higher memory requirements and does not scale to our RoBERTa-base and AASIST settings within the same computational budget.

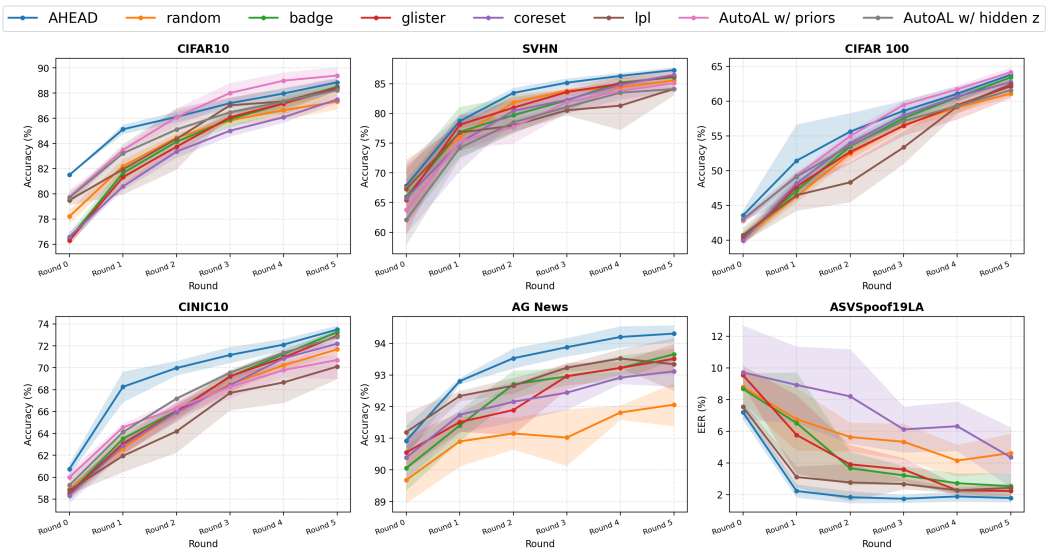

Figure 2: Active learning performance across six datasets spanning three modalities. Lower values indicate better performance on ASVspoof2019, where Equal Error Rate (EER) is the evaluation metric. AutoAL is omitted on AG News and ASVspoof2019 because it does not scale to the RoBERTa-base and AASIST settings within the same computational budget.

modeling the evolution of future loss provides a more informative signal of importance than relying solely on the current optimization state, and that the proposed look-ahead mechanism generalizes across vision, language, and audio tasks.

Furthermore, the effectiveness of our method across diverse backbone architectures demonstrates its architectural compatibility and plug-and-play property. In particular, RoBERTa-base, a pre-trained language model with a substantially different architecture from convolutional image models, achieves improved performance with our proposed predictor, demonstrating that the proposed mechanism extends beyond vision architectures. Meanwhile, AASIST is trained as an end-to-end classifier rather than being used as a fixed feature extractor. The successful integration with AASIST demonstrates that the proposed mechanism is compatible with end-to-end neural classifiers beyond conventional image architectures. Overall, these results verify that the proposed look-ahead mechanism generalizes across different modalities, including vision, language, and audio tasks.

4.2.2 Active Learning Performance

We further evaluate the effectiveness of the proposed method in active learning scenarios. Figure 2 illustrates the performance trajectories of different acquisition strategies over multiple acquisition rounds, where the shaded regions represent one standard deviation across different runs. The proposed method achieves competitive or improved performance across evaluated datasets, covering vision, text, and audio classification tasks.

Compared with conventional acquisition strategies, including Random, CoreSet (Sener & Savarese, 2017), BADGE (Ash et al., 2020), GLISTER (Killamsetty et al., 2021), LPL (Yoo & Kweon, 2019), and AutoAL (Wang et al., 2025), our method achieves faster performance improvements during the early acquisition stages. This advantage is particularly important under limited labeling budgets, where selecting more beneficial samples has a greater impact on model performance. For example, on CINIC10 and CIFAR10, our method achieves clear improvements over existing approaches in the initial acquisition rounds, indicating that the proposed look-ahead mechanism provides more effective guidance than strategies relying on current uncertainty, representation coverage, or acquisition models based on existing training information.

