

Active Learning to Classify Macromolecular Structures in situ with Fewer Supervision in Cryo-Electron Tomography

Xuefeng Du¹, Haohan Wang³, Zhenxi Zhu⁵, Xiangrui Zeng², Yi-Wei Chang⁶,
Jing Zhang⁷, Eric Xing⁴ and Min Xu^{2,*}

¹ Department of Computer Science,
University of Wisconsin-Madison, Madison, 53706, USA

² Computational Biology Department,
Carnegie Mellon University, Pittsburgh, 15213, USA

³ Language Technologies Institute,
Carnegie Mellon University, Pittsburgh, 15213, USA

⁴ Machine Learning Department,
Carnegie Mellon University, Pittsburgh, 15213, USA

⁵ Department of Computer Science,
Beijing University of Posts and Telecommunications, 100876, China

⁶ Department of Biochemistry and Biophysics,
University of Pennsylvania, Philadelphia, 19104, USA

⁷ Department of Computer Science,
University of California - Irvine, Irvine, 92697, USA

* Corresponding author emails: mxu1@cs.cmu.edu

Abstract

Motivation: Cryo-Electron Tomography (cryo-ET) is a 3D bioimaging tool that visualizes the structural and spatial organization of macromolecules at a near-native state in single cells, which has broad applications in life science. However, the systematic structural recognition and recovery of macromolecules captured by cryo-ET are difficult due to high structural complexity and imaging limits. Deep learning based subtomogram classification have played critical roles for such tasks. However, as supervised approaches, their performance relies on sufficient and laborious annotation on a large training dataset.

Results: To alleviate this major labeling burden, we proposed a Hybrid Active Learning (HAL) framework for querying subtomograms for labelling from a large unlabeled subtomogram pool. Firstly, HAL adopts uncertainty sampling to select the subtomograms that have the most uncertain predictions. This strategy enforces the model to be aware of the inductive bias during classification and subtomogram selection, which satisfies the *discriminateness* principle in AL literature. Moreover, to mitigate the sampling bias caused by such strategy, a discriminator is introduced to judge if a certain subtomogram is labeled

or unlabeled and subsequently the model queries the subtomogram that have higher probabilities to be unlabeled. Such query strategy encourages to match the data distribution between the labeled and unlabeled subtomogram samples, which essentially encodes the *representativeness* criterion into the subtomogram selection process. Additionally, HAL introduces a subset sampling strategy to improve the diversity of the query set, so that the information overlap is decreased between the queried batches and the algorithmic efficiency is improved. Our experiments on subtomogram classification tasks using both simulated and real data demonstrate that we can achieve comparable testing performance (on average only 3% accuracy drop) by using less than 30% of the labeled subtomograms, which shows a promising result for subtomogram classification task with limited labeling resources.

1 Introduction

Cellular processes are generally governed by macromolecules. To accurately understand these processes, Cryo-Electron Tomography (cryo-ET) has been developed recently to enable a systematic 3D visualization of subcellular structures in single cells at sub-molecular resolution and in native state. However, due to the structural content complexity of the captured tomograms and imaging limitations, it is difficult to classify macromolecules in subtomograms (A *subtomogram* is a subvolume of a tomogram that is likely to contain a single macromolecule) for structural recovery via manual inspections. Given that subtomogram classification is essentially a 3D image classification problem, supervised deep learning has recently become a major approach thanks to its ability to extract complex image composition rules from big image data. However, even though different approaches have been developed on either 2D or 3D cryo-ET data [7, 8, 59, 33], few of them emphasize the labeling burden, which is very time-consuming and requires structural biology expertise. This situation impedes the off-the-shelf deployment of these algorithms. For instance, even for just 1,000 real subtomograms were used in [34], it already introduced a time-consuming labeling work for domain experts.

Under such circumstances, we resort *active learning*, which selects a subset of subtomogram samples, if labeled and used for training, will best improve the model’s performance under the same labeling budget [43, 17] (Fig. 1). Two main principles for such unlabeled sample selection are proposed [9] and they both have limitations: *discriminativeness* and *representativeness*. The discriminativeness principle aims to find the most discriminative samples for the current classifier, which will shrink the space of candidate classifiers as rapidly as possible [56]. The popular proposed criteria are uncertainty rule [62, 54], expected error reduction [21] and query by committee [46, 16]. In this case, the samples are selected based on specific criterion instead of being i.i.d. sampled. Such sampling bias prevents active learning from finding a classifier with good generalization performance and query efficiency [56], which becomes even severe for high-dimensional and complex 3D medical images. The representativeness principle aims to address this problem by querying the samples which can represent the overall patterns or statistics of the unlabeled data, such as by clustering [39] and generative models [67, 52, 27, 48, 25]. Such methods perform better when fewer initial labeled data is provided. However, they will become inefficient with the increase of queried classes, as they solely rely on data distributions and do not fully use the label information [56].

Since using either type of principle alone is not enough to guarantee the optimal result, in this paper, we approach this task by integrating the discriminativeness and representativeness in one optimization formulation, namely the Hybrid Active Learning (HAL) framework. To satisfy the principle of data representativeness, we start with a small labeled set and a large unlabeled set and train a supervised Convolutional Neural Network (CNN) on the labeled set. We then extract the feature representations of both the labeled and unlabeled set. Inspired by the distribution alignment techniques [15], in each iteration, we train a discriminator on these representations and predict how likely each subtomogram sample is labeled or unlabeled. Then, we select and label those subtomogram samples in the unlabeled dataset which is predicted to have higher probabilities of coming from unlabeled dataset. This alternative optimization scheme effectively improves the representativeness of the labeled training set. Moreover, since the subtomograms captured by cryo-ET are highly heterogeneous, a large selected batch is likely to contain redundant subtomogram samples, which leads to a significant information overlap and thus an inefficient querying process. Therefore, we apply a sub-sampling strategy to enlarge the query batch without losing diversity. For the discriminativeness principle, We additionally introduce the label




