

Active Learning in Physics: From 101, to Progress, and Perspective

Yongcheng Ding,* José D. Martín-Guerrero, Yolanda Vives-Gilabert, and Xi Chen

Active learning (AL) is a family of machine learning (ML) algorithms that predates the current era of artificial intelligence. Unlike traditional approaches that require labeled samples for training, AL iteratively selects unlabeled samples to be annotated by an expert. This protocol aims to prioritize the most informative samples, leading to improved model performance compared to training with all labeled samples. In recent years, AL has gained increasing attention, particularly in the field of physics. This paper presents a comprehensive and accessible introduction to the theory of AL reviewing the latest advancements across various domains. Additionally, the potential integration of AL is explored with quantum ML, envisioning a synergistic fusion of these two fields rather than viewing AL as a mere extension of classical ML into the quantum realm.

1. Background

With the remarkable progress of machine learning (ML), the notion of “artificial intelligence for science” has become deeply ingrained in the minds of computer scientists. We are constantly enticed by the idea of tackling challenging physics problems with an omniscient model that seems to exist in our imagination. However, a significant obstacle stands in our path, and it primarily consists in the expensive labeling process in supervised learning. Labels are indeed expensive by any degree. For instance, in fields like computer vision or language processing, training sets often require a multitude of figures to be painstakingly labeled, typically done manually. This

has led to a demand for hiring individuals solely dedicated to this task or even the infamous reCAPTCHA tests we have endured for years. As physicists, our situation is even more daunting. The labels for our samples are usually derived from experiments and data analysis. This means we rely on highly skilled human annotators, research funding, and considerable time investments for both numerical simulations and laboratory experiments. To overcome this challenge, active learning (AL) emerges as a family of ML algorithms designed to alleviate the labeling cost, delivering astonishing results in various industries. The underlying concept is based on the hypothesis that not every sample in the dataset needs to be labeled to train a supervised ML model effectively. Instead, only a small subset of samples, specifically the most informative ones, can be labeled, resulting in a highly efficient training process with a significantly reduced training set. This naturally raises the question: how do we determine that samples are the most informative ones? Although we are initially uncertain about the labels of these samples, we can rely on the ML model itself to assess their informativeness. The model can evaluate each sample's potential for providing valuable information, and the specific criteria for making such judgments will be discussed later. With this workflow in mind, we can embark on the AL journey to train our models. It begins with initializing the model using a small set of labeled samples. Through an iterative process, we select the most informative unlabeled sample from the pool, present it to an oracle (e.g., a human annotator) for labeling, expand the training set for supervised learning, and update the model's parameters.^[1–3] This process continues until convergence is achieved or until the cost of acquiring another label becomes prohibitive. By embracing AL, we can surmount the challenges posed by expensive labeling and unleash the full potential of our ML models in the realm of physics and beyond.

Y. Ding, X. Chen
Department of Physical Chemistry
University of the Basque Country UPV/EHU
Apartado 644, Bilbao 48080, Spain
E-mail: jonzen.ding@gmail.com

Y. Ding
International Center of Quantum Artificial Intelligence for Science and
Technology (QuArtist) and Department of Physics
Shanghai University
Shanghai 200444, P. R. China

J. D. Martín-Guerrero, Y. Vives-Gilabert
Intelligent Data Analysis Laboratory (IDAL)
Department of Electronic Engineering
ETSE-UV
Universitat de València
Avinguda de la Universitat, s/n 46100 Burjassot, Valencia 46100, Spain
J. D. Martín-Guerrero
Valencian Graduate School and Research Network of Artificial Intelligence
(ValgrAI)
Valencia 46022, Spain

X. Chen
EHU Quantum Center
University of the Basque Country UPV/EHU
Barrio Sarriena, s/n 48940 Leioa, Biscay 48940, Spain

 The ORCID identification number(s) for the author(s) of this article can be found under <https://doi.org/10.1002/qute.202300208>

© 2023 The Authors. Advanced Quantum Technologies published by Wiley-VCH GmbH. This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

Correction added on 30 November 2023, after first online publication: The copyright line was changed.