The proposed method also maintains stable improvements throughout the entire acquisition process and achieves the best final performance on most datasets. Specifically, our method obtains the high-

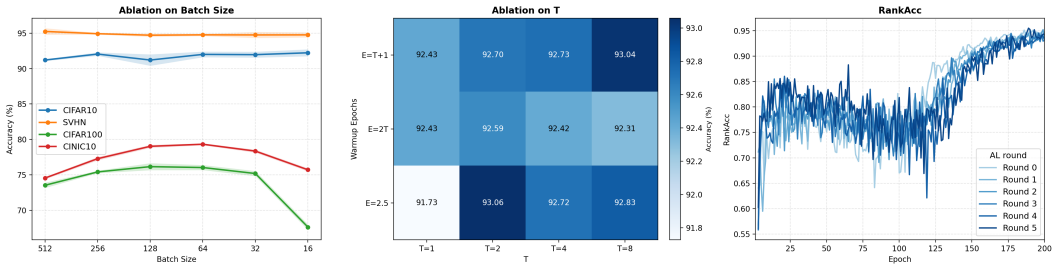

Figure 3: Ablation studies of AHEAD. The left and middle panels analyze the effects of acquisition batch size B, look-ahead horizon T, and warm-up length E, respectively. The right panel presents the RankNet predictor's training dynamics across acquisition rounds, measured by future-loss ranking accuracy.

est final accuracy on CIFAR10, SVHN, CIFAR100, CINIC10, and AG News. On ASVSpoof2019, where a lower EER indicates better detection performance, our method achieves the lowest error rate among all compared approaches. For AutoAL, we compare two variants that use prior information and hidden-state representations z. Our method consistently outperforms the hidden-state-based variant across all image datasets and achieves competitive performance with the prior-based variant. These results demonstrate that explicitly modeling the evolution of future loss provides more effective acquisition signals than relying solely on current model representations.

4.3 Ablation Studies

We investigate the influence of three factors in the proposed framework: acquisition batch size B, look-ahead horizon T, and warm-up length E. Figure 3 presents the results, where shaded regions indicate 95% confidence intervals across multiple runs.

The proposed method remains stable under different acquisition batch sizes, while the optimal granularity varies across datasets. For example, on CIFAR100, reducing B from 512 to 64 improves the final accuracy from 73.53% to 76.03%, whereas further reducing it to 16 degrades performance, suggesting that appropriate acquisition granularity is important for stable sample selection.

We further analyze the effect of T and E. Look-ahead prediction benefits from sufficient training trajectories: longer horizons remain effective when enough warm-up observations are available (e.g., $E \geq T + 1$), while increasing T alone does not consistently improve performance. The RankNet training dynamics on SVHN further show that the predictor reliably learns future-loss ordering, with ranking accuracy exceeding 70% for most epochs and approaching 95% before acquisition.

Finally, we compare the learned predictor with the analytical solution from Proposition 3.1. Without look-ahead supervision ($T = 0$), the average accuracy on CIFAR10 is 84.88%. With look-ahead supervision, the analytical solution achieves 91.37% and 91.68% accuracy for $T = 1$ and $T = 3$, respectively, while the learned predictor further improves these results to 92.08% and 92.23%. These results validate the effectiveness of look-ahead supervision and show that learning future loss dynamics provides additional benefits beyond the analytical weighting solution.

5 Conclusion

In this work, we identify a temporal mismatch in existing active learning and sample reweighting methods, where current difficulty is used to guide decisions that affect future training states. To address this issue, we propose AHEAD, a look-ahead learning framework that predicts future sample difficulty and uses it as supervision for learning importance functions. By modeling the evolution of loss dynamics, AHEAD provides a temporally aligned signal that unifies training-time reweighting and label-free active acquisition. Extensive experiments across vision, language, and audio classification tasks demonstrate the effectiveness and generality of future-aware sample selection, highlighting the importance of considering training dynamics when evaluating sample value.

AI USE STATEMENT

In this work, we used generative AI tools, including ChatGPT and Cursor, for language refinement of the manuscript and assistance in code implementation. We have not used generative AI tools for developing the research ideas, designing the proposed method, designing experiments, conducting theoretical analysis, generating experimental results, or preparing data, and these tasks are not applicable to AI-assisted disclosure in this work. Additionally, we used generative AI tools for improving the clarity of the paper title and manuscript presentation.

All AI-assisted outputs were carefully reviewed and verified by the authors. In particular, AI-assisted code was manually inspected line by line and tested to ensure correctness before being used in experiments. The authors take full responsibility for the final content of this work, including all text, claims, implementations, and experimental results.

ETHICS STATEMENT

This work studies active learning algorithms for improving the efficiency of data annotation in machine learning systems. We do not collect new human subject data or release new datasets. All experiments are conducted on publicly available benchmark datasets, which are widely used in the machine learning community and follow their respective data usage policies.

