

UC Irvine

UC Irvine Previously Published Works

Title

Active learning to classify macromolecular structures in situ for less supervision in cryo-electron tomography

Permalink

<https://escholarship.org/uc/item/48p267z3>

Journal

Bioinformatics, 37(16)

ISSN

1367-4803

Authors

Du, Xuefeng
Wang, Haohan
Zhu, Zhenxi
[et al.](#)

Publication Date

2021-08-25

DOI

10.1093/bioinformatics/btab123

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Structural bioinformatics

Active learning to classify macromolecular structures *in situ* for less supervision in cryo-electron tomography

Xuefeng Du¹, Haohan Wang², Zhenxi Zhu³, Xiangrui Zeng⁴, Yi-Wei Chang⁵,
Jing Zhang⁶, Eric Xing⁷ and Min Xu ^{4,*}

¹Department of Computer Science, University of Wisconsin-Madison, Madison, WI 53706, USA, ²Language Technologies Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA, ³Department of Computer Science, Beijing University of Posts and Telecommunications, Beijing 100876, China, ⁴Computational Biology Department, Carnegie Mellon University, Pittsburgh, PA 15213, USA, ⁵Department of Biochemistry and Biophysics, University of Pennsylvania, Philadelphia, PA 19104, USA, ⁶Department of Computer Science, University of California—Irvine, Irvine, CA 92697, USA and ⁷Machine Learning Department, Carnegie Mellon University, Pittsburgh, PA 15213, USA

*To whom correspondence should be addressed.

Associate Editor: Jinbo Xu

Received on November 27, 2020; revised on January 14, 2021; editorial decision on February 17, 2021; accepted on February 22, 2021

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 has played critical roles for such tasks. As supervised approaches, however, 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 labeling 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 very promising result for subtomogram classification task with limited labeling resources.

Availability and implementation: <https://github.com/xulabs/aitom>.

Contact: mxu1@cs.cmu.edu

Supplementary information: [Supplementary data](#) are available at *Bioinformatics* online.

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 (Che *et al.*, 2018; Du *et al.*, 2019; Liu *et al.*, 2019; Xu *et al.*, 2017), 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 1000 real subtomograms were used in (Liu *et al.*, 2018), 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 (Gissin and Shalev-Shwartz, 2019; Sener and Savarese, 2018) (Fig. 1). Two main principles for such unlabeled sample selection are proposed (Dasgupta, 2011) and they both have limitations: *discriminateness* and *representativeness*. The discriminateness principle aims to find the most discriminative samples for the current classifier, which will shrink the space of candidate classifiers as rapidly as possible (Wang and Ye, 2013). The popular proposed criteria are uncertainty rule (Wang *et al.*, 2018; Yang and Loog, 2018), expected error reduction (Huang *et al.*, 2016) and query by committee (Gilad-Bachrach *et al.*, 2005; Seung *et al.*, 1992). 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 (Wang and Ye, 2013), 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 (Nguyen and Smeulders, 2004) and generative models (Kim *et al.*, 2020; Lee and Kim, 2019; Sinha *et al.*, 2019; Tran *et al.*, 2019; Zhu and Bento,

2017). 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 (Wang and Ye, 2013).

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 discriminateness 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 (Ganin and Lempitsky, 2015), 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 discriminateness principle, We additionally introduce the label information by using the entropy of predictions as selection criterion. Such heuristic is a strong active learning baseline, namely uncertainty sampling (Yang and Loog, 2018). 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 (Ash *et al.*, 2020; Yin *et al.*, 2017) and we defer the discussion in the Supplementary Section S3. The contributions of this paper is summarized as follows: (i) 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). (ii) 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. (iii) 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. (iv) The empirical results for subtomogram classification using both simulated and

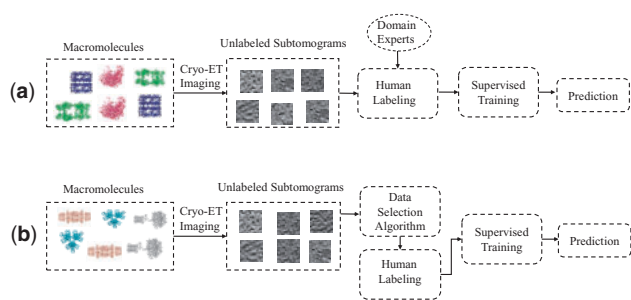


Fig. 1. Exemplary illustration of the active learning approach in cryo-ET classification. (a) The existing passive recognition pipeline where sufficient labeling is required where (b) demonstrates the recognition scenario guided by active learning

