
AutoAL: Automated Active Learning with Differentiable Query Strategy Search

Yifeng Wang¹ Xueying Zhan² Siyu Huang³

Abstract

As deep learning continues to evolve, the need for data efficiency becomes increasingly important. Considering labeling large datasets is both time-consuming and expensive, active learning (AL) provides a promising solution to this challenge by iteratively selecting the most informative subsets of examples to train deep neural networks, thereby reducing the labeling cost. However, the effectiveness of different AL algorithms can vary significantly across data scenarios, and determining which AL algorithm best fits a given task remains a challenging problem. This work presents the first differentiable AL strategy search method, named **AutoAL**, which is designed on top of existing AL sampling strategies. AutoAL consists of two neural nets, named SearchNet and FitNet, which are optimized concurrently under a differentiable bi-level optimization framework. For any given task, SearchNet and FitNet are iteratively co-optimized using the labeled data, learning how well a set of candidate AL algorithms perform on that task. With the optimal AL strategies identified, SearchNet selects a small subset from the unlabeled pool for querying their annotations, enabling efficient training of the task model. Experimental results demonstrate that AutoAL consistently achieves superior accuracy compared to all candidate AL algorithms and other selective AL approaches, showcasing its potential for adapting and integrating multiple existing AL methods across diverse tasks and domains. Code is available at: <https://github.com/haizailache999/AutoAL>.

¹Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, USA ²Ray and Stephanie Lane Computational Biology Department, Carnegie Mellon University, Pittsburgh, USA ³Visual Computing Division, School of Computing, Clemson University, Clemson, USA. Correspondence to: Xueying Zhan <xueyingz@andrew.cmu.edu>.

1. Introduction

With the development of deep learning (DL) techniques, large volumes of high-quality labeled data have become increasingly crucial as the model training processes are usually data hungry (Ren et al., 2021; Zhan et al., 2022). Active learning (AL) addresses this challenge by iteratively selecting the most informative unlabeled samples for annotation, thereby boosting model efficiency (Xie et al., 2021; Wang et al., 2023). Deep active learning strategies are typically divided into two categories: uncertainty-based and representativeness/diversity-based approaches (Sener & Savarese, 2017; Zhu & Bento, 2017; Ash et al., 2019). Uncertainty-based methods focus on querying samples the model is most uncertain about, while diversity-based strategies aim to select a diverse subset of samples that effectively represent the entire dataset (Citovsky et al., 2021). Both strategies have limitations, especially in batch selection and varying dataset complexities, leading to the inconsistency of the strategy efficiency (Ren et al., 2021). Hybrid strategies, which balance uncertainty and diversity (Zhan et al., 2021a;b), offer partial solutions. However, these strategies rely on strategy choices and subjective experiment settings, which can be ineffective in real-life applications.

Previous works (Baram et al., 2004; Hsu & Lin, 2015; Pang et al., 2018) have explored selecting among different AL strategies to achieve optimal selections without human effort. SelectAL (Hacohen & Weinshall, 2024) introduces a method for dynamically choosing AL strategies for different deep learning tasks and budgets by estimating the relative budget size of the problem. It involves evaluating the impact of removing small subsets of data points from the unlabeled pool to predict how different AL strategies would perform. However, this strategy depends on approximating the reduction in generalization error on small subsets, which would not fully capture the complexities of real-world tasks. Zhang et al. (2024) proposes to select the optimal batch of AL strategy from hundreds of candidates, aiming to maximize future cumulative rewards based on noisy past reward observations of each candidate algorithm. However, both the computational cost and the lack of differentiability make the optimization of these works inefficient.

To address these challenges, we propose **AutoAL**, an automated active learning framework that leverages differen-

tiable query strategy search. AutoAL offers two key advantages: 1) it enables automatic and efficient updates during optimization, allowing the model to adapt seamlessly; 2) it is inherently data-driven, leveraging the underlying data distribution to guide the AL strategy selection. However, integrating multiple AL strategies into the candidate pool makes achieving differentiability difficult. To overcome this issue, AutoAL leverages a bi-level optimization framework that smoothly integrates and optimizes different AL strategies. We design two neural networks: SearchNet and FitNet. FitNet is trained on the labeled dataset to adaptively yield the informativeness of each unlabeled sample. Guided by FitNet’s task loss, SearchNet is efficiently optimized to accurately select the best AL strategy. Instead of searching over a discrete set of candidate AL strategies, we relax the search space to a continuous domain, allowing SearchNet to be optimized via gradient descent based on FitNet’s output. The gradient-based optimization, which is more data-efficient than traditional black-box searches, enables AutoAL to achieve state-of-the-art performance in strategy selection with significantly reduced computational costs. Our key contributions are as follows:

- We introduce a novel AutoAL framework for differentiable AL query strategy search. To the best of our knowledge, AutoAL is the first automatic query strategy search algorithm that can be trained in a differentiable way, achieving highly efficient updates of query strategy in a data-driven manner.
- AutoAL is built upon existing AL strategies (i.e., uncertainty-based and diversity-based). By incorporating most AL strategies into its search space, AutoAL leverages the advantages of these strategies by finding the most effective ones for specific AL settings, objectives, or data distributions.
- We evaluate AutoAL across various AL tasks. Given the importance of AL in domain-specific tasks, particularly where data quality and imbalance are concerns, we conduct extensive experiments on both nature image datasets and medical datasets. The results demonstrate that AutoAL consistently outperforms all candidate strategies and other state-of-the-art AL methods across these datasets.

2. Related Work

Active Learning. Active Learning has been studied for decades to improve the model performance using small sets of labeled data, significantly reducing the annotation cost (Cohn et al., 1996; Settles, 2009; Zhan et al., 2021b). Most AL research focuses on pool-based AL, which identifies and selects the most informative samples from a large unlabeled pool iteratively. Pool-based AL sampling strategies are generally categorized into two primary branches:

diversity-based and uncertainty-based. The diversity-based methods select the batch of samples that can best represent the entire dataset distribution. CoreSet (Sener & Savarese, 2017) selects a batch of representative points from a core set, which is a subsample of the dataset that effectively serves as a proxy for the entire set. In contrast, uncertainty-based methods focus on selecting the samples with high uncertainty, which is due to the data generation or the modeling/learning process. Bayesian Active Learning by Disagreements (BALD) (Gal et al., 2017) chooses data points that are expected to maximize the mutual information between predictions and model posterior. Meta-Query Net (Park et al., 2022) trains an MLP that receives one open-set score and one AL score as input and outputs a balanced meta-score for sample selection in the open-set dilemma. Yoo & Kweon (2019) incorporates a loss prediction strategy by appending a compact parametric module designed to forecast the loss of unlabeled inputs relative to the target model. This is achieved by minimizing the discrepancy between the predicted loss and the actual target loss. It requires subjective judgment and strategic selection of active learning methods to be effective on specific datasets in real-world settings.