DOI: 10.1002/qute.202300208

Delving into the details of the strategies used to select unlabeled samples for querying. It is important to note that there are numerous methods available for this task. However, in this discussion, we will focus on beginner-level strategies and reserve the exploration of advanced approaches for specific scenarios. By understanding these fundamental strategies, we can gain valuable insights into the initial steps of AL and lay a solid foundation for further exploration. The starting hypothesis is that we have a ML model characterized by its parameters θ , which is either initialized by randomization, or trained by a set of labeled samples $X = \{x_i, y_i\}_{i=1}^l$ with x_i be vectors as inputs and y_i be labels. We also have a pool of unlabeled samples $U = \{x_i\}_{i=l+1}^{l+u}$, which is much larger than the size of the training set.

1.1. Uncertainty Sampling

The most basic strategy in AL revolves around the belief that the sample with the greatest uncertainty will provide the most valuable information once labeled.^[4] To implement this strategy, we calculate the uncertainty of all unlabeled samples in the pool based on the model's estimation and select the candidate with the highest uncertainty for labeling by a human annotator. The key question here is: how do we quantify uncertainty using a ML model? One intuitive approach is known as the "least confidence" method, where the model estimates the probabilities associated with each potential label for a given sample. We query the sample according to the criteria

$$x_{LC} = \operatorname{argmax}_x [1 - P_\theta(y|x)] \quad (1)$$

$$\hat{y} = \operatorname{argmax}_y [P_\theta(y|x)] \quad (2)$$

where $\hat{y}$ is the class with the most confidence estimated by the model θ . Here, $P_\theta(\hat{y}|x)$ means the probability estimated by the model θ that a sample $x \in U$, belongs to the class y . Once the sample x_{LC} is selected and labeled by the human annotator, we transfer it from the pool U to the training set X , and update the parameters θ of the model.

1.2. Margin Sampling

One immediately notices that the least confidence only considers the information of the most probable class $\hat{y}_1$, i.e., information of other labels is lost. By taking the second most probable class $\hat{y}_2$ into account, one has the margin sampling as a more informative strategy

$$x_M = \operatorname{argmin}_x [P_\theta(\hat{y}_1|x) - P_\theta(\hat{y}_2|x)] \quad (3)$$

which stands for the experience that it is harder for one to classify ambiguous samples separated by smaller margin. By querying the sample with the minimal margin, one introduces more information to the training set to discriminate among all classes.^[4]

1.3. Entropy Sampling

According to Shannon's theory, one can characterize the amount of information by entropy after excluding miscellaneous information and averaging all uncertainties.^[4] Thus, one can employ

the information entropy that takes information from all classes as:

$$x_E = \operatorname{argmax}_x \left[- \sum_i P_\theta(\hat{y}_i|x) \log P_\theta(\hat{y}_i|x) \right] \quad (4)$$

For uncertainty sampling, these three strategies exhibit different behaviors across multiple class problems and various pools and training sets, particularly when the samples are imbalanced. However, it can be proven that they are equivalent when dealing with binary classification problems. To illustrate this point, let's consider a default support-vector-machine that seeks to separate the parameter space for binary classification. All these strategies would converge to select the sample that is closest to the decision boundary. This demonstrates the model-free nature of AL.

In **Figure 1**, we present the workflow of pool-based active learning with uncertainty sampling.