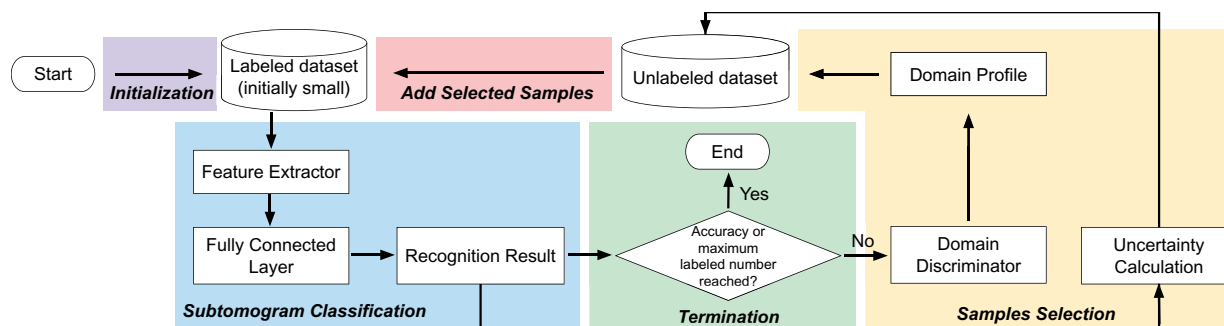


Fig. 2. The full active learning scheme for subtomogram classification in our HAL framework

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 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 $\ell\mathcal{D}(t)$ and the unlabeled subtomograms as $UD(t)$. Then we have $\ell\mathcal{D}(t) \cup UD(t) = \mathcal{D}$ and $\ell\mathcal{D}(t) \cap UD(t) = \emptyset$. From the domain adaptation point of view, we treat $\ell\mathcal{D}(t)$ and $UD(t)$ as two separate domains, namely the source domain 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 $\ell\mathcal{D}(t)$ and $UD(t)$ 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:

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

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 UD , then it should be different from $\ell\mathcal{D}$, 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 UD are indistinguishable from $\ell\mathcal{D}$, then we successfully represent the distribution with $\ell\mathcal{D}$. 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 (Shui et al., 2020; Tang and Huang, 2019).

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 (Ganin and Lempitsky, 2015) 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 (Azimi et al., 2012). 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 Figure 3.

2.1.2 Theoretical analysis

The motivation of the representativeness principle is to label the most appropriate data from the unlabeled subset UD that can represent the distribution of the training (or the entire) dataset as well as possible. In this case, a classifier trained on $\ell\mathcal{D}$ should perform similarly compared to that trained with the entire dataset \mathcal{D} labeled. Naturally, we are interested in how to measure the distribution




Fig. 3. The model architecture of our subtomogram classifier with detailed layer configuration. ‘ $64 \times 3 \times 3 \times 3$ -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

difference between two observations $x \sim \ell$ 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 (Kifer *et al.*, 2004) $d_{\mathcal{H}\Delta\mathcal{H}}(\ell, \mathcal{U})$ 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 (Ben-David *et al.*, 2010) and substituting the source and target domain as ℓ and \mathcal{U} , we get

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

Then following the derivation in (Ganin and Lempitsky, 2015), we replace $d_{\mathcal{H}\Delta\mathcal{H}}(\ell, \mathcal{U})$ by its upper bound

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

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 (Barron, 1993; Funahashi, 1989).

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 effective in literature (Beluch *et al.*, 2018; Gal *et al.*, 2017; Wang *et al.*, 2017). Thus, we propose a hybrid query method by selecting discriminative subtomogram samples with uncertainty sampling (Yang and Loog, 2018).

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 (Yang and Loog, 2018).

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:

$$E(t) = \operatorname{argmax}_{x \in \mathcal{U}D(t)} \left[- \sum_{y \in C} P(y|x) \log P(y|x) \right], \quad (4)$$

where C denotes the class space. $P(y|x)$ denotes the conditional probability of y given x in the multi-class classifier. We evaluate the quality of the model prediction $P(y|x)$ by comparing the class with the highest probability against the ground truth labels given by domain experts. We implement this evaluation in the neural networks by using the softmax loss function, which is the common practice in image classification.