Figure 1: Exemplary illustration of the active learning approach in cryo-ET classification. (a) shows the existing passive recognition pipeline where sufficient labeling is required where (b) demonstrates the recognition scenario guided by active learning.

information by using the entropy of predictions as selection criterion. Such heuristic is a strong active learning baseline, namely uncertainty sampling [62]. In each sampling iteration, we use both principles to score the current unlabeled subtomogram samples and then ensemble the two scores for final ranking and selection. We then add all queried subtomogram samples into the labeled dataset and repeat until the labeling budget is reached. The overall learning and querying steps are summarized in Figure 2. Note that the hybrid querying heuristics are also proposed in literature [64, 2] and we defer the discussion in the appendix. The contributions of this paper is summarized as follows: 1) We propose a 3D HAL framework to query unlabeled subtomogram samples and expand the training dataset, such that deep models can be trained with significantly lower labeling cost while incurring minimal prediction accuracy drop. We provide a theoretical analysis of the expected classification risk of our framework (Equation 2). 2) HAL is the first active learning work to address the issue of labeling cost in cryo-ET analysis tasks, which integrates two principled query heuristics in one optimization framework to make the queried subtomograms both representative and discriminative. 3) In HAL, we adopt several effective strategies to improve the performance, such as proposing a convolutional discriminator to learn the comparative metric of representations from shallower layers, introducing sub-sampling to improve the diversity of every query batch. 4) The empirical results for subtomogram classification using both simulated and real data demonstrate that we are able to achieve comparable testing performance (on average only $< 3\%$ prediction accuracy drop) while significantly reducing the labeling burden by over 70%.

2 Proposed Approach

Our Hybrid Active Learning method integrates two principles. For the representativeness principle, we develop an alternative optimization scheme for training the multi-class subtomogram classification model and a discriminator. Specifically, given a small initial set of labeled subtomograms, the classification model is firstly trained in a supervised way. Then, the hidden representations of both the labeled and unlabeled subtomogram samples are extracted to train a binary classification model (i.e. the discriminator). Thus, the probability scores of the unlabeled subtomogram samples are obtained from the predictions of the discriminator. Meanwhile, the uncertainty score (i.e. the entropy of the predictions from the




Figure 2: The full active learning scheme for subtomogram classification in our HAL framework.

multi-class classification model) is further fused with the discriminator score to produce the final query metric for ranking the unlabeled subtomograms. Afterwards, the top subtomograms are selected and labeled for iterative training until the budget is reached.

2.1 Representativeness Principle

In this part, the multi-class subtomogram classification model starts with a sparsely labeled dataset $\mathcal{D}$. Within the dataset, we denote the labeled subtomograms at iteration t as $\mathcal{LD}(t)$ and the unlabeled subtomograms as $\mathcal{UD}(t)$. Then we have $\mathcal{LD}(t) \cup \mathcal{UD}(t) = \mathcal{D}$ and $\mathcal{LD}(t) \cap \mathcal{UD}(t) = \emptyset$. From the domain adaptation point of view, we treat $\mathcal{LD}(t)$ and $\mathcal{UD}(t)$ as two separate domains, namely the source domain $\mathcal{L}$ and target domain $\mathcal{U}$ respectively. $M(\cdot)$ is defined as the feature extractor and $D(\cdot)$ is the introduced discriminator which aims to distinguish these two domains.

At each iteration t , to enhance the sample representativeness, we first train the main classifier using the softmax cross entropy loss. Then we extract the representations from the intermediate layers on both $\mathcal{LD}(t)$ and $\mathcal{UD}(t)$ and regard them as inputs to the discriminator. Next, we train this discriminator by a binary classification task so that it can discriminate the labeled and unlabeled subtomograms well. If we assume the output of the discriminator $D(\cdot)$ to be 0 for labeled class and 1 for unlabeled class, then we select and label a batch of subtomogram samples $B(t)$ which satisfy:

$$B(t) = \arg \max_{x \in \mathcal{UD}(t)} \Pr(D(M(x)) = 1 | M(x)), \quad (1)$$

where $B(t)$ is the queried unlabeled batch at iteration t .

Why a discriminator? The reason behind is if we can determine with high probability that an unlabeled subtomogram is from $\mathcal{UD}$, then it should be different from $\mathcal{LD}$, which is helpful for improving the information encoded in the labeled dataset and thus better for the model to generalize on the remaining unlabeled subtomogram examples after we label it. Otherwise, if the subtomogram examples from $\mathcal{UD}$ are indistinguishable from $\mathcal{LD}$, then we successfully represent the distribution with $\mathcal{LD}$. This motivates us to design a discriminator D for such probability estimation and alignment. Moreover, the introduced discriminator is expected to provide more flexibility during classification and the subtomogram sample selection since it has a learnable metric for separating the labeled and unlabeled subtomogram examples, which is better than hand-crafted metric designs [47, 50].

2.1.1 Task-specific Designs

In this section, we proposed two task-specific designs to further refine the capacity of the representativeness principle, namely the convolutional discriminator and the subset sampling strategy.

Commonly, the discriminator for unsupervised domain adaptation [15] often regards the outputs from the fully connected layers as the input. They claim such design will focus on more fine-grained information for feature adaptation since these layers of the network extract and propagate more specific

features. However, these fine-grained features are more suitable for multi-class classification which is usually biased for the discriminator, especially for highly heterogeneous 3D cryo-ET data. Instead, we propose to use the output of the last max-pooling layer as the input for the discriminator and enhance the discriminator with the more flexible convolutional operations. Specifically, the convolutional discriminator consists of two convolutional layers followed by two fully connected layers (Fig. 3). These convolution operations enable our model to learn a flexible representation space and a task-specific comparison metric for binary classification, which is helpful for querying valuable subtomograms more effectively.