1.4. Query by Committee

Instead of relying on uncertainty sampling with a single model, the query by committee approach involves using a committee of models, assuming they are available for the problem. In this strategy, the candidate is selected based on consensus rather than uncertainty. Similar to entropy sampling, the committee utilizes voting entropy to determine the minimum consensus as:

$$x_{VE} = \operatorname{argmax}_x \left[- \sum_i \frac{V(y_i)}{C} \log \frac{V(y_i)}{C} \right] \quad (5)$$

where $V(y_i)$ is the vote results for label y_i from the committee of C models. One can also use the Kullback–Leibler divergence as the measure of minimal consensus as:

$$x_{KL} = \operatorname{argmax}_x \left[\frac{1}{C} \sum_{i=1}^C D(P|P_C) \right] \quad (6)$$

$$D(P_{\theta^{(i)}}|P_C) = \sum_i P_{\theta^{(i)}}(y_i|x) \log \frac{P_{\theta^{(i)}}(y_i|x)}{P_C(y_i|x)} \quad (7)$$

where $\theta^{(i)}$ is the i -th model in the committee and $P_C(y_i|x) = \frac{1}{C} \sum_{c=1}^C P_{\theta^{(c)}}(y_i|x)$ characterizes the probability of making an agreement on that y_i is the correct label.

The Kullback–Leibler divergence serves as an information-theoretic metric to quantify the disparity between two probability distributions. Hence, this measure of discrepancy prioritizes the most informative query by assessing the highest average distinction between the label distributions of any individual committee member and the consensus.^[5]

Based on the experience, query by committee usually outperforms uncertainty sampling by accuracy if the maximum size of the training set is bounded at the cost of extra computational resources required by the committee. Meanwhile, we have to emphasize that the models in the committee should be suitable for the task. Otherwise, near-to-random predictions by models degenerate the strategy to random sampling, as the baseline for




Figure 1. Schematic diagram of pool-based AL with uncertainty sampling or query by committee.

benchmarking the validity of AL, or even worse, a committee of wrong models might result in prediction that is overtaken by blind-guessing.

In Figure 1 we show the workflow of AL that covers the concept of uncertainty sampling or query by committee, which is also known as “pool-based AL”. In brief, the annotator for labeling the selected sample is compulsory. Later, we will see the variants that challenge the requirement, resulting in possible arguments over terminology.

2. Variants of Active Learning

According to the definition, an AL algorithm consists of a sampling strategy, iteratively updated training set and model, as well as an oracle. The term “active” gives rise to different variants of ML algorithms and has sparked impressive research in the field. While this review does not aim to engage in a debate about terminology, it is important to include and differentiate these variants to enhance our understanding of the various protocols.

2.1. Reinforcement Learning

An agent in reinforcement learning (RL) can take actions once it observes states based on the policy. It is quite natural that this algorithm is employed for solving quantum control problems such as state preparation or quantum gate design. Meanwhile, RL also allows the agent to actively propose quantum experiments to explore intricate quantum phenomena. To be more specific, researchers introduce an autonomous learning model that learns to design these experiments without relying on prior knowledge or human intuition.^[6] The model not only outperformed previous

approaches in terms of efficiency but also discovered new experimental techniques. By using the projective simulation model, they successfully designed photonic quantum experiments generating high-dimensional entangled multiphoton states, a topic of significant interest in modern quantum experiments. The AI system exhibited creativity by autonomously rediscovering experimental techniques that are now considered standard in contemporary quantum optical experiments. By showcasing the capability of ML, this study revealed how it has the potential to revolutionize the generation of experiments. Additionally, it emphasized the significant role played by intelligent machines in advancing scientific research through their assistance in experimental design, akin to other ML algorithms.^[7]

2.2. Semisupervised Learning

Semisupervised learning (SSL) follows a similar protocol to AL, starting with a small training set that is iteratively updated as samples are selected from the pool and labeled. It shares the goal of reducing labeling costs, but there is a crucial difference: instead of querying the most uncertain samples, it selects samples with the highest confidence in their labels based on the model’s estimation or cluster algorithms. These labeled samples are then incorporated into the training set. This approach allows semisupervised learning to train a model at a lower cost compared to AL, as it eliminates the need for external oracles, typically, human annotators. For example, regarding evaporative cooling experiments, researchers developed a ML-based approach to optimize experimental control.^[8] The method utilizes neural networks to learn the relationship between control parameters and the desired control goal, enabling the determination of optimal control parameters. A key challenge in this approach is the limited