Our proposed method focuses on improving sample selection during model training and does not involve generating, modifying, or inferring sensitive personal information. However, as with other machine learning approaches, the effectiveness of active learning methods may depend on the characteristics and potential biases present in the underlying datasets and models. Therefore, the reported results should be interpreted within the scope of the evaluated datasets and tasks, and further evaluation may be required before deployment in safety-critical or high-impact applications.

We follow standard research practices by providing accurate descriptions of our methodology, experimental settings, and results. We use publicly available datasets and existing model architectures in accordance with their licenses and acknowledge all relevant prior work.

REPRODUCIBILITY STATEMENT

We provide detailed descriptions of the proposed framework, including the look-ahead loss prediction mechanism and the active learning acquisition strategy, in Section 3. The implementation details, experimental protocols, model configurations, and training settings are provided in Appendix D. All datasets used in our experiments are publicly available benchmarks, and their descriptions, preprocessing procedures, and evaluation protocols are provided in the corresponding sections and the supplementary materials. The supplementary materials include the complete training code, hyperparameter configurations, and experiment scripts required to reproduce our results.

REFERENCES

Krizhevsky Alex. Learning multiple layers of features from tiny images. *https://www. cs. toronto. edu/kriz/learning-features-2009-TR. pdf*, 2009.

Jordan T. Ash, Chicheng Zhang, Akshay Krishnamurthy, John Langford, and Alekh Agarwal. Deep batch active learning by diverse, uncertain gradient lower bounds. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=ryghZJBKPS.

Robert Baldock, Hartmut Maennel, and Behnam Neyshabur. Deep learning through the lens of example difficulty. *Advances in Neural Information Processing Systems*, 34:10876–10889, 2021.

William H Beluch, Tim Genewein, Andreas Nürnberger, and Jan M Köhler. The power of ensembles for active learning in image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 9368–9377, 2018.

Samuel Budd, Emma C Robinson, and Bernhard Kainz. A survey on active learning and human-in-the-loop deep learning for medical image analysis. *Medical image analysis*, 71:102062, 2021.

Chris Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Greg Hullender. Learning to rank using gradient descent. In *Proceedings of the 22nd international conference on Machine learning*, pp. 89–96, 2005.

Gui Citovsky, Giulia DeSalvo, Claudio Gentile, Lazaros Karydas, Anand Rajagopalan, Afshin Rostamizadeh, and Sanjiv Kumar. Batch active learning at scale. *Advances in Neural Information Processing Systems*, 34:11933–11944, 2021.

Cody Coleman, Edward Chou, Julian Katz-Samuels, Sean Culatana, Peter Bailis, Alexander C Berg, Robert Nowak, Roshan Sumbaly, Matei Zaharia, and I Zeki Yalniz. Similarity search for efficient active learning and search of rare concepts. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pp. 6402–6410, 2022.

Luke N Darlow, Elliot J Crowley, Antreas Antoniou, and Amos J Storkey. Cinic-10 is not imagenet or cifar-10. *arXiv preprint arXiv:1810.03505*, 2018.

Shivin Dass, Alaa Khaddaj, Logan Engstrom, Aleksander Madry, Andrew Ilyas, and Roberto Martín-Martín. DataMIL: Selecting data for robot imitation learning with datamodels. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=AcTsKglDdh.

Rui Ding, Jianguo Liu, Kang Hua, Xuebin Wang, Xiaoben Zhang, Minhua Shao, Yuxin Chen, and Junhong Chen. Leveraging data mining, active learning, and domain adaptation in a multi-stage, machine learning-driven approach for the efficient discovery of advanced acidic oxygen evolution electrocatalysts. *Science Advances*, April 2025. URL https://www.science.org/doi/10.1126/sciadv.adr9038.

Rui Ding, Zixin Ding, Rodrigo Pires Ferreira, Yuxin Chen, and Junhong Chen. Text-twin-translation (t3): A full-stack machine learning framework for functional material-device systems discovery. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2*, KDD '26, pp. 10831–10841, New York, NY, USA, 2026. Association for Computing Machinery. ISBN 9798400722592. doi: 10.1145/3770855.3819013. URL https://doi.org/10.1145/3770855.3819013.

Zixin Ding, Si Chen, Ruoxi Jia, and Yuxin Chen. Learning to rank for active learning via multi-task bilevel optimization. In *The 40th Conference on Uncertainty in Artificial Intelligence*, 2024. URL https://openreview.net/forum?id=o5Iw3kN9Eg.