In addition, recall that querying unlabeled subtomogram samples requires iteratively training the multi-class classification model and expert annotation in a loop. Therefore, the querying efficiency is important. One simple solution is to query subtomograms in larger batches instead of one subtomogram at a time, which reduces the waiting time until the classifier finishes training [3]. However, since we select the data in a large batch which are all predicted by the discriminator with a high probability to be unlabeled, they tend to have a similar distribution, especially for the subtomograms captured with a higher noise level. This scenario causes significant information overlap. To mitigate this, we emphasize on the diversity of the sampled subtomograms in a batch by assuming consecutive mini-queries will be less likely to contain similar instances. We split the original queried batch $B(t)$ into m sub-batches. Suppose we desire to select K subtomograms at iteration t , we first train the discriminator on the representations until convergence and label the top $\frac{K}{m}$ subtomograms. Then we repeat the process by interleaving the discriminator training and subtomogram selection until K subtomogram samples are queried. During this process, we only train the main classifier once but train the discriminator for m times which is more efficient. The detailed architecture of our model is demonstrated in Fig. 3.

2.1.2 Theoretical Analysis

The motivation of the representativeness principle is to label the most appropriate data from the unlabeled subset $\mathcal{UD}$ that can represent the distribution of the training (or the entire) dataset as well as possible. In this case, a classifier trained on $\mathcal{LD}$ should perform similarly compared to that trained with the entire dataset $\mathcal{D}$ labeled. Naturally, we are interested in how to measure the distribution difference between two observations $x \sim \mathcal{L}$ and $x \sim \mathcal{U}$ and see if the design of a discriminator can achieve less classification error on $\mathcal{U}$. Without loss of generality, we use $\mathcal{H}\Delta\mathcal{H}$ divergence [24] $d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{L}, \mathcal{U})$ for distribution difference estimation, which measures the maximum difference of the probabilities for inconsistent prediction.

Denote ϵ_d , ϵ_U and ϵ_L to be the classification error of the discriminator and the multi-class classification model on the unlabeled and labeled subtomogram samples, respectively, we argue ϵ_U is bounded by a term related to ϵ_L and ϵ_d by Theorem 1.

Theorem 1. Assume the complexity of the discriminator is more than a XOR function, given $f(\cdot)$ a multi-class candidate classifier, the classification error on the unlabeled dataset is bounded by: $\epsilon_U(f) \leq \epsilon_L(f) + \epsilon_d + C$. C is an uncorrelated constant.

Proof Sketch. Following the proof and assumptions made in [6] and substituting the source and target domain as $\mathcal{L}$ and $\mathcal{U}$, we get

$$\epsilon_U(f) \leq \epsilon_L(f) + \frac{1}{2}d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{L}, \mathcal{U}) + C. \quad (2)$$

Then following the derivation in [15], we replace $d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{L}, \mathcal{U})$ by its upper bound

$$2 \sup_{\eta \in \mathcal{H}_d} |\Pr_{\mathcal{L}}[z : \eta(z) = 1] + \Pr_{\mathcal{U}}[z : \eta(z) = 0] - 1|, \quad (3)$$

which can be seen as $2\epsilon_d$. Here $\mathcal{H}_d$ is the function space for the discriminator. And after substitution, the theorem is proven. Here the assumption is easily satisfied since the discriminator is implemented by a neural network which is complex enough according to the Universal Approximation Theorem [4, 13].

Given such a guarantee, if $\epsilon_L(f)$ and ϵ_d is minimized, the classification error on the unlabeled dataset is bounded.

2.2 Discriminateness Principle

In addition to the introduced discriminator for improving the sample representativeness, we argue that the useful label information (i.e. inductive bias) is missing in the current query strategy, which is shown to be




Figure 3: The model architecture of our subtomogram classifier with detailed layer configuration. “ $64 \times 3 \times 3 \times 3$ -1-same” means 64 convolutional filters with $3 \times 3 \times 3$ kernel, 1 stride and “same” padding. “FC layer (1024)” means 1024 units in the fully connected layer. The shape of the final FC layer depends on the number of classes denoted as “FC layer (...)”. Max-Pooling (2-1) means $2 \times 2 \times 2$ filters with 1 stride. Dropout (0.7) means the dropout rate is 0.7 in that layer. ReLU and Softmax denote the activation function.

effective in literature [55, 14, 5]. Thus, we propose a hybrid query method by selecting discriminative subtomogram samples with uncertainty sampling [62].

The intuition is as follows: the representativeness principle assumes the unlabeled pool is large enough to represent the true distribution. However, the data from the sparse regions of distribution will be sampled because the unlabeled set gradually becomes not representative due to its decreasing size. Conversely, uncertainty sampling can keep a balance between labeled and unlabeled subtomograms on the representation space by selecting subtomograms corresponding to data density such that the classification will not easily be biased by the sparse region of the manifold [62].

On the other hand, uncertainty sampling is designed to sample the most uncertain instance which is closest to the decision boundary. Since the number of subtomograms in the initial stage is limited, the estimated decision boundary is far from the actual one. Therefore, it may select noisy instances and stuck at sub-optimal solutions due to a lack of exploration. In contrast, the semi-supervised setting in the discriminator-based query strategy can avoid this drawback by observing the entire dataset. Therefore, during training, the representativeness principle by the discriminator-based query and the discriminativeness principle by uncertainty sampling assist each other and further enhance the stability of query and classification performance.

Specifically, we use the entropy of the predictions to measure the uncertainty, which is formulated as:

$$E(t) = \arg \max_{x \in \mathcal{UD}(t)} \left[- \sum_{y \in C} P(y|x) \log P(y|x) \right], \quad (4)$$

where C denotes the class space. $P(y|x)$ denotes the conditional probability of y given x in the multi-class classifier.

2.3 Hybrid Active Learning

In order to tradeoff between the two principles, at each iteration, we design a ranking score for final selection criteria:

$$S(t) = \Pr(D(M(x)) = 1|M(x)) + \lambda E(t), \quad (5)$$

where λ is the weighting hyperparamter for balancing different scores. The other notations have the same meaning as Eqn. [1] and [4]. While there remain other score fusion methods, we argue our implementation is simple and effective enough to achieve sufficient application purposes.