Figure 2. Schematic diagrams for a) RL and b) SSL, which are often considered as AL in literatures.

availability of labeled data from experiments. To address this challenge, the authors employed semisupervised learning (according to our definition), which allows for overcoming data scarcity. The effectiveness of their scheme was demonstrated through the control of cooling experiments in cold atoms. Initially, the method was tested using simulated data, followed by its application to real experiments. The results showcased that the proposed method

achieved superior performance within a few hundred experimental runs. Notably, their approach does not require prior knowledge of the specific experimental system and holds applicability across different systems for experimental control.

In **Figure 2**, we illustrate the typical workflows of RL and SSL, which both significantly differ from pool-based AL in our definition. For RL, the agent interacts with the environment after

observing the state to alter its action for accumulating reward iteratively. In this way, it resembles the creature's behavior by being "active". For SSL, the job of annotator for labeling is replaced by the machine learning model itself according to its estimation with maximal confidence.

3. Progress

In recent years, AL has found applications in the field of physical science, particularly in material science and chemistry. It is worth noting that the connection between AL and quantum physics was initially proposed for accelerating AL through the principles of quantum mechanics.^[9] However, it is important to clarify that the definition of AL is rooted in RL, as we discussed in the previous section. Another interesting development is the application of quantum support vector machines for sampling against adversarial attacks, demonstrating the potential of quantum AL algorithms with polylogarithmic complexity.^[10] While the utilization of quantum algorithms to enhance AL is a compelling topic, our focus in this paper is on the applications of AL in the realm of physical science. Physicists often face challenges in optimizing classical algorithms with their physical insights, making it more feasible for them to leverage well-established theories from computer science for their research endeavors.

3.1. Quantum Information

Before delving into the practical applications of AL in the field of physical science, it is important to first introduce quantum information retrieval with AL. This particular area of research is well-suited for AL with the precise definition as presented in this paper. Moreover, it serves as a natural playground for AL due to the inherent cost of labeling, which can be characterized as fidelity loss resulting from quantum measurements. By exploring quantum information retrieval with AL, we can gain valuable insights and establish a solid foundation for further applications in other domains within the realm of physical science. Problems are proposed as binary and multinomial classification,^[11] which are easy to be solved by support vector machines as toy models. The only difference between the quantum version and the classical version is that the label of each sample is determined by the amplitudes on eigenbases, being encoded in two-level or multi-level quantum systems as quantum information for extraction via measurements with apparatuses. Although the problems are very simple, these allow one to benchmark AL of common strategies with random sampling as the baseline, demonstrating the feasibility by the rate estimation over 90% with only 5% samples labeled. Moreover, it offers a tool to characterize and compare the information extraction by von Neumann measurement and weak measurement. The von Neumann measurement extracts the information with higher signal-to-noise-ratio, collapses the wave function, and induces significant fidelity loss. In the contrary, weak measurement extracts less effective information with less perturbation on the wave function. By bounding the fidelity loss, AL trains the model with different information extraction protocol, and compares them by accuracy rate as figure-of-merit, implying the trade-off between copies, threshold of fidelity loss, and measurement strength. In the realm of more practical applications,

an active-learning-inspired algorithm has been proposed to estimate the state fidelity or gate fidelity by optimizing the settings of quantum measurements.^[12] This approach can be integrated into the training of variational quantum algorithms, enhancing their performance. Additionally, standard AL has made significant progress in the field of quantum state tomography,^[13] which is considered a milestone in quantum information retrieval. In this context, the sampling strategy involves the use of query by committee, while the ML model is based on a restricted Boltzmann machine. The success of this approach is evident in its adaptive reconstruction of the ground state of the XXZ model and kinetically constrained spin chain model, highlighting the power of AL in addressing complex problems in many-body physics.