Yarin Gal, Riashat Islam, and Zoubin Ghahramani. Deep bayesian active learning with image data. In *International conference on machine learning*, pp. 1183–1192. PMLR, 2017.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

Andrew Ilyas, Sung Min Park, Logan Engstrom, Guillaume Leclerc, and Aleksander Madry. Datamodels: Understanding predictions with data and data with predictions. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 9525–9587. PMLR, 17–23 Jul 2022. URL https://proceedings.mlr.press/v162/ilyas22a.html.

Jee-weon Jung, Hee-Soo Heo, Hemlata Tak, Hye-jin Shim, Joon Son Chung, Bong-Jin Lee, Ha-Jin Yu, and Nicholas Evans. Aasist: Audio anti-spoofing using integrated spectro-temporal graph attention networks. In *ICASSP 2022-2022 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pp. 6367–6371. IEEE, 2022.

Angelos Katharopoulos and François Fleuret. Not all samples are created equal: Deep learning with importance sampling. In *International conference on machine learning*, pp. 2525–2534. PMLR, 2018.

Krishnateja Killamsetty, Durga Sivasubramanian, Ganesh Ramakrishnan, and Rishabh Iyer. Glister: Generalization based data subset selection for efficient and robust learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pp. 8110–8118, 2021.

Andreas Kirsch, Sebastian Farquhar, Parmida Atighehchian, Andrew Jesson, Frédéric Branchaud-Charron, and Yarin Gal. Stochastic batch acquisition: A simple baseline for deep active learning. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL https://openreview.net/forum?id=vcHwQyNBjW. Expert Certification.

Ksenia Konyushkova, Raphael Sznitman, and Pascal Fua. Learning active learning from data. *Advances in neural information processing systems*, 30, 2017.

Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pp. 2980–2988, 2017.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.

Turab Lookman, Prasanna V Balachandran, Dezhen Xue, and Ruihao Yuan. Active learning in materials science with emphasis on adaptive sampling using uncertainties for targeted design. *npj Computational Materials*, 5(1):21, 2019.

Ilya Loshchilov and Frank Hutter. Online batch selection for faster training of neural networks. *arXiv preprint arXiv:1511.06343*, 2015.

Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016.

Shreshth A Malik, Tiarnan Doherty, Panagiotis Tigas, Muhammed Razzak, S Roberts, Aron Walsh, and Yarin Gal. MADE: Benchmark environments for closed-loop materials discovery. In *Forty-third International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=nrXxVDYMMF.

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 *International Conference on Machine Learning*, pp. 15630–15649. PMLR, 2022.

Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Baolin Wu, Andrew Y Ng, et al. Reading digits in natural images with unsupervised feature learning. In *NIPS workshop on deep learning and unsupervised feature learning*, volume 2011, pp. 4. Granada, 2011.

Mansheej Paul, Surya Ganguli, and Gintare Karolina Dziugaite. Deep learning on a data diet: Finding important examples early in training. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, 2021. URL https://openreview.net/forum?id=Uj7pF-D-YvT.

Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. In *International conference on machine learning*, pp. 4334–4343. PMLR, 2018.

Giuseppe Riccardi and Dilek Hakkani-Tur. Active learning: Theory and applications to automatic speech recognition. *IEEE transactions on speech and audio processing*, 13(4):504–511, 2005.

Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach. *arXiv preprint arXiv:1708.00489*, 2017.

Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach. In *International Conference on Learning Representations*, 2018. URL https://openreview.net/forum?id=H1aIuk-RW.

Jun Shu, Qi Xie, Lixuan Yi, Qian Zhao, Sanping Zhou, Zongben Xu, and Deyu Meng. Meta-weight-net: Learning an explicit mapping for sample weighting. *Advances in neural information processing systems*, 32, 2019.

Samarth Sinha, Sayna Ebrahimi, and Trevor Darrell. Variational adversarial active learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 5972–5981, 2019.

Massimiliano Todisco, Xin Wang, Ville Vestman, Md Sahidullah, Héctor Delgado, Andreas Nautsch, Junichi Yamagishi, Nicholas Evans, Tomi Kinnunen, and Kong Aik Lee. Asvspoof 2019: Future horizons in spoofed and fake audio detection. *arXiv preprint arXiv:1904.05441*, 2019.