Adaptive Sample Selection in AL. Neither diversity-based nor uncertainty-based methods are perfect in all data scenarios. Diversity-based methods tend to perform better when the dataset contains rich category content and large batch size, while uncertainty-aware methods typically perform better in opposite settings (Zhan et al., 2021b; Citovsky et al., 2021). Several past works have explored adaptive selection in active learning to solve this problem. A common approach is to choose the best AL strategy from a set of candidate methods and use it to query unlabeled data, such as (Hacohen & Weinshall, 2024) and (Zhang et al., 2024) introduced in Section 1. Active Learning By Learning (ALBL) (Hsu & Lin, 2015) adaptively selects from a set of existing algorithms based on their estimated contribution to learning performance. It frames each candidate as a bandit machine, converting the problem into a multi-armed bandit scenario. However, none of these methods leverage the differentiable framework for automatic AL selection. Compared to the methods mentioned above, AutoAL first relaxes the search space to be continuous, thus can optimize automatically via gradient descent. This framework enables AutoAL to seamlessly integrate multiple existing AL strategies and rapidly adapt to select the optimal strategy based on the labeled pool. The simple gradient-based optimization is much more data-efficient than traditional black-box searches. Therefore, AutoAL can significantly reduce computational costs compared to these methods.

Bilevel Optimization. Bilevel optimization is a hierarchical mathematical framework where the feasible region of one optimization task is constrained by the solution set of

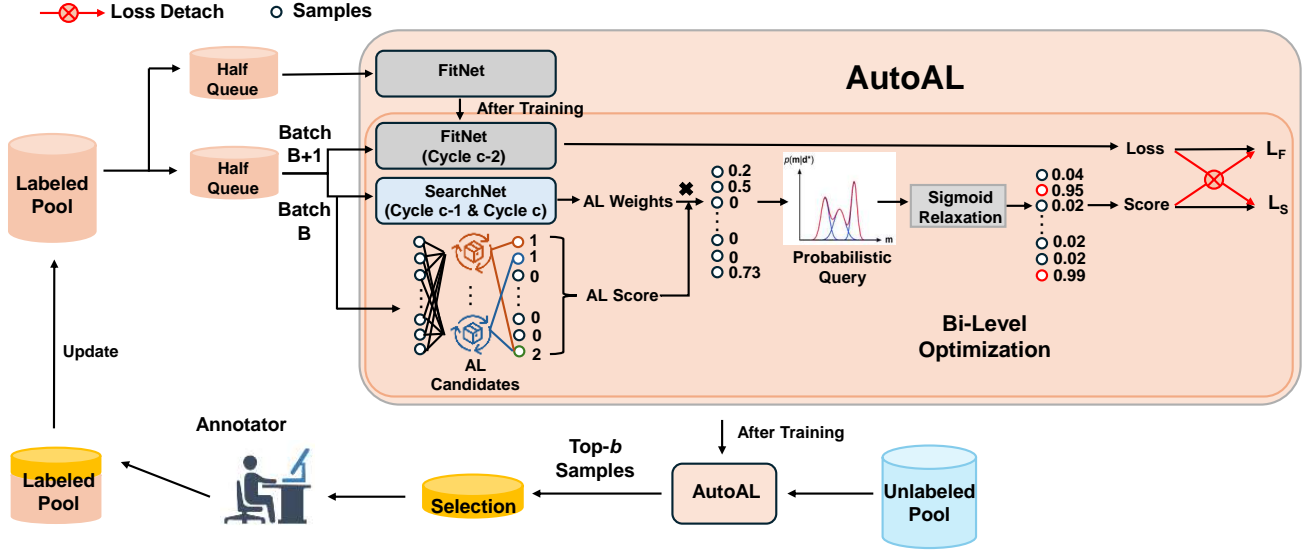


Figure 1. Overall framework of differentiable query strategy search for automated active learning (AutoAL). AutoAL leverages labeled pool to train the FitNet and SearchNet in a bi-level optimization mode. Data samples with the largest search score are then selected from the unlabeled pool.

another optimization task (Liu et al., 2021). It has gained significant attention due to its nested problem structure, which allows the two tasks to optimize jointly. Bilevel optimization adapts to many DL applications (Chen et al., 2022), such as hyperparameter optimization (Chen et al., 2019; MacKay et al., 2019), meta-knowledge extraction (Finn et al., 2017), neural architecture search (Liu et al., 2018; Hu et al., 2020), and active learning (Sener & Savarese, 2017). For instance, Sener & Savarese (2017) formulates AL as a bi-level coresampling problem and designs a 2-approximation method to select a subset of data that represents the entire dataset. In this work, we novelly formulate the AL query strategy search as a bi-level optimization problem, i.e., iteratively updating FitNet and using its task loss to guide the optimization of SearchNet. We further reformulate it as a differentiable learning paradigm to enable a more efficient selection of examples across various objectives and data distributions.

3. Methodology

3.1. Problem Setting

Pool-based AL focuses on selecting the most informative data iteratively from a large pool of unlabeled independent and identically distributed (*i.i.d.*) data samples until a fixed budget is exhausted or the expected model performance has been reached. Specifically, in AL processes, assuming that we have an initial seed labeled set $L = \{(x_j, y_j)\}_{j=1}^M$ and a large unlabeled data pool $U = \{x_i\}_{i=1}^N$, where $M \ll N$, $y_i \in \{0, 1\}$ is the class label of x_i for binary classification and $y_i \in \{1, \dots, p\}$ for multi-class classification. In each

iteration, we select a batch of the most informative data $Q^* = \arg \max_{x \in U} \alpha(x, \Omega)$ with batch size b from U based on the basic learned model Ω and the acquisition function $\alpha(x, \Omega)$, and query their labels from oracle/annotator. With these samples, the expected loss function f of the basic learned model can be minimized. L and U are then updated.