3 Experiments and Results

3.1 Dataset and Preparation

We evaluate our method on two simulated and three real cryo-ET datasets. For simulation datasets, we utilize the PDB2VOL program [58] to generate 23 classes of subtomograms which have the same class space as [59] at two Signal-to-Noise Ratio (SNR) levels, including 0.03 (S1) and 0.05 (S2). These datasets are realistically simulated by approximating the true cryo-ET image reconstruction process through a tilt-angle of ± 60 , including the Contrast Transfer Function and Modulation Transfer Function. Each class contains 1,000 subtomograms with size of 40^3 voxels. These simulated datasets are used in our 23-class classification tasks.

For real datasets, we use a set of rat neuron tomograms from [18] (R1). For one tomogram, we manually select 1,800 subtomogram samples which contain particles of 28^3 voxels from 5424 subtomograms extracted by Difference of Gaussian (DoG) [36] (R1a). We evaluate the particle picking task for determining whether or not a sample contains a particle. This is formulated as a binary classification task for the multi-class classification model. We also extract 2,394 subtomograms with size of 40^3 in the same tomogram set. The 2,394 subtomograms contain 6 classes detected and classified by template matching (R1b) [18]. We evaluate the 6-class classification task on it. In addition, we process a 7-class dataset [40] (R2) from EMPIAR [22]. Each class contains 400 subtomograms with size of 28^3 . A 7-class classification task is evaluated on R2. The 2D $x - z$ center slice of the 3D images is demonstrated in Fig.

4.




Figure 4: Examples of used subtomograms in the form of iso-surface (bottom row) and center-sliced density map in parallel with the $x-z$ plane (first two rows). 5 out of 23 classes of simulated subtomograms are shown for simplicity. Their PDB IDs are below each image.

3.2 Results and Comparisons

In this section, we report the subtomogram classification result for both the simulated and real data. We start with 3% of the entire dataset as the labeled subtomogram samples for the simulated datasets $S1$, $S2$ and 3%, 4%, 4% for real datasets $R1a$, $R1b$ and $R2$, respectively. The effect of the number of the initially labeled subtomogram samples is shown in the next section. The query batch size is empirically fixed to 800 for the simulated datasets and 32 for the real datasets, which follows the common active learning setting [52]. In terms of the subset sampling for the simulated datasets, we report the model performance with the number of subset to be 20 while the number is 4, 8, 4 for the real datasets $R1a$, $R1b$ and $R2$, respectively. The effect of the number of the subset and subset size is deferred for discussion in the next

Method/Dataset	S1	S2	R1a	R1b	R2
Supervised Training	83.68 $\pm$ 0.24	95.36 $\pm$ 0.11	87.24 $\pm$ 0.09	81.53 $\pm$ 0.37	99.00 $\pm$ 0.02
HAL	80.23$\pm$1.47	93.96$\pm$0.42	85.48$\pm$0.56	74.80$\pm$0.33	95.00$\pm$0.94
Random Query [57]	74.77 $\pm$ 2.34	77.66 $\pm$ 0.98	67.00 $\pm$ 1.13	67.30 $\pm$ 0.22	85.85 $\pm$ 1.53
Uncertainty Query [23]	77.32 $\pm$ 1.00	90.78 $\pm$ 1.21	66.29 $\pm$ 0.80	70.00 $\pm$ 0.75	73.00 $\pm$ 0.28
Bayesian Query [14]	73.23 $\pm$ 2.64	85.29 $\pm$ 1.99	77.00 $\pm$ 0.89	70.70 $\pm$ 0.74	78.00 $\pm$ 0.55
CoreSet Query [43]	63.59 $\pm$ 2.62	63.48 $\pm$ 1.88	78.17 $\pm$ 0.42	62.60 $\pm$ 0.59	42.00 $\pm$ 0.92
BGAL [52]	78.23 $\pm$ 0.65	85.32 $\pm$ 1.01	82.18 $\pm$ 0.99	71.22 $\pm$ 1.33	88.34 $\pm$ 0.93
VAAL [48]	75.67 $\pm$ 0.91	86.51 $\pm$ 1.23	83.33 $\pm$ 1.08	69.72 $\pm$ 0.86	85.22 $\pm$ 0.85
EE-BMAL [64]	79.32 $\pm$ 0.77	89.11 $\pm$ 1.35	84.34 $\pm$ 1.00	71.62 $\pm$ 2.01	90.67 $\pm$ 1.05
BADGE [2]	79.46 $\pm$ 0.41	91.00 $\pm$ 0.73	82.63 $\pm$ 0.81	73.01 $\pm$ 0.54	93.21 $\pm$ 1.11
Labeled Percentage	23.87%	16.91%	11.89%	9.35%	12.00%

Table 1: Comparison of HAL and the baseline AL methods on five different datasets (results are the classification accuracy in %). The same labeling budget is used among different methods. The standard deviation is reported at the top right corner.

S1			S2			R1a			R1b			R2		
Config	Acc	T	Config	Acc	T	Config	Acc	T	Config	Acc	T	Config	Acc	T
400/2	75.96	1.9 h	400/2	83.62	1.6 h	32/1	80.25	0.2 h	32/1	63.50	0.2 h	32/1	90.86	0.4 h
80/10	77.96	3.8 h	80/10	85.20	2.7 h	16/2	81.23	0.4 h	16/2	71.74	0.4 h	16/2	90.69	0.4 h
40/20	80.23	5.4 h	40/20	93.96	4.6 h	8/4	85.48	0.5 h	8/4	70.08	0.4 h	8/4	95.00	0.5 h
20/40	78.82	8.0 h	20/40	86.65	6.6 h	4/8	83.34	0.6 h	4/8	74.80	0.5 h	4/8	94.11	0.6 h
10/80	78.56	9.1 h	10/80	89.73	8.4 h	2/16	84.62	0.9 h	2/16	74.09	0.7 h	2/16	93.28	0.7 h
2/400	79.49	14.7 h	2/400	91.08	10.6 h	1/32	80.19	1.0 h	1/32	73.21	0.8 h	1/32	88.62	0.7 h

Table 2: Comparative Results of different subset configurations on HAL (in %). T refers to the overall training time. Config is in the form of Subset Size/Number of Subset. h denotes hours.

section. We report the classification results after 7,5,6,5 and 8 query iterations in dataset S1, S2, R1a, R1b and R2, respectively since we empirically found more iterations will not bring significant improvement on HAL. We run all the baselines under the same setting and report all metrics using an average of 10 runs with random seed from 1 to 10.