Mariya Toneva, Alessandro Sordoni, Remi Tachet des Combes, Adam Trischler, Yoshua Bengio, and Geoffrey J. Gordon. An empirical study of example forgetting during deep neural network learning. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=BJlxm30cKm.

Yifeng Wang, Xueying Zhan, and Siyu Huang. AutoAL: Automated active learning with differentiable query strategy search. In *Forty-second International Conference on Machine Learning*, 2025. URL https://openreview.net/forum?id=TkUxFAbpfQ.

Donggeun Yoo and In So Kweon. Learning loss for active learning. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 93–102. IEEE, 2019.

Xiang Zhang, Junbo Zhao, and Yann LeCun. Character-level convolutional networks for text classification. *Advances in neural information processing systems*, 28, 2015.

A Proof of Proposition 3.1

We prove the proposition by applying the Karush-Kuhn-Tucker (KKT) conditions. Consider the entropy-regularized optimization problem:

$$\max_q \sum_i q_i d_i - \tau_q \sum_i q_i \log q_i \tag{14}$$

subject to the probability simplex constraint:

$$\sum_i q_i = 1, \qquad q_i \geq 0. \tag{15}$$

The Lagrangian is

$$\mathcal{L}(q, \lambda, \mu) = \sum_i q_i d_i - \tau_q \sum_i q_i \log q_i + \lambda(1 - \sum_i q_i) + \sum_i \mu_i q_i, \tag{16}$$

where λ is the multiplier for the normalization constraint and μ_i corresponds to the non-negativity constraints.

For an interior optimum with $q_i > 0$, complementary slackness gives $\mu_i = 0$. The stationarity condition becomes

$$\frac{\partial \mathcal{L}}{\partial q_i} = d_i - \tau_q(1 + \log q_i) - \lambda = 0. \tag{17}$$

Rearranging gives

$$\log q_i = \frac{d_i - \lambda - \tau_q}{\tau_q}, \tag{18}$$

and therefore

$$q_i = \exp\left(\frac{d_i}{\tau_q}\right) \exp\left(-\frac{\lambda + \tau_q}{\tau_q}\right). \tag{19}$$

The second term is a normalization constant shared by all samples. Applying the simplex constraint,

$$\sum_i q_i = 1, \tag{20}$$

we obtain

$$\exp\left(-\frac{\lambda + \tau_q}{\tau_q}\right) = \frac{1}{\sum_j \exp(d_j/\tau_q)}. \tag{21}$$

Substituting back yields the unique optimum:

$$q_i^\star(d) = \frac{\exp(d_i/\tau_q)}{\sum_j \exp(d_j/\tau_q)}. \tag{22}$$

Finally, because the exponential function is strictly increasing,

$$d_i > d_j \Rightarrow \exp(d_i/\tau_q) > \exp(d_j/\tau_q), \tag{23}$$

which implies

$$d_i > d_j \Rightarrow q_i^\star > q_j^\star. \tag{24}$$

The objective is strictly concave in q because the entropy term is strictly concave for $\tau_q > 0$. Therefore, the KKT solution is the unique global optimum.

B Training Protocol

The proposed framework jointly trains three components: the learner, the importance function, and the look-ahead loss predictor. We first warm up the learner for E epochs using uniform sample weights to collect initial training trajectories. During this phase, we record intermediate states and losses for constructing delayed future-loss supervision.

After warm-up, training proceeds online. At epoch t, the look-ahead predictor receives the current state s_i^t and estimates the relative future difficulty r_i^{t+T}. The predicted future difficulty supervises the update of the importance function, which produces sample weights for the subsequent learner update. The learner is then optimized using the detached importance weights.

The predictor receives delayed supervision: a state recorded at epoch t is used to update the predictor only after the corresponding loss at epoch $t + T$ becomes available. To stabilize this delayed target, we maintain an EMA copy of the predictor for target generation. This procedure separates the optimization timescale of future-loss prediction from the supervision timescale of importance learning. More details were shown in Algorithm 1.

C Stabilizing Look-Ahead Learning

Both the learner and future predictor evolve during training, so the target presented to the importance function may itself fluctuate. We use an EMA predictor, blockwise target updates, and a sample-wise trajectory regularizer to stabilize this temporal learning process.