The performance of existing AL methods relies on the choice of query strategies, which should be carefully adapted to different tasks or applications. To achieve a robust and adaptive AL performance, this work novelly creates an automated AL strategy selection framework. AutoAL integrates two neural nets, FitNet Ω_F and SearchNet Ω_S upon each basic learned model Ω to select the best AL strategy before selecting Q . To facilitate automatic selection, AutoAL incorporates them into a bi-level optimization framework and relaxes the search space to enable differentiable updates. We discuss more details of AutoAL framework in the following sections.

3.2. Automated Active Learning

We propose Automated Active Learning (AutoAL), which aims to adaptively deliver the optimal query strategy for each sample in a given unlabeled data pool. Specifically, AutoAL consists of two neural networks: FitNet Ω_F and SearchNet Ω_S . Ω_S selects the optimal AL strategy from a set \mathcal{A} , which contains K candidate AL sampling strategies $\{\mathcal{A}_\kappa\}_{\kappa \in [K]}$ (e.g., BALD, Maximum Entropy, etc). Ω_F models the data distribution within the unlabeled dataset and guides the training of SearchNet Ω_S . Since annotations

for the unlabeled data cannot be accessed directly by the model Ω , AutoAL requires *no* data from the unlabeled pool U , using only the labeled dataset L for training.

In AutoAL, each AL iteration consists of \mathcal{C} cycles. In each cycle c , the labeled data L is randomly split into two equal-sized subsets: training and validation sets. FitNet Ω_F computes the cross-entropy loss \mathcal{L} (Zhang & Sabuncu, 2018) for data. Meanwhile, SearchNet Ω_S generates scores for the samples, which are used to rank them. This raises a key question: If only the training set is available, without access to the unlabeled pool, can Ω_S still effectively select the optimal samples?

Our key design is intuitive: SearchNet Ω_S treats the training batch as the labeled pool and the validation batch as the unlabeled pool. This allows Ω_S to simulate the process of AL selection, without direct access to the actual unlabeled data pool. The output of Ω_S is shown in Figure 1, a sample-wise score that aggregates the candidate AL strategy selection. The output of Ω_S is the informative loss of each sample. The optimization goal of Ω_S is to select data with higher losses, as they are expected to provide more informative updates to the neural network. For Ω_F , the training objective is to minimize the loss of the selected samples, particularly those *prioritized* by Ω_S . AutoAL is formulated as a bi-level optimization problem, as the optimization of Ω_F is dependent only on itself, but the optimization of Ω_S depends on the optimal Ω_F .

$$\Omega_S^* = \arg \max_{\Omega_S} \sum_{j=1}^{M/2} \mathcal{L}_S((x_j, y_j), \Omega_S, \Omega_F^*), \quad (1)$$

$$\text{s.t. } \Omega_F^* = \arg \min_{\Omega_F} \sum_{j=1}^{M/2} \mathcal{L}_F((x_j, y_j), \Omega_F).$$

\mathcal{L}_F and \mathcal{L}_S represent the losses of Ω_F and Ω_S , respectively. The neural nets Ω_F and Ω_S are optimized jointly as outlined in Eq. 1. At the lower level of the nested optimization framework, Ω_F is optimized to approach the distribution of the dataset. At the upper level, Ω_S is optimized to output the informativeness of each data, based on the optimal distribution returned from Ω_F .

3.3. Differentiable Query Strategy Optimization

Although Eq. 1 shows an effective framework for automatically deriving the optimal query strategy, solving the bi-level optimization problem is inefficient in practice. To address this, we derive a probabilistic query strategy coupled with a differentiable optimization framework, improving the efficiency of the optimization process.

Probabilistic query strategy. During training, for each labeled data x_j , we query a score $\mathcal{S}_\kappa(x_j)$ from each candidate AL sampling strategy \mathcal{A}_i , where $\mathcal{S}_\kappa(x_j) \in [0, 1]$ indicates whether the sample is selected. To model the overall score

$\mathcal{S}(x_j)$, which is a combination of all the AL search scores $\mathcal{S}_\kappa(x_j)$, we adopt a *Gaussian Mixture* approach (Reynolds et al., 2009), which has been proved efficient in AL (Zhao et al., 2020):

$$p(\mathcal{S}) = \sum_{k=1}^K \pi_k \mathcal{N}(\mathcal{S} | \mu_k, \Sigma_k), \quad (2)$$

$$\{\hat{\pi}_k, \hat{\mu}_k, \hat{\Sigma}_k\} = \arg \max_{\pi_k, \mu_k, \Sigma_k} \prod_{j=1}^M p(\mathcal{S}(x_j)), \quad (3)$$

where π_k is the weight of the k -th gaussian component and \mathcal{N} is the probability density function. Then we generate samples according to the Gaussian Mixture model and get the t -th maximum value, where t is the ratio of batch size b to the total pool size $M + N$. $W_{\kappa,j}$ is the probability of each query strategy for each sample predicted from the neural network Ω_S .

$$S_{\text{sample}} \sim \sum_{k=1}^K \hat{\pi}_k \mathcal{N}(\mathcal{S} | \hat{\mu}_k, \hat{\Sigma}_k), \quad (4)$$

$$\hat{\mathcal{S}}_\kappa(x_j, \Omega_S) = (\mathcal{S}_\kappa(x_j) - \vartheta_t(S_{\text{sample}})) * W_{\kappa,j}. \quad (5)$$

A sample with a higher score $\hat{\mathcal{S}}$ indicates that it has been selected by more candidate AL strategies, and therefore should be given higher priority for labeling.