In Tab. 1, we firstly compare with supervised training with the entire dataset labeled. In dataset S2, we use 16.91% labeled training data to achieve 93.86% test accuracy compared to 95.36% in fully supervised training. In dataset R1a, we use 11.89% of training data to achieve 85.48% test accuracy compared to 87.24% in fully supervised training. Moreover, we compared with 8 representative active learning baselines, including methods using a single query principle, namely, Random Query [57], Uncertainty Query [23], CoreSet Query [43], Bayesian Query [14], Bayesian Generative Active Learning [52] (BGAL) and hybrid query heuristics, exploration-exploitation BMAL [64] (EE-BMAL), VAAL [48] and BADGE [2]. As shown in Tab. 1, our method achieves a superior performance on all 5 different datasets under the same labeling budget. For single query principle, BGAL performs the best compared to others, especially on dataset R1a, which achieves a 82.18% final accuracy. Surprisingly, even with a theoretical guarantee, the core-set sampling performs the worst among baselines, which is possibly caused by the complex data distribution that makes it harder to cover the entire dataset with the constructed core-sets. Moreover, the baselines that adopt a hybrid query strategy usually performs better because of the mutual benefits of different criteria. However, they still underperforms our HAL which explicitly trade-offs the representativeness and the discriminativeness principle.

3.3 Ablation Study

To validate the effect of our task-specific designs and the hybrid query strategy, we did a controlled experiment that removes the convolutional layers in the discriminator (Variant 1), the subset sampling

Model/Dataset	S1	S2	R1a	R1b	R2
Variant 1 (V1)	78.45	91.87	83.65	72.27	94.28
Variant 2 (V2)	75.77	91.92	80.25	63.50	90.86
Variant 3 (V3)	75.32	88.63	79.99	70.40	91.38
HAL	80.23	93.96	85.48	74.80	95.00

Table 3: Ablation study results (In %) on the convolutional discriminator, subset sampling and the uncertainty sampling

Ratio/Dataset	S1	S2	R1a	R1b	R2
0.01	77.24	89.99	72.13	43.61	88.32
0.02	80.01	93.36	81.62	47.77	93.29
0.03	80.23	93.96	85.48	73.95	94.37
0.04	80.96	94.01	84.93	74.80	95.00
0.05	81.11	93.99	86.05	74.95	95.23

Table 4: Comparative results (In %) on the number of initially labeled subtomograms for five datasets.

(Variant 2) and the uncertainty sampling strategy (Variant 3), which is shown in Tab.3. Note that we did not remove the representativeness principle because that degenerates the model to the baseline method of Uncertainty Query. The training setting is the same as the previous section. According to Tab.3, removing any of the three parts will lead to performance drop. For example, removing the subset sampling will destabilize the training which decreases the accuracy from 74.80% to 63.50% on *R1b*.

3.4 The effect of initially labeled subtomograms

To observe the effect of the number of initially labeled subtomogram samples, we test our HAL under different ratios from 1% to 5% with the interval of 1% on five datasets while keeping the other training configuration unchanged. The comparative result is shown in Tab.4. As shown in Tab.4, the number of the initially labeled subtomogram samples have considerable effects on the final classification accuracy. The model is trained towards a sub-optimal direction if this ratio is too small, leading to much lower accuracy even though more subtomogram samples are labeled in the later stages. However, the initially labeled subtomogram samples are practically expensive to obtain. Thus, observing that the final accuracy does not increase too much if we keep increasing the number of initially labeled subtomogram samples, we empirically fix the ratio to be 3% for dataset *R1a*, *S1*, *S2*, and 4% for dataset *R1b*, *R2*.

3.5 The effect of subset configuration

During subset sampling, it is important to determine the optimal combination of the number of subsets m and the subset size $\frac{K}{m}$ in order to balance the training time and the diversity of the subset. Specifically, we test 6 different subset size for five datasets whose results are summarized in Tab.2. As can be observed, a balanced configuration of subset size and numbers can help achieve better performance. Meanwhile, the time complexity increases dramatically if the subset size is smaller since much more iterations are trained on the discriminator. Therefore, to balance time and accuracy, a moderate subset size is preferable.

3.6 Query Visualization

For a comprehensive discussion, we plot the comparative query process in Fig.5. Here we demonstrate the accuracy versus the number of labeled subtomogram samples during training. We can see the stability and the final classification accuracy of HAL is better without accuracy decrease or stagnation along the sampling procedure.

3.7 Implementation Details

Firstly, we normalize the data in every dataset. For classification on both simulated and real data, we randomly split the entire dataset to the train and test set by a ratio of 3:1. We set the learning rate and the batch size of the multi-class subtomogram classifier and the discriminator as 0.001, 128 and 0.01, 256 without further fine-grained tuning. We train the discriminator with early stopping when the accuracy reaches 98% in order to prevent overfitting. we set the λ in Eqn.5 to 1 because both scores have the same value range.




Figure 5: Comparative querying process with baselines (a) and ablations (b). The shaded area means the standard deviation. For simplicity, three of five datasets are shown.