Algorithm 1 Training Procedure of Look-Ahead Importance Learning

Require: warm-up epochs E, look-ahead horizon T

Initialize classifier parameters θ, importance function ϕ, predictor parameters ψ
Initialize EMA predictor $\bar{\psi} \leftarrow \psi$
Warm-up phase:
for $t = 1, \ldots, E$ **do**
    Train classifier with uniform weights $q_i = 1/|\mathcal{B}|$; Store states s_i^t and losses ℓ_i^t
end for
Look-ahead training:
while not converged **do**
    Compute representations z_i^t and importance scores a_i^t; Compute importance weights q_i^t;
    Predict future difficulty:

$$r_i^{t+T} = g_\psi(s_i^t)$$

    Update importance function q_i^t by maximizing predicted future difficulty
    Update classifier using detached weights:

$$(\lambda_q q_i^t + \lambda_{base})\ell_i^t$$

    Then the delayed losses ℓ_i^t become available. Update predictor ψ using cached states s_i^{t-T}
    Update EMA predictor:

$$\bar{\psi} \leftarrow \rho\bar{\psi} + (1-\rho)\psi$$

end while

Table 2: Dataset Details

Dataset	Training	Validating	Test	Type	Initial	Acquisition
CIFAR10	45000	5000	10000	Image	5000	2000
CIFAR100	45000	5000	10000	Image	5000	2000
CINIC10	90000	90000	90000	Image	5000	2000
SVHN	65932	7325	26032	Image	500	200
AG News	108000	12000	7600	Text	2500	1000
ASVSpoof2019 LA	25380	24844	71237	Audio	5000	2000

Definition C.1 (Blockwise look-ahead target)**.** Let t_b denote the first epoch of a look-ahead block. We generate its target using an EMA predictor $g_{\bar{\psi}}$,

$$\bar{\psi} \leftarrow \rho\bar{\psi} + (1-\rho)\psi, \tag{25}$$

$$\tilde{r}^{\,t_b+T} = \mathrm{Norm}\left(g_{\bar{\psi}}(S_{t_b})\right). \tag{26}$$

The target is fixed for

$$t \in \{t_b, t_b+1, \ldots, t_b+T-1\}, \qquad \tilde{r}^t_{\mathrm{target}} = \tilde{r}^{\,t_b+T}. \tag{27}$$

A new target is generated only at the beginning of the next block.

The online predictor g_ψ continues to receive delayed supervision during the block; only the target used to optimize q_ϕ is held fixed. This separates the predictor's update timescale from the importance function's supervision timescale.

During the first T epochs, before any delayed future-loss targets become available, we warm up the learner and collect the initial state-loss pairs. Predictive reweighting is activated once the first delayed targets are observed.

D Implementation Details

We implement our framework using PyTorch and evaluate it under a standard multi-round active learning protocol. For all experiments, we initialize the active learning process with a randomly

selected labeled subset and iteratively acquire additional samples according to the proposed acquisition strategy. Unless otherwise specified, all experiments are repeated with over 5 random seeds, and the averaged performance is reported.

For image classification experiments, we use ResNet-18 (He et al., 2016) as the backbone without pretraining. For text classification, we adopt RoBERTa-base (Liu et al., 2019) with pretrained weights. For audio deepfake detection, we employ AASIST (Jung et al., 2022) without pretraining. The proposed framework is inserted before the final acquisition stage and can be applied to different classification architectures without modifying their backbone structures.

The look-ahead loss predictor is trained online during the active learning process using the observed training dynamics. Specifically, we maintain historical loss information and use it to predict future sample-wise losses after a look-ahead interval of T training steps. The predictor is optimized by minimizing the ranking objective described in Section 3, where informative sample pairs are selected according to their loss differences. The warm-up period, look-ahead horizon, predictor architecture, and pair sampling strategy are configured consistently across experiments unless otherwise stated.

For each acquisition round, the model is trained on the current labeled set for a fixed number of epochs before selecting new samples from the unlabeled pool. The selected samples are added to the labeled set, and the process continues until reaching the predefined labeling budget. Detailed training configurations, including optimizer settings, learning rates, batch sizes, and all hyperparameters, are provided in the supplementary materials together with the complete implementation. Core-Set, BADGE, GLISTER, and random sampling[2], and AutoAL[3] were derived from open-sourced versions on GitHub. All training and backbone settings are the same across all algorithms mentioned in this paper. Due to the substantial memory overhead of differentiable acquisition strategy search, we were unable to run AutoAL on AG News and ASVspoof2019 under our available computational resources, which run out of memory on 8x24G NVIDIA RTX 4090 by FSDP. Therefore, AutoAL results are reported only on the datasets where the method could be successfully executed.

[2]https://github.com/decile-team/distil

[3]https://github.com/haizailache999/AutoAL