Differentiable acquisition function optimization. If we limit the individual query score $\mathcal{S}_\kappa(x_j)$ to discrete values $\{0, 1\}$, the bi-level optimization objective of Eq. 1 becomes non-differentiable and still difficult to optimize. To enable differentiable optimization, we relax the categorical selection of strategies into a continuous space, we apply a *Sigmoid* function over all possible strategies.:

$$\bar{\mathcal{S}}(x_j) = \sum_{\kappa \in K} \frac{\lambda}{1 + \exp(-\Theta_{\mathcal{S}'_\kappa}^{(j)})} \hat{\mathcal{S}}_\kappa(x_j, \Omega_S), \quad (6)$$

where the strategy mixing weights for each sample j is parameterized by a vector $\Theta^{(j)}$ of dimension $|\mathcal{A}|$. λ is a scaling vector. $\bar{\mathcal{S}}(x_j)$ is the final differentiable learning objective function, which can be optimized with back-propagation. Therefore, Ω_F and Ω_S can be efficiently optimized under Eq. 1.

3.4. Learning and Algorithm of AutoAL

With differentiable query strategy optimization in Sec. 3.3, we reformulate the bi-level optimization problem in Sec. 3.2 as efficient optimization objectives as follows:

$$\mathcal{L}_F = \frac{1}{B} \sum_{j'=1+(n+1)B}^{(n+2)B} \bar{\mathcal{S}}_{\text{detach}}(x'_j) \cdot \mathcal{L}(x'_j, y'_j) + \bar{\lambda} \mathcal{L}_{re}(t, B), \quad (7)$$

$$\mathcal{L}_S = -\frac{1}{B} \sum_{j=1+nB}^{(n+1)B} \bar{S}(x_j) \cdot \mathcal{L}_{detach}(x_j, y_j) - \bar{\lambda} \mathcal{L}_{re}(t, B), \quad (8)$$

As described in section 3.2, in cycle $c-2$, Ω_F is optimized to get the data distribution. The final loss is calculated according to Eq. 7. In cycle $c-1$, AutoAL optimizes Ω_S to get the data sample with highest loss. The loss function from Ω_F is detached to avoid updating itself, and the final loss is calculated according to Eq. 8. Here we use negative loss to achieve gradient ascent. For $n \in \{0, 1, \dots, \lfloor \frac{M}{2B} \rfloor\}$, AutoAL additionally integrates a loss prediction module according to (Yoo & Kweon, 2019) to help update Ω_S , as

$$\mathcal{L}_{re}(t, B) = \left(\frac{1}{1 + \exp(0.5 \cdot |\alpha - t \cdot B|)} - 0.5 \right). \quad (9)$$

\mathcal{L}_{re} represents a regularization loss that limits the number of selected samples, as shown in Eq. 9. Here, α is the number of selected samples, t is the ratio of the query batch size b in each AL iteration to the total pool size $M + N$, and $\bar{\lambda}$ is a scaling vector. The final learning objective is to select the most informative samples while effectively leveraging the underlying data distribution.

The AutoAL algorithm is illustrated in Algorithm 1. In each AL iteration, the labeled set is first used to train the FitNet and SearchNet as described in section 3.2. Then the optimal SearchNet Ω_S^* is applied to the unlabeled pool to select a new batch of data, where Ω_S^* is used to generate score \bar{S} for each sample x_i , and $Q^* = \arg \max_{x \in U} \bar{S}(x_i)$ with top- b highest scoring samples are selected and labeled. The labeled set L and the unlabeled set U are updated as $L = L + Q^*$ and $U = U - Q^*$, respectively. The updated labeled set is then used for task model training.

Algorithm 1 AutoAL: Automated Active Learning with Differentiable Query Strategy Search

Input: K candidate algorithms $\mathcal{A} = \{\mathcal{A}_\kappa\}_{\kappa \in [K]}$, labeled pool $L = \{(x_j, y_j)\}_{j=1}^M$ and unlabeled pool $U = \{x_i\}_{i=1}^N$, total number of AL rounds \mathcal{R} , batch size b , task model \mathcal{T} .

Output: task model \mathcal{T}

- 1: Initialize \mathcal{T}
 - 2: **for** $c=1, \dots, \mathcal{R}$ **do**
 - 3: Optimize Ω_F^* , Ω_S^* according to Eq. 1;
 - 4: Calculate $\bar{S}(x_i)$ by Eq. 6 for all $i \in N$;
 - 5: Select the top b samples with the highest score $\bar{S}(x_j)$;
 - 6: Query label y_i for all $i \in b$; Update L and U ;
 - 7: Train task model \mathcal{T} by using L ;
 - 8: **end for**
 - 9: return \mathcal{T}
-

4. Experiments

We empirically evaluate the performance of AutoAL on a wide range of datasets with both nature and medical images. These experiments demonstrate that AutoAL outperforms

the baselines by fully automating the training within the differentiable bi-level optimization framework. Additionally, we conduct ablation studies to analyze the contribution of each module and examine the effects of the candidate AL strategies, including different numbers of candidates.

4.1. Experiment Settings

Datasets and Baselines. We conduct AL experiments on seven datasets: Cifar-10 and Cifar-100 (Krizhevsky et al., 2009), SVHN (Netzer et al., 2011), TinyImageNet (Le & Yang, 2015) in the nature image domain, and OrganCM-NIST, PathMNIST, and TissueMNIST from MedMNIST database (Yang et al., 2023) in the medical image domain. The details of datasets are presented in Appendix Table 1. We compare our method against multiple AL sampling strategies, including Maximum Entropy Sampling (Shannon, 1948), Margin Sampling (Netzer et al., 2011), Least Confidence (Wang & Shang, 2014), KMeans Sampling (Ahmed et al., 2020), Bayesian Active Learning by Disagreements (BALD) (Gal et al., 2017), Variation Ratios (VarRatio) (Freeman, 1965) and Mean Standard Deviation (MeanSTD) (Kampffmeyer et al., 2016). We also consider the state-of-the-art deep AL methods as baselines, including Batch Active learning by Diverse Gradient Embeddings (BADGE) (Ash et al., 2019), Loss Prediction Active Learning (LPL) (Yoo & Kweon, 2019), Variational Adversarial Active Learning (VAAL) (Sinha et al., 2019), Core-set Selection (Coreset) (Sener & Savarese, 2017), Ensemble Variance Ratio Learning (ENSvarR) (Beluch et al., 2018), Deep Deterministic Uncertainty (DDU) (Mukhoti et al., 2023). We also include ALBL (Hsu & Lin, 2015) to compare AutoAL with existing selective AL strategies.