4 Conclusion

Computational analysis, deep learning approaches in particular, has played an increasingly important role for obtaining molecular machinery insights from cryo-ET data. However, the heavy labeling work behind data-driven methods presents obstacles for biologists to use them as assistant approaches. In this paper, we present a novel active learning tool in the cryo-ET domain with concerns for limited labeling resources, which approaches the active learning objective by querying both representative and discriminative subtomogram samples. Our experimental results on both simulated and real data demonstrate it produces significantly improved test performance compared to baselines under the same labeling budget. Our method represents an important step towards fully utilizing deep learning for *in situ* recognition of macromolecules inside single cells captured by cryo-ET. It can potentially also be very useful for other biomedical research with limited labeling resources.

Funding

This work was supported in part by U.S. National Institutes of Health (NIH) grants P41GM103712, R01GM134020, and K01MH123896, U.S. National Science Foundation (NSF) grants DBI-1949629 and IIS-2007595, Mark Foundation 19-044-ASP, and AMD COVID-19 HPC Fund. XZ was supported by a fellowship from Carnegie Mellon University’s Center for Machine Learning and Health.

A. APPENDIX

A.1 Related Works

Deep Learning based cryo-ET Analysis

Since 2017, supervised deep learning started getting popular for cryo-ET analysis. [8] proposed a 2D CNN segmentation model for segmenting ultrastructures on 2D slices of a 3D tomogram. Supervised deep learning based methods have also been proposed for 3D subtomogram classification task thanks to its high-throughput processing capability. Specifically, [59] was the first to propose a deep learning approach to separate different structures in subtomograms. [34] introduced multi-task learning to learn a feature space for simultaneous classification, segmentation, and structural recovery of macromolecules. Other techniques, such as domain adaptation [65], open-set recognition [10] and semi-supervised learning [35] have also been applied for reliable and deployable analysis result. However, few of them take the labeling burden into account, which provided a time-consuming labeling work for domain experts and impeded their off-the-shelf usage.

Active Learning

Active learning has been popular prior to deep learning. It was usually used on small models. The uncertainty-based methods usually measure uncertainty by posterior probability of the predicted classes [29, 28] or the margin between the first predicted class and the second confidently predicted class [23, 41]. Other methods measure uncertainty via entropy [44, 37, 23] or the distance to the decision boundary for SVM [31, 51, 53]. Another direction that is based on the discriminative principle, such as query by committee, usually develops multiple models as a committee and uses the disagreement between them as the criteria for uncertainty estimation [38, 46]. The representativeness principle selects samples that can cover the distribution of the entire dataset by clustering [39], discrete optimization [63, 19, 11]. Other popular methods either focus on the neighboring information between samples [20, 1, 60] or the expected model change [45, 42, 12] for sample selection.

When deep learning comes in, several query heuristics are proposed on larger models and datasets, such as uncertainty based methods [32, 55, 14, 5], core-set sampling [43], generative models [48]. However, some of them, especially core-set sampling, are not scalable on very large datasets compared to our HAL because a large distance matrix from unlabeled samples is needed which makes the query procedure highly expensive. Some hybrid query approaches are also proposed. For instance, [64] selected the data by uncertainty sampling and random sampling while we use a discriminator to query representative samples in order to reduce sampling bias in uncertainty sampling. [2] selected samples whose gradients span a diverse directions, which did not explicitly consider different query heuristics. [47] used the Wasserstein metric for measuring the representativeness, which is hand-crafted metric compared to the learnable discriminator in HAL. [30] proposed to combine uncertainty with density but only queried one sample a time. The matrix inverse operation in density estimation is expensive for large 3D images. The others integrated generative models, such as adversarial training [67] or variational approaches [48] with other heuristics while HAL does not involve any adversarial learning procedure, which would lead to unstable training and data selection results.

There are also several works that applied active learning for biomedical images [49, 61, 66, 26] that are built upon conventional active learning approaches for natural images. They are either not considering the tradeoff between the two query principles or not applicable for 3D subtomogram classification tasks with high noise and transformation variations, which will lead to a sub-optimal subtomogram selection and classification performance.

References

- [1] Oisín Mac Aodha, Neill D. F. Campbell, Jan Kautz, and Gabriel J. Brostow. Hierarchical subquery evaluation for active learning on a graph. *arXiv preprint arXiv: 1504.08219*, 2015.