3.2. High Energy Physics

The introduction of AL techniques holds significant potential in the field of high energy physics, encompassing disciplines such as particle physics, nuclear physics, and plasma physics (high-energy-density physics for the latter two). The experimental cost for studies in this field, either laboratory or numerical, is usually enormous, e.g., the accelerators, colliders, as well as the supercomputers. Meanwhile, one can train a ML model with AL to reduce the cost. For example, query by committee and its variant with dropout has been proposed to constrain the parameters of supersymmetry model in 19 dimensions,^[14] identify uncertain regions and steering new searches. For going beyond the Standard Model and Higgs bosons, an AL approach is presented for predicting the compatibility between new physics theories and existing experimental data obtained from particle colliders.^[15] This approach has achieved over 90% rate estimation using significantly fewer computing resources, accounting for <10% of the computational requirements of earlier methods, enabling the examination of previously untestable models and facilitates large-scale evaluations of new physics theories. In nuclear physics, margin sampling is applied to map out the thermodynamics stability of the quantum chromodynamics equation,^[16] which is developed by the Beam Energy Scan Theory collaboration. It utilizes a nonuniversal linear mapping of 3D Ising model variables onto the phase diagram, which involves four free parameters. Within the resulting 4D parameter space, certain combinations lead to unstable or acausal realizations of the equation. The AL framework focuses on the most crucial regions in the input parameter space. This approach efficiently identifies and eliminates unphysical equation of state instances with high accuracy. For soft matter physics, query by committee strategy with an ensemble of neural networks efficiently simulates the near-equilibrium plasma flows at multiscale.^[17] To overcome the computational cost of gathering information from molecular dynamics with kinetic theory, AL is employed to train neural networks using a small randomly sampled subset of the parameter space. The method is applied to investigate a plasma interfacial mixing problem relevant to warm dense matter, demonstrating significant computational efficiency compared to the full kinetic-molecular-dynamics approach. The results indicate that this approach enables the exploration of Coulomb coupling physics across a wide range of temperatures, and densities that are currently beyond the reach of existing theoretical models.

3.3. Condensed Matter Physics

AL has gained significant traction in this field, particularly due to its relevance in computational physics and material physics experiments. As the dimension of the Hilbert space expands exponentially with the system's size, computational resources become increasingly vital for Monte Carlo algorithms, density functional theory, and molecular dynamics. Moreover, sample preparation in material science is both costly and time-consuming. A subfield that benefits from AL is the computational approach to many-body physics. In addition to employing state tomography, the technique of query by committee is proposed to economically label data by utilizing an ensemble of neural networks to effectively fit the multidimensional function.^[18] Researchers outline the overall protocol of their fitting scheme and provide a detailed procedure for computing physical observables using the fitted functions. To demonstrate the effectiveness of the method, they present two examples: the quantum three-body problem in atomic physics and the calculation of anomalous Hall conductivity in condensed matter physics, yielding satisfactory outcomes in both cases. AL has shown promise in estimating phase boundaries and predicting exotic phenomena such as quantum phase transitions.^[19] In the context of the antiferromagnetic Ising model on a triangular lattice under a transverse field, support vector machines, employing Gaussian kernel and uncertainty sampling, successfully predict the boundaries between the ordered phase, Kosterlitz–Thouless phase, and paramagnetic phases. By leveraging only a small number of data points obtained from the quantum Monte Carlo algorithm, the predicted phase boundaries closely resemble the analytical solutions, demonstrating a remarkable level of similarity. In material science, the well-celebrated study has achieved the uniformly accurate interatomic potentials for materials simulation with deep potential generator as the variant of AL,^[20] which comprises three key components: exploration, generation of precise reference data, and training. The application of this procedure to sample systems such as Al, Mg, and Al–Mg alloys, showcases the capability of the model to create uniformly accurate potential energy surface models using a minimal amount of reference data. Furthermore, by combining AL and element embedding approach, the formation energies of oxygen vacancy layers, lattice parameters, and their statistical correlations in infinite-layer and perovskite oxides across the periodic table have been explored.^[21] By considering the Kullback–Leibler divergence, they show that the neural networks can predict these observables with 30% of the data in the pool, achieving high precision, while the AL algorithm compose the training set without human knowledge of chemistry. AL has estimated the material property curves and surfaces,^[22] focusing on the challenge of determining material properties in relation to independent variables, which typically involves time-consuming experiments or calculations. They have compared various sampling strategies based on directed exploration using a Kriging-based model across different materials problems of varying complexity. The atomic cluster models^[23] has also been studied: one based on the D-optimality criterion and the other utilizing ensemble learning. The extrapolation grade indicator facilitates active exploration of new structures, potentially leading to the automated discovery of rare-event configurations. The study also demonstrates the applicability of AL in exploring lo-