Implementation Details. For Ω_F and Ω_S , we build the backbone using ResNet-18 (He et al., 2016). We also employ ResNet-18 as the classification model on all baselines and AutoAL for fair comparison. Since AutoAL is built upon existing AL strategies and focuses on selecting the optimal strategy, we integrate seven AL methods into AutoAL: Maximum Entropy, Margin Sampling, Least Confidence, KMeans, BALD, VarRatio, and MeanSTD. For the optimization of Ω_S and loss prediction module, we use SGD optimizer (Ruder, 2016). For the optimization of Ω_F , we use Adam optimizer (Kingma, 2014). Both use 0.005 as the learning rate. While training, FitNet will first update for 200 epochs using the validation queue, then Ω_F , Ω_S and the loss prediction module will update iteratively with a total of 400 epochs. All experiments are repeated *three times* with different randomly selected initial labeled pools, reporting mean and standard deviation.

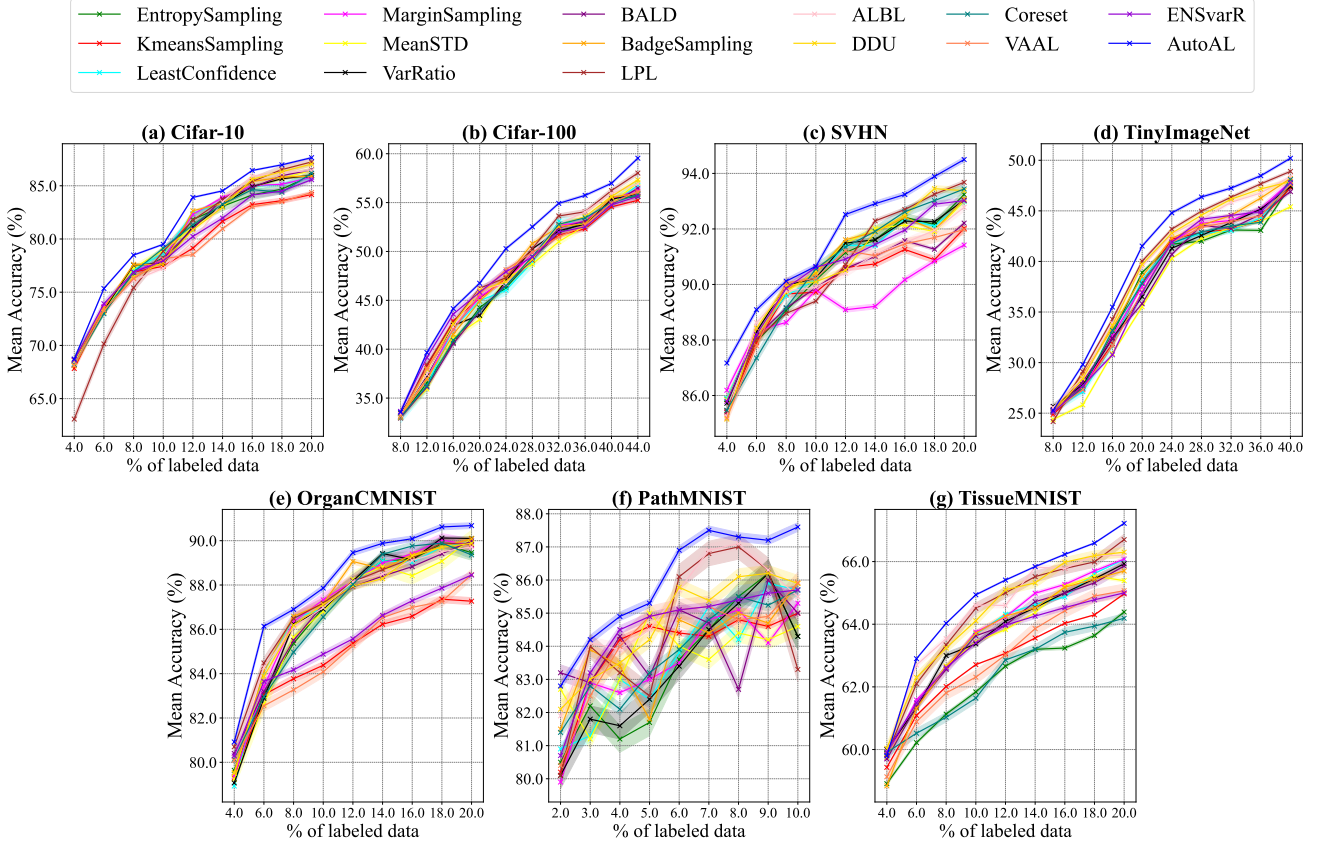


Figure 2. Overall performance on seven benchmark datasets: natural image datasets (top) and medical image datasets (bottom).

4.2. Evaluation of Active Learning Performance

Fig. 2 shows the overall performance comparison of different AL methods, where AutoAL consistently outperforms the baselines across all datasets. Additionally, we have the following observations:

1. Our method consistently outperforms the other approaches in terms of rounds, a crucial attribute for a successful active learning method. This adaptability is particularly valuable in real-world applications, where the labeling budget may vary significantly across different tasks. For instance, one might only have the resources to annotate 10% of the data, rather than 20% or 40%.
2. Our method demonstrates robust performance not only on easier datasets but also on more challenging ones, such as Cifar-100, OrganCMNIST and TinyImageNet. Cifar-100 and TinyImageNet has significantly more classes than other datasets, while OrganCMNIST has smaller data pool, with only 12,975 images. These challenges make active learning more difficult, yet our method’s superior performance across these datasets highlights its robustness and generalizability.
3. The labeled data vs. accuracy curves of our method are relatively smooth compared to those of other methods. The baseline methods sometimes show significant accuracy drops during certain active learning rounds. This mainly dues to harmful data selection (Koh & Liang, 2017) or overfitting problem. However, by selecting the optimal strategy thus selecting the informative data, AutoAL can alleviate the problem and make the curve much more smooth than baseline strategies.
4. Our method shows smaller standard deviations, indicating that it is robust across repeated experiments. This robustness is crucial for AL applications, especially in a real-world setting, where the AL process is typically conducted only once.
5. Different AL strategies produce varying results across datasets. For instance, Margin Sampling underperforms on the SVHN dataset, while KMeans Sampling and VAAL yield the worst performance on the OrganCMNIST and Cifar-10 datasets. This variability highlights the significance of our approach, which aims to identify and select the optimal strategy for each dataset.