- [2] Jordan T. Ash, Chicheng Zhang, Akshay Krishnamurthy, John Langford, and Alekh Agarwal. Deep batch active learning by diverse, uncertain gradient lower bounds. In *ICLR*, 2020.
- [3] Javad Azimi, Alan Fern, Xiaoli Zhang Fern, Glencora Borradaile, and Brent Heeringa. Batch active learning via coordinated matching. In *ICML*, 2012.
- [4] Andrew R. Barron. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information Theory*, pages 930–945, 1993.
- [5] William H. Beluch, Tim Genewein, Andreas Nürnberger, and Jan M. Köhler. The power of ensembles for active learning in image classification. In *CVPR*, pages 9368–9377, 2018.
- [6] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine Learning*, 79(1-2):151–175, 2010.
- [7] Chengqian Che, Ruogu Lin, Xiangrui Zeng, Karim Elmaaroufi, John M. Galeotti, and Min Xu. Improved deep learning-based macromolecules structure classification from electron cryo-tomograms. *Machine Vision and Applications*, 29(8):1227–1236, 2018.
- [8] Muyuan Chen, Wei Dai, Stella Y Sun, Darius Jonasch, Cynthia Y He, Michael F Schmid, Wah Chiu, and Steven J Ludtke. Convolutional neural networks for automated annotation of cellular cryo-electron tomograms. *Nature Methods*, 14(10):983, 2017.
- [9] Sanjoy Dasgupta. Two faces of active learning. *Theoretical Computer Science*, 412(19):1767–1781, 2011.
- [10] Xuefeng Du, Xiangrui Zeng, Bo Zhou, Alex Singh, and Min Xu. Open-set recognition of unseen macromolecules in cellular electron cryo-tomograms by soft large margin centralized cosine loss. In *BMVC*, 2019.
- [11] Ehsan Elhamifar, Guillermo Sapiro, Allen Y. Yang, and S. Shankar Sastry. A convex optimization framework for active learning. In *ICCV*, pages 209–216, 2013.
- [12] Alexander Freytag, Erik Rodner, and Joachim Denzler. Selecting influential examples: Active learning with expected model output changes. In *ECCV*, pages 562–577, 2014.
- [13] Ken-ichi Funahashi. On the approximate realization of continuous mappings by neural networks. *Neural Networks*, pages 183–192, 1989.
- [14] Yarin Gal, Riashat Islam, and Zoubin Ghahramani. Deep bayesian active learning with image data. In *ICML*, pages 1183–1192, 2017.
- [15] Yaroslav Ganin and Victor S. Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, pages 1180–1189, 2015.
- [16] Ran Gilad-Bachrach, Amir Navot, and Naftali Tishby. Query by committee made real. In *NIPS*, pages 443–450, 2005.
- [17] Daniel Gissin and Shai Shalev-Shwartz. Discriminative active learning. *arXiv preprint arXiv:1907.06347*, 2019.
- [18] Qiang Guo, Carina Lehmer, Antonio Martinez-Sanchez, Till Rudack, Florian Beck, Hannelore Hartmann, Manuela Prez-Berlanga, Frdric Frottin, Mark S. Hipp, and F. Ulrich Hartl. In situ structure of neuronal c9orf72 poly-ga aggregates reveals proteasome recruitment. *Cell*, 172(4):696, 2018.
- [19] Yuhong Guo. Active instance sampling via matrix partition. In *NIPS*, pages 802–810, 2010.
- [20] Mahmudul Hasan and Amit K. Roy-Chowdhury. Context aware active learning of activity recognition models. In *ICCV*, pages 4543–4551, 2015.
- [21] Jiaji Huang, Rewon Child, Vinay Rao, Hairong Liu, Sanjeev Satheesh, and Adam Coates. Active learning for speech recognition: the power of gradients. *arXiv preprint arXiv:1612.03226*, 2016.
- [22] Andrii Iudin, Paul K Korir, José Salavert-Torres, Gerard J Kleywegt, and Ardan Patwardhan. Empiar: a public archive for raw electron microscopy image data. *Nature methods*, page 387, 2016.

- [23] Ajay J. Joshi, Fatih Porikli, and Nikolaos Papanikolopoulos. Multi-class active learning for image classification. In *CVPR*, pages 2372–2379, 2009.
- [24] Daniel Kifer, Shai Ben-David, and Johannes Gehrke. Detecting change in data streams. In *VLDB*, pages 180–191, 2004.
- [25] Kwan-Young Kim, Dongwon Park, Kwang In Kim, and Se Young Chun. Task-aware variational adversarial active learning. *arXiv preprint arXiv: 2002.04709*, 2020.
- [26] Weicheng Kuo, Christian Häne, Esther L. Yuh, Pratik Mukherjee, and Jitendra Malik. Cost-sensitive active learning for intracranial hemorrhage detection. In *MICCAI*, pages 715–723, 2018.
- [27] Sun-Kyung Lee and Jong-Hwan Kim. Bald-vae: Generative active learning based on the uncertainties of both labeled and unlabeled data. In *International Conference on Robot Intelligence Technology and Applications*, pages 6–11, 2019.
- [28] David D. Lewis. A sequential algorithm for training text classifiers: Corrigendum and additional data. *SIGIR*, pages 13–19, 1995.
- [29] David D. Lewis and Jason Catlett. Heterogeneous uncertainty sampling for supervised learning. In *ICML*, pages 148–156, 1994.
- [30] Xin Li and Yuhong Guo. Adaptive active learning for image classification. In *CVPR*, pages 859–866, 2013.
- [31] Xin Li and Yuhong Guo. Multi-level adaptive active learning for scene classification. In *ECCV*, pages 234–249, 2014.
- [32] Liang Lin, Keze Wang, Deyu Meng, Wangmeng Zuo, and Lei Zhang. Active self-paced learning for cost-effective and progressive face identification. *TPAMI*, pages 7–19, 2018.
- [33] Chang Liu, Xiangrui Zeng, Ruogu Lin, Xiaodan Liang, Zachary Freyberg, Eric Xing, and Min Xu. Deep learning based supervised semantic segmentation of electron cryo-subtomograms. In *ICIP*, pages 1578–1582, 2018.
- [34] Chang Liu, Xiangrui Zeng, Kaiwen Wang, Qiang Guo, and Min Xu. Multi-task learning for macromolecule classification, segmentation and coarse structural recovery in cryo-tomography. In *BMVC*, page 271, 2018.
- [35] Siyuan Liu, Xuefeng Du, Rong Xi, Fuya Xu, Xiangrui Zeng, Bo Zhou, and Min Xu. Semi-supervised macromolecule structural classification in cellular electron cryo-tomograms using 3d autoencoding classifier. In *BMVC*, 2019.
- [36] Pei Long, Min Xu, Zachary Frazier, and Frank Alber. Simulating cryo electron tomograms of crowded cell cytoplasm for assessment of automated particle picking. *BMC Bioinformatics*, 17(1):405, 2016.
- [37] Wenjie Luo, Alexander G. Schwing, and Raquel Urtasun. Latent structured active learning. In *NIPS*, pages 728–736, 2013.
- [38] Andrew McCallum and Kamal Nigam. Employing EM and pool-based active learning for text classification. In *ICML*, pages 350–358, 1998.
- [39] Hieu Tat Nguyen and Arnold W. M. Smeulders. Active learning using pre-clustering. In *ICML*, 2004.
- [40] Alex J Noble, Venkata P Dandey, Hui Wei, Julia Brasch, Jillian Chase, Priyamvada Acharya, Yong Zi Tan, Zhening Zhang, Laura Y Kim, Giovanna Scapin, Micah Rapp, Edward T Eng, William J Rice, Anchi Cheng, Carl J Negro, Lawrence Shapiro, Peter D Kwong, David Jeruzalmi, Amedee des Georges, Clinton S Potter, and Bridget Carragher. Routine single particle cryoem sample and grid characterization by tomography. *bioRxiv*, 2017.
- [41] Dan Roth and Kevin Small. Margin-based active learning for structured output spaces. In *ECML*, pages 413–424, 2006.
- [42] Nicholas Roy and Andrew McCallum. Toward optimal active learning through sampling estimation of error reduction. In *ICML*, pages 441–448, 2001.