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 hyperparameter 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 (Wriggers *et al.*, 1999) to generate 23 classes of subtomograms which have the same class space as (Xu *et al.*, 2017) 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 (Guo *et al.*, 2018) (R1). For one tomogram, we manually select 1800 subtomogram samples which contain particles of 28^3 voxels from 5424 subtomograms extracted by Difference of Gaussian (DoG) (Long *et al.*, 2016) (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 2394 subtomograms with size of 40^3 in the same tomogram set. The 2394 subtomograms contain 6 classes detected and classified by template matching (R1b) (Guo *et al.*, 2018). We evaluate the 6-class classification task on it. In addition, we process a 7-class dataset (Noble *et al.*, 2017) (R2) from EMPIAR (Iudin *et al.*, 2016). Each class contains 400 subtomograms with size of 28^3 . Following common practice in the active learning literature, no data augmentation techniques are used during training. The effect of data augmentation remains to be further explored.

A 7-class classification task is evaluated on R2. The 2D $x-z$ center slice of the 3D images and the iso-surface of the simulated datasets are demonstrated in Figure 4.

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 (Tran *et al.*, 2019). 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 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 Table 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

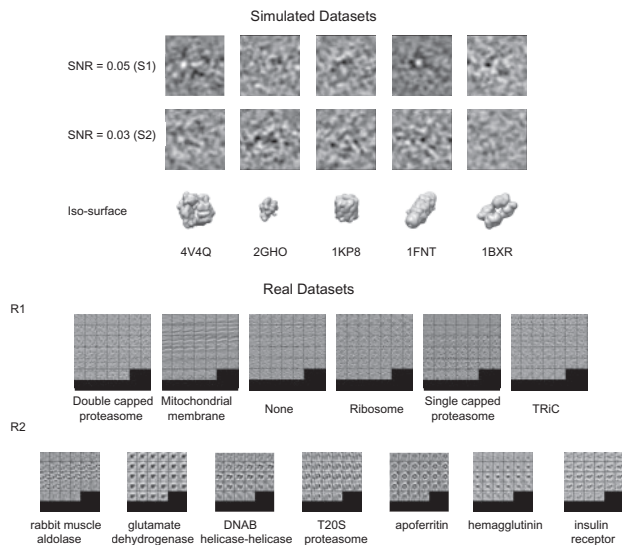


Fig. 4. Examples of used subtomograms. For simulated datasets, 5 out of 23 classes of simulated subtomograms are shown for simplicity. We plot in the form of iso-surface (bottom row) and center-sliced density map in parallel with the x - z plane (the first two rows). Their PDB IDs are below each image. For real datasets, we visualize the center-sliced density map in parallel with the x - z plane with the name of each macromolecule class below each image

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 (Woo and Park, 2012), Uncertainty Query (Joshi et al., 2009), CoreSet Query (Sener and Savarese, 2018), Bayesian Query (Gal et al., 2017), Bayesian Generative Active Learning (Tran et al., 2019) (BGAL) and hybrid query heuristics, exploration-exploitation BMAL (Yin et al., 2017) (EE-BMAL), VAAL (Sinha et al., 2019) and BADGE (Ash et al., 2020). As shown in Table 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 (Variant 2) and the uncertainty sampling strategy (Variant 3), which is shown in Table 2. Note that we did not remove the repre-

Table 2. Ablation study results (in %) on the convolutional discriminator, subset sampling and the uncertainty 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 1. Comparison of HAL and the baseline AL methods on five different datasets (results are the classification accuracy in %)

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 (Woo and Park, 2012)	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 (Joshi et al., 2009)	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 (Gal et al., 2017)	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 (Sener and Savarese, 2018)	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 (Tran et al., 2019)	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 (Sinha et al., 2019)	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 (Yin et al., 2017)	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 (Ash et al., 2020)	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%

Note: The same labeling budget is used among different methods. The standard deviation is reported at the top right corner.

sentativeness 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 Table 2, 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 Table 3.

As shown in Table 3, the number of the initially labeled subtomogram samples have considerable effects on the final classification accuracy. The model is trained toward a sub-optimal direction if

Table 3. Comparative results (in %) on the number of initially labeled subtomograms for five datasets

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

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 Table 4. 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 Figure 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.

Table 4. Comparative Results of different subset configurations on HAL (in %)

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

Note: T refers to the overall training time. Config is in the form of Subset Size/Number of Subset. h denotes hours.