Figure 3. Ablation study on three components of AutoAL. ‘ResNet Backbone’: without the loss prediction module, updating only the ResNet. ‘Loss Prediction Module’: without updating the ResNet backbone, optimizing solely with the loss prediction module. ‘ResNet+Loss Prediction’: the full AutoAL pipeline.




Figure 4. Ablation study on the size of the candidate pool in AutoAL strategy selection.

4.3. Ablation Studies

We perform the ablation study on two key components of AutoAL: the SearchNet architecture and the AL candidate strategies. Experiments are conducted on three datasets: Cifar-100, SVHN, and OrganCMNIST, which are chosen for their mix of natural and medical images, as well as the increased difficulty they present for active learning applications.

SearchNet Architecture. Our SearchNet consists of two main components: a ResNet-18 backbone and a loss prediction module for sample loss prediction (Yoo & Kweon, 2019). In this ablation study, we disable one component at a time and record the results. Fig. 3 illustrates the performance, revealing the effectiveness of both ResNet-18 backbone and the loss prediction module. However, when only the loss prediction module is used to optimize the network, performance drops significantly. This aligns with our intuition, as although the loss prediction module can assist

with network optimization, the main focus shifts from selecting the best AL strategy to merely minimizing the sample loss. This misalignment will result in incorrect AL strategy weight estimation (See Fig. 1).

Size of Active Learning Candidate Pool. Since AutoAL is built on multiple AL strategies, we study how the size of the AL candidate pool impacts performance. We start from the Maximum Entropy as the sole candidate for the initial baseline, then expand the candidate pool to three strategies by adding Margin Sampling and Least Confidence. For the pool of five candidates, we further include KMeans and BALD. As shown in Fig. 4, for all three datasets, the single candidate variation performs the worst but still outperforms the Entropy Sampling baseline in Fig. 2. This is because AutoAL does not directly query data samples with maximum entropy but queries the data with the highest score from SearchNet, which has been optimized with the help of the loss prediction module and FitNet. For Cifar-100 and

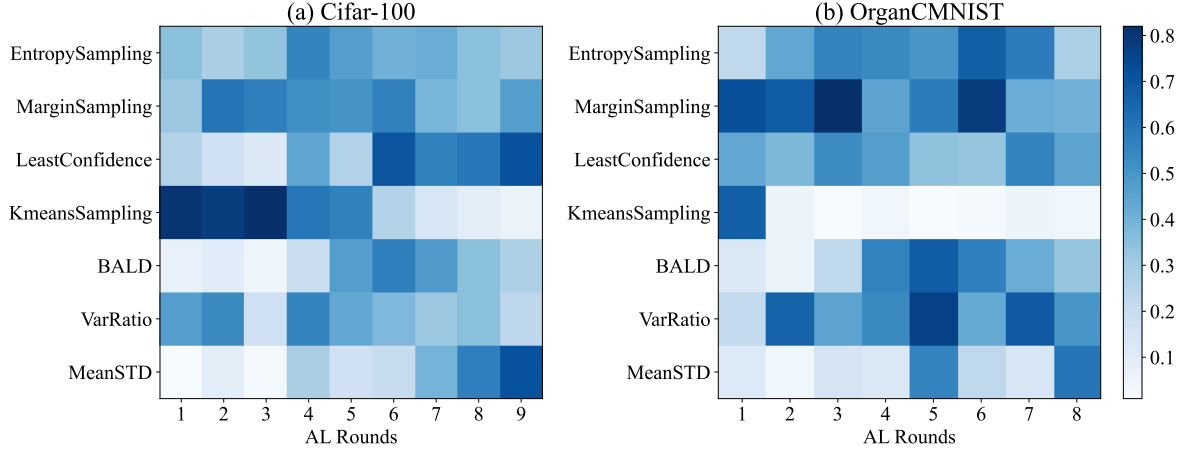


Figure 5. AL strategy scores across different AL rounds on CIFAR-100 and OrganCMNIST datasets.

OrganCMNIST, AutoAL can perform well when there are at least three candidates in the selection pool. However, for SVHN, more AL candidates result in better performance. This indicates that the upper bound of candidate pool size varies across different datasets. In certain datasets, incorporating more active learning strategies yields better results, but this is not the case for other datasets. Additionally, pools with more candidates exhibit lower standard deviations. This demonstrates AutoAL’s strong capability to select the best AL strategy. It can consistently choose the optimal one and exclude the less effective options.

4.4. AL Selection Strategy

Here, we examine which AL strategy is prioritized by AutoAL in different active learning rounds. For each AL strategy candidate, we calculate the AL score of each image according to Eq. 6. Notice that we don’t sum all the scores according to the dimension of each active learning, but compute the average value of all images over a round, and normalize the results for visualization. As illustrated in Fig. 5, in both experiments, KMeans dominate the sample selection during the initial rounds, such as the first round, but be de-prioritized in subsequent rounds. This indicates that diversity-based measures, which select samples that represent diverse regions of the data distribution, are more critical in the early stages of active learning. In the early rounds, when only a small percentage of the data has been labeled, the model’s understanding of the entire data distribution is limited. As a result, selecting samples that cover a broad spectrum of the dataset becomes crucial to provide the model with a more comprehensive view of the data.

On Cifar-100 dataset, both Least Confidence and MeanSTD, two uncertainty-based measures, demonstrate improved performance in the final rounds. AutoAL assigns them higher priority to enhance overall performance. As the AL pro-

cess progresses and more diverse data become available, the model attempts to develop a better understanding of the underlying structure of the target data. Uncertainty-based measures are more effective now because the major challenge switches to refining the decision boundaries. The transition from diversity to uncertainty-based strategies is consistent with the model’s evolving needs. In early rounds, the model requires diversity to build a broad understanding of the data, but once that foundation is established, its focus shifts toward uncertainty, which helps fine-tune model decision-making ability. This adaptive approach ensures the model targets the most diverse subsets at each stage of learning process, from broad exploration in the early rounds to fine-grained exploitation in the later rounds. By automatically adjusting strategy priorities, AutoAL improves the overall accuracy and shows stronger robustness.