- [43] Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach. In *ICLR*, 2018.
- [44] Burr Settles and Mark Craven. An analysis of active learning strategies for sequence labeling tasks. In *EMNLP*, pages 1070–1079, 2008.
- [45] Burr Settles, Mark Craven, and Soumya Ray. Multiple-instance active learning. In *NIPS*, pages 1289–1296, 2007.
- [46] H. Sebastian Seung, Manfred Opper, and Haim Sompolsky. Query by committee. In *COLT*, pages 287–294, 1992.
- [47] Changjian Shui, Fan Zhou, Christian Gagné, and Boyu Wang. Deep active learning: Unified and principled method for query and training. In *AISTATS*, pages 1308–1318, 2020.
- [48] Samarth Sinha, Sayna Ebrahimi, and Trevor Darrell. Variational adversarial active learning. In *ICCV*, pages 5971–5980, 2019.
- [49] Asim Smailagic, Pedro Costa, Hae Young Noh, Devesh Walawalkar, Kartik Khandelwal, Adrian Galdran, Mostafa Mirshekari, Jonathon Fagert, Susu Xu, Pei Zhang, and Aurélio J. C. Campilho. Medal: Accurate and robust deep active learning for medical image analysis. In *ICMLA*, pages 481–488, 2018.
- [50] Ying-Peng Tang and Sheng-Jun Huang. Self-paced active learning: Query the right thing at the right time. In *AAAI*, pages 5117–5124, 2019.
- [51] Simon Tong and Daphne Koller. Support vector machine active learning with applications to text classification. *JMLR*, pages 45–66, 2001.
- [52] Toan Tran, Thanh-Toan Do, Ian D. Reid, and Gustavo Carneiro. Bayesian generative active deep learning. In *ICML*, pages 6295–6304, 2019.
- [53] Sudheendra Vijayanarasimhan and Kristen Grauman. Large-scale live active learning: Training object detectors with crawled data and crowds. In *CVPR*, pages 1449–1456, 2011.
- [54] Hanmo Wang, Xiaojun Chang, Lei Shi, Yi Yang, and Yi-Dong Shen. Uncertainty sampling for action recognition via maximizing expected average precision. In *IJCAI*, pages 964–970, 2018.
- [55] Keze Wang, Dongyu Zhang, Ya Li, Ruimao Zhang, and Liang Lin. Cost-effective active learning for deep image classification. *TCSVT*, pages 2591–2600, 2017.
- [56] Zheng Wang and Jieping Ye. Querying discriminative and representative samples for batch mode active learning. In *SIGKDD*, pages 158–166, 2013.
- [57] Hoyoung Woo and Cheong Hee Park. An efficient active learning method based on random sampling and backward deletion. In *IScIDE*, pages 683–691, 2012.
- [58] Willy Wriggers, Ronald A. Milligan, and J. Andrew McCammon. Situs: A package for docking crystal structures into low-resolution maps from electron microscopy. *Journal of Structural Biology*, 125(2):185 – 195, 1999.
- [59] Min Xu, Xiaoqi Chai, Hariank Muthakana, Xiaodan Liang, Ge Yang, Tzviya Zeev-Ben-Mordehai, and Eric P. Xing. Deep learning-based subdivision approach for large scale macromolecules structure recovery from electron cryo tomograms. *Bioinformatics*, 33(14):i13–i22, 2017.
- [60] Zhao Xu, Kristian Kersting, and Thorsten Joachims. Fast active exploration for link-based preference learning using gaussian processes. In *ECML PKDD*, pages 499–514, 2010.
- [61] Lin Yang, Yizhe Zhang, Jianxu Chen, Siyuan Zhang, and Danny Z. Chen. Suggestive annotation: A deep active learning framework for biomedical image segmentation. In *MICCAI*, pages 399–407, 2017.
- [62] Yazhou Yang and Marco Loog. A benchmark and comparison of active learning for logistic regression. *Pattern Recognition*, 83(C):401–415, 2018.
- [63] Yi Yang, Zhigang Ma, Feiping Nie, Xiaojun Chang, and Alexander G. Hauptmann. Multi-class active learning by uncertainty sampling with diversity maximization. *IJCV*, pages 113–127, 2015.

- [64] Changchang Yin, Buyue Qian, Shilei Cao, Xiaoyu Li, Jishang Wei, Qinghua Zheng, and Ian Davidson. Deep similarity-based batch mode active learning with exploration-exploitation. In *ICDM*, pages 575–584, 2017.
- [65] Liangyong Yu, Ran Li, Xiangrui Zeng, Hongyi Wang, Jie Jin, Ge Yang, Rui Jiang, and Min Xu. Few shot domain adaptation for in situ macromolecule structural classification in cryo-electron tomograms. *arXiv preprint arXiv: 2007.15422*, 2020.
- [66] Zongwei Zhou, Jae Y. Shin, Lei Zhang, Suryakanth R. Gurudu, Michael B. Gotway, and Jianming Liang. Fine-tuning convolutional neural networks for biomedical image analysis: Actively and incrementally. In *CVPR*, pages 4761–4772, 2017.
- [67] Jia-Jie Zhu and José Bento. Generative adversarial active learning. *arXiv preprint arXiv: 1702.07956*, 2017.