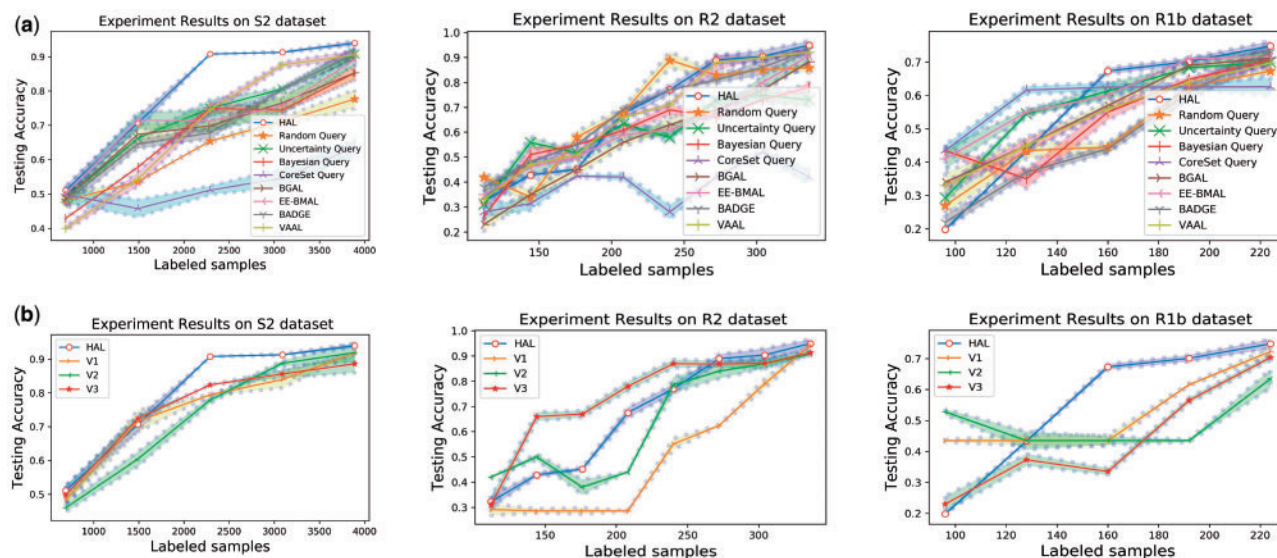


Fig. 5. Comparative querying process with baselines (a) and ablations (b). The shaded area means the standard deviation. For simplicity, three of five datasets are shown

Additional ablation study on the effect of λ in Eqn. 5 and implementation details can be found in Supplementary Sections S1 and S2.

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 toward 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) [P41GM103712, R01GM134020 and K01MH123896], U.S. National Science Foundation (NSF) [DBI-1949629 and IIS-2007595], Mark Foundation For Cancer Research 19-044-ASP and AMD COVID-19 HPC Fund. X.Z. was supported by a fellowship from Carnegie Mellon University's Center for Machine Learning and Health.

Conflict of Interest: none declared.