5. Conclusion

In this paper, we have introduced AutoAL, the first automated search framework for deep active learning. AutoAL employs two neural networks, SearchNet and FitNet, integrated into a bilevel optimization framework to enable joint optimizations. To efficiently solve the bi-level optimization problem, we have proposed a probabilistic query strategy that relaxes the search space from discrete to continuous, enabling a differentiable and data-driven learning paradigm for AutoAL. This framework not only accelerates the search process but also ensures that AutoAL can generalize across a wide range of tasks and datasets. Furthermore, AutoAL is flexible, allowing for easy integration of most of the existing active learning strategies into the candidate pool, making it adaptable to evolving active learning techniques. Our extensive empirical evaluation, conducted across both natural and medical image datasets, has demonstrated the robustness and generalizability of AutoAL. The results show that

AutoAL consistently outperforms baselines, highlighting its capability to optimize sample selection and enhance model performance. In future, we plan to apply AutoAL to boarder machine learning tasks such as structured prediction and consider more sophisticated AL method combinations.

Acknowledgements

The authors acknowledge research support from Clemson University with a generous allotment of computation time on the Palmetto cluster.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning by introducing a novel automated search framework for deep active learning. Our work has many potential societal consequences, including reducing the annotation cost of training deep learning models and improving the efficiency of data annotation, etc. While our proposed method has the potential to promote decision-making in various downstream tasks (e.g., natural and medical image data selection), we acknowledge that our AutoAL has been carefully designed to mitigate potential biases and ensure fair selection.

References

- Ahmed, M., Seraj, R., and Islam, S. M. S. The k-means algorithm: A comprehensive survey and performance evaluation. *Electronics*, 9(8):1295, 2020.
- Ash, J. T., Zhang, C., Krishnamurthy, A., Langford, J., and Agarwal, A. Deep batch active learning by diverse, uncertain gradient lower bounds. *arXiv preprint arXiv:1906.03671*, 2019.
- Baram, Y., Yaniv, R. E., and Luz, K. Online choice of active learning algorithms. *Journal of Machine Learning Research*, 5(Mar):255–291, 2004.
- Beluch, W. H., Genewein, T., Nürnberger, A., and Köhler, J. M. 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.
- Chen, C., Chen, X., Ma, C., Liu, Z., and Liu, X. Gradient-based bi-level optimization for deep learning: A survey. *arXiv preprint arXiv:2207.11719*, 2022.
- Chen, Y., Chen, B., He, X., Gao, C., Li, Y., Lou, J.-G., and Wang, Y. λ opt: Learn to regularize recommender models in finer levels. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 978–986, 2019.
- Citovsky, G., DeSalvo, G., Gentile, C., Karydas, L., Rajagopalan, A., Rostamizadeh, A., and Kumar, S. Batch active learning at scale. *Advances in Neural Information Processing Systems*, 34:11933–11944, 2021.
- Cohn, D. A., Ghahramani, Z., and Jordan, M. I. Active learning with statistical models. *Journal of artificial intelligence research*, 4:129–145, 1996.
- Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pp. 1126–1135. PMLR, 2017.
- Freeman, L. C. Elementary applied statistics: for students in behavioral science. (*No Title*), 1965.
- Gal, Y., Islam, R., and Ghahramani, Z. Deep bayesian active learning with image data. In *International conference on machine learning*, pp. 1183–1192. PMLR, 2017.
- Hacohen, G. and Weinshall, D. How to select which active learning strategy is best suited for your specific problem and budget. *Advances in Neural Information Processing Systems*, 36, 2024.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Hsu, W.-N. and Lin, H.-T. Active learning by learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29, 2015.
- Hu, Y., Wu, X., and He, R. Tf-nas: Rethinking three search freedoms of latency-constrained differentiable neural architecture search. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XV 16*, pp. 123–139. Springer, 2020.
- Kampffmeyer, M., Salberg, A.-B., and Jenssen, R. Semantic segmentation of small objects and modeling of uncertainty in urban remote sensing images using deep convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pp. 1–9, 2016.
- Kingma, D. P. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Koh, P. W. and Liang, P. Understanding black-box predictions via influence functions. In *International conference on machine learning*, pp. 1885–1894. PMLR, 2017.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. 2009.