References

- Ash, J.T. et al. (2020) Deep batch active learning by diverse, uncertain gradient lower bounds. In: *International Conference on Learning Representations*. Addis Ababa, Ethiopia.
- Azimi, J. et al. (2012) Batch active learning via coordinated matching. In: *International Conference on Machine Learning*. Edinburgh, Scotland, UK.
- Barron, A.R. (1993) Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Trans. Inf. Theory*, **39**, 930–945.
- Beluch, W.H. et al. (2018) The power of ensembles for active learning in image classification. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Salt Lake City, UT, USA. pp. 9368–9377.
- Ben-David, S. et al. (2010) A theory of learning from different domains. *Mach. Learn.*, **79**, 151–175.
- Che, C. et al. (2018) Improved deep learning-based macromolecules structure classification from electron cryo-tomograms. *Mach. Vis. Appl.*, **29**, 1227–1236.
- Dasgupta, S. (2011) Two faces of active learning. *Theor. Comput. Sci.*, **412**, 1767–1781.
- Du, X. et al. (2019) Open-set recognition of unseen macromolecules in cellular electron cryo-tomograms by soft large margin centralized cosine loss. In: *British Machine Vision Conference*, Cardiff, UK. p. 148.
- Funahashi, K. (1989) On the approximate realization of continuous mappings by neural networks. *Neural Netw.*, **2**, 183–192.
- Gal, Y. et al. (2017) Deep bayesian active learning with image data. In: *International Conference on Machine Learning*, Sydney, NSW, Australia. pp. 1183–1192.
- Ganin, Y. and Lempitsky, V.S. (2015) Unsupervised domain adaptation by backpropagation. In: *International Conference on Machine Learning*, Lille, France. pp. 1180–1189.
- Gilad-Bachrach, R. et al. (2005) Query by committee made real. In: *Advances in Neural Information Processing Systems*, Vancouver, British Columbia, Canada. pp. 443–450.
- Gissin, D. and Shalev-Shwartz, S. (2019) Discriminative active learning. *arXiv preprint arXiv: 1907.06347*.
- Guo, Q. et al. (2018) In situ structure of neuronal c9orf72 poly-ga aggregates reveals proteasome recruitment. *Cell*, **172**, 696–705.e12.
- Huang, J. et al. (2016) Active learning for speech recognition: the power of gradients. *arXiv preprint arXiv: 1612.03226*.
- Iudin, A. et al. (2016) Empiar: a public archive for raw electron microscopy image data. *Nat. Methods*, **13**, 387–388.
- Joshi, A.J. et al. (2009) Multi-class active learning for image classification. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Miami, Florida, USA. pp. 2372–2379.
- Kifer, D. et al. (2004) Detecting change in data streams. In: *International Conference on Very Large Data Bases*, Toronto, Canada. pp. 180–191.
- Kim, K. et al. (2020) Task-aware variational adversarial active learning. *arXiv preprint arXiv: 2002.04709*.
- Lee, S.-K. and Kim, J.-H. (2019) Bald-vae: Generative active learning based on the uncertainties of both labeled and unlabeled data. In: *International Conference on Robot Intelligence Technology and Applications*, Daejeon, Korea (South). pp. 6–11.
- Liu, C. et al. (2018) Multi-task learning for macromolecule classification, segmentation and coarse structural recovery in cryo-tomography. In: *British Machine Vision Conference*, Newcastle, UK. p. 271.
- Liu, S. et al. (2019) Semi-supervised macromolecule structural classification in cellular electron cryo-tomograms using 3d autoencoding classifier. In: *British Machine Vision Conference*, Cardiff, UK. p. 30.
- Long, P. et al. (2016) Simulating cryo electron tomograms of crowded cell cytoplasm for assessment of automated particle picking. *BMC Bioinformatics*, **17**, 405.
- Nguyen, H.T. and Smeulders, A.W.M. (2004) Active learning using pre-clustering. In: *International Conference on Machine Learning*. Banff, Alberta, Canada.
- „Noble, A.J. et al. (2018) Routine single particle CryoEM sample and grid characterization by tomography. *Elife*, **7**, e34257.
- Sener, O. and Savarese, S. (2018) Active learning for convolutional neural networks: a core-set approach. In: *International Conference on Learning Representations*. Vancouver, BC, Canada.
- Seung, H.S. et al. (1992). Query by committee. In: *Annual Conference on Learning Theory*, Pittsburgh, PA, USA. pp. 287–294.
- Shui, C. et al. (2020) Deep active learning: Unified and principled method for query and training. In: *International Conference on Artificial Intelligence and Statistics*, Palermo, Sicily, Italy. pp. 1308–1318.
- Sinha, S. et al. (2019) Variational adversarial active learning. In: *International Conference on Computer Vision*, Seoul, Korea (South). pp. 5971–5980.
- Tang, Y. and Huang, S. (2019) Self-paced active learning: Query the right thing at the right time. In: *AAAI Conference on Artificial Intelligence*, Honolulu, Hawaii, USA. pp. 5117–5124.
- Tran, T. et al. (2019) Bayesian generative active deep learning. In: *International Conference on Machine Learning*, Long Beach, California, USA. pp. 6295–6304.
- Wang, H. et al. (2018) Uncertainty sampling for action recognition via maximizing expected average precision. In: *International Joint Conferences on Artificial Intelligence*, Stockholm, Sweden. pp. 964–970.
- Wang, K. et al. (2017) Cost-effective active learning for deep image classification. *IEEE Trans. Circuits Syst. Video Technol.*, **27**, 2591–2600.
- Wang, Z. and Ye, J. (2013) Querying discriminative and representative samples for batch mode active learning. In: *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Chicago, IL, USA. pp. 158–166.
- Woo, H. and Park, C.H. (2012) An efficient active learning method based on random sampling and backward deletion. In: *International Conference on Intelligence Science and Big Data Engineering*, Nanjing, China. pp. 683–691.
- Wriggers, W. et al. (1999) Situs: a package for docking crystal structures into low-resolution maps from electron microscopy. *J. Struct. Biol.*, **125**, 185–195.
- Xu, M. et al. (2017) Deep learning-based subdivision approach for large scale macromolecules structure recovery from electron cryo tomograms. *Bioinformatics*, **33**, i13–i22.
- Yang, Y. and Loog, M. (2018) A benchmark and comparison of active learning for logistic regression. *Pattern Recogn.*, **83**, 401–415.
- Yin, C. et al. (2017) Deep similarity-based batch mode active learning with exploration-exploitation. In: *International Conference on Data Mining*, New Orleans, LA, USA. pp. 575–584.
- Zhu, J. and Bento, J. (2017) Generative adversarial active learning. *arXiv preprint arXiv: 1702.07956*.