- Le, Y. and Yang, X. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015.
- Liu, H., Simonyan, K., and Yang, Y. Darts: Differentiable architecture search. *arXiv preprint arXiv:1806.09055*, 2018.
- Liu, R., Gao, J., Zhang, J., Meng, D., and Lin, Z. Investigating bi-level optimization for learning and vision from a unified perspective: A survey and beyond. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(12):10045–10067, 2021.
- MacKay, M., Vicol, P., Lorraine, J., Duvenaud, D., and Grosse, R. Self-tuning networks: Bilevel optimization of hyperparameters using structured best-response functions. *arXiv preprint arXiv:1903.03088*, 2019.
- Mukhoti, J., Kirsch, A., van Amersfoort, J., Torr, P. H., and Gal, Y. Deep deterministic uncertainty: A new simple baseline. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 24384–24394, 2023.
- Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., Ng, A. Y., 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.
- Pang, K., Dong, M., Wu, Y., and Hospedales, T. M. Dynamic ensemble active learning: A non-stationary bandit with expert advice. In *2018 24th International Conference on Pattern Recognition (ICPR)*, pp. 2269–2276. IEEE, 2018.
- Park, D., Shin, Y., Bang, J., Lee, Y., Song, H., and Lee, J.-G. Meta-query-net: Resolving purity-informativeness dilemma in open-set active learning. *Advances in Neural Information Processing Systems*, 35:31416–31429, 2022.
- Ren, P., Xiao, Y., Chang, X., Huang, P.-Y., Li, Z., Gupta, B. B., Chen, X., and Wang, X. A survey of deep active learning. *ACM computing surveys (CSUR)*, 54(9):1–40, 2021.
- Reynolds, D. A. et al. Gaussian mixture models. *Encyclopedia of biometrics*, 741(659-663), 2009.
- Ruder, S. An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747*, 2016.
- Sener, O. and Savarese, S. Active learning for convolutional neural networks: A core-set approach. *arXiv preprint arXiv:1708.00489*, 2017.
- Settles, B. Active learning literature survey. 2009.
- Shannon, C. E. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948.
- Sinha, S., Ebrahimi, S., and Darrell, T. Variational adversarial active learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 5972–5981, 2019.
- Wang, D. and Shang, Y. A new active labeling method for deep learning. In *2014 International joint conference on neural networks (IJCNN)*, pp. 112–119. IEEE, 2014.
- Wang, Y., Tu, Z., Xiang, Y., Zhou, S., Chen, X., Li, B., and Zhang, T. Rapid image labeling via neuro-symbolic learning. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 2467–2477, 2023.
- Xie, Y., Tomizuka, M., and Zhan, W. Towards general and efficient active learning. *arXiv preprint arXiv:2112.07963*, 2021.
- Yang, J., Shi, R., Wei, D., Liu, Z., Zhao, L., Ke, B., Pfister, H., and Ni, B. Medmnist v2-a large-scale lightweight benchmark for 2d and 3d biomedical image classification. *Scientific Data*, 10(1):41, 2023.
- Yoo, D. and Kweon, I. S. Learning loss for active learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 93–102, 2019.
- Zhan, X., Li, Q., and Chan, A. B. Multiple-criteria based active learning with fixed-size determinantal point processes. *arXiv preprint arXiv:2107.01622*, 2021a.
- Zhan, X., Liu, H., Li, Q., and Chan, A. B. A comparative survey: Benchmarking for pool-based active learning. In *IJCAI*, pp. 4679–4686, 2021b.
- Zhan, X., Wang, Q., Huang, K.-h., Xiong, H., Dou, D., and Chan, A. B. A comparative survey of deep active learning. *arXiv preprint arXiv:2203.13450*, 2022.
- Zhang, J., Shao, S., Verma, S., and Nowak, R. Algorithm selection for deep active learning with imbalanced datasets. *Advances in Neural Information Processing Systems*, 36, 2024.
- Zhang, Z. and Sabuncu, M. Generalized cross entropy loss for training deep neural networks with noisy labels. *Advances in neural information processing systems*, 31, 2018.
- Zhao, J., Sun, S., Wang, H., and Cao, Z. Promoting active learning with mixtures of gaussian processes. *Knowledge-Based Systems*, 188:105044, 2020.
- Zhu, J.-J. and Bento, J. Generative adversarial active learning. *arXiv preprint arXiv:1702.07956*, 2017.

A. Appendix

Table 1. A summarization of datasets used in the experiments. The Imbalance Ratio (IR) is the ratio of the number of images in the majority class to the number of images in the minority class. $IR = 1$ represents balance dataset.

	DataSet	Training Size	Test Size	Class Numbers	Imbalance Ratio
Nature Images	Cifar-10	50,000	10,000	10	1.0
	Cifar-100	50,000	10,000	100	1.0
	SVHN	73,257	26,032	10	3.0
	TinyImageNet	100,000	10,000	200	1.0
Medical Images	OrganCMNIST	12,975	8,216	11	5.0
	PathMNIST	89,996	7,180	9	1.6
	TissueMNIST	165,466	47,280	8	9.1

A.1. Complexity Analysis

Table 2. The average runtime of the entire active learning (AL) process (including training and querying), where the runtime of Entropy-Sampling is the **unit** time. We separate total time cost of AutoAL into three main parts: candidate ALs querying AL score for images (AutoAL-ComputeALScore), SearchNet and FitNet updating (AutoAL-Search), and querying samples and training classification model after AutoAL search (AutoAL-TrainResNet).

	CIFAR-10	CIFAR-100	SVHN	TinyImageNet	OrganCMNIST	PathMNIST	TissueMNIST
EntropySampling	1.00	1.00	1.00	1.00	1.00	1.00	1.00
MarginSampling	0.99	1.02	5.97	0.99	2.01	1.01	1.84
LeastConfidence	0.85	1.01	1.03	0.96	1.85	0.97	1.81
KMeansSampling	1.20	1.24	1.59	0.62	1.33	0.93	2.37
MeanSTD	1.46	1.18	0.89	0.23	2.07	0.85	2.35
VarRatio	1.40	1.17	0.87	0.87	1.87	0.78	2.30
BALD	1.48	1.18	0.90	0.92	2.04	1.02	1.86
BadgeSampling	1.83	4.09	2.34	8.32	1.21	1.49	2.52
LPL	2.03	0.87	0.98	0.37	1.87	1.35	1.77
ALBL	1.08	0.57	1.01	0.82	2.10	0.88	1.42
DDU	1.48	1.87	2.13	1.62	1.32	2.32	1.68
Coreset	1.68	1.42	1.07	1.34	1.25	1.47	2.05
VAAL	1.24	1.27	7.02	2.05	5.28	4.36	7.68
ENSvarR	4.21	4.52	6.58	12.27	8.33	7.63	10.46
AutoAL-ComputeALScore	2.34	2.03	2.80	3.40	3.79	3.24	1.74
AutoAL-Search	0.12	0.11	0.15	0.18	0.20	0.17	0.09
AutoAL-TrainResNet	1.57	1.36	1.88	2.28	2.55	2.18	1.15
AutoAL-Total	4.03	3.50	4.83	5.86	6.54	5.59	2.95

In this subsection, we perform additional analysis experiments to showcase AutoAL’s efficiency and generalizability. We mainly use the average running time to verify the results, and all the experiments was done on one Nvidia A100 GPU. We use the running time of EntropySampling as a benchmark and calculate other ALs running costs relative to it. We find:

1. Our AutoAL demonstrates its ability to generalize to large datasets, and the time cost is compatible to some of the baselines such as ENSvarR.
2. We separate the total time cost of our method into three main parts, and the AL strategy sampling will cost the most time and the update of SearchNet and FitNet will nearly cost none. This means AutoAL will time cost rely heavily on the candidate pool. When the candidate ALs are fast enough, AutoAL will be in low cost and complexity.
3. We conducted the experiment on changing the candidate pool size, and we found that when there are 3 candidates, the time cost is just 1.3x compared to EntropySampling, and our method with three candidates can also gain a high

improvement on labeling accuracy. This will give real-world users flexibility in choosing the number of candidate ALs considering both accuracy and complexity.