cal atomic environments using large-scale molecular dynamics simulations.

4. Perspective: Quantum Machine Learning

Here, we present our perspective on the extension of AL to the emerging field of quantum ML. The term “quantum machine learning” encompasses ML using quantum devices or quantum computers, including the utilization of artificial neural networks to accelerate ML models.^[24] Quantum ML holds great promise for enhancing the efficiency of ML processes. For example, it can be employed to analyze outputs from quantum physics experiments as a post-processing step, capitalizing on the availability of quantum states as inputs. Additionally, leveraging this paradigm to address classical problems and explore potential quantum advantages becomes intriguing when classical data can be encoded into quantum systems as quantum information. In both scenarios, AL plays a crucial role, as quantum states are fragile, being inevitable to collapse or perturbate when extracting information for labeling; besides, they are non-replicable due to the fundamental principles of quantum mechanics. Generating sufficient copies of quantum states for labeling or training necessitates repetitive data encoding or experimentation and storage in quantum memories, or on-demand creation upon query. Therefore, any technique that reduces the number of samples and labels required for training a quantum device or quantum neural network represents the silver bullet to practical quantum ML. To our best knowledge, the pioneering quantum AL is performed in a programmable photonic quantum processor.^[25] In this study, two AL-enabled variational quantum classifiers were designed and implemented. A programmable free-space photonic quantum processor was employed to execute these classifiers, comparing their performance with and without the query-by-committee strategy. The results demonstrated the significant advantages of AL in quantum ML, with labeling efforts reduced by up to 85% and computational efforts reduced by >90% in a data classification task. These findings highlight its potential and effectiveness in quantum ML, inspiring further applications in large-scale settings to enhance training efficiency and explore practical quantum advantages in quantum physics and real-world applications.

As the field of quantum ML continues to evolve, the development of specialized strategies becomes crucial for harnessing the full potential of AL in quantum systems.

4.1. Expected Model Change

We propose introducing an additional strategy in AL called “expected model change”. This strategy focuses on selecting samples that are expected to induce the greatest change in the model, provided we have knowledge of their labels. A specific example of this strategy is the “expected gradient length,” which can be applied to any learning problem that utilizes gradient-based training. Consequently, it can be effectively combined with a variational quantum circuit as the learning model, which is typically trained by minimizing a loss function based on the expectation values of observables using gradient methods. In other words, the AL process aims to query a sample, denoted as x , which, when

labeled and added to the training set X , will lead to the largest magnitude change in the training gradient. Let $\nabla l_\theta(X)$ be the gradient of the loss function with respect to the gate parameter, and $\nabla l_\theta(X \cup \{x, y\})$ be the updated gradient once the labeled sample $\{x, y\}$ is transferred to the training set. We select the sample the with maximal gradient length as:

$$x_{\text{EGL}} = \operatorname{argmax}_x \sum_i P_\theta(y_i|x) \operatorname{norm} [\nabla l_\theta(X \cup \{x, y_i\})] \quad (8)$$

where norm calculates the Euclidean norm of each gradient vector. The criteria can be further simplified by $\nabla l_\theta(X \cup \{x, y_i\}) \approx \nabla l_\theta(\{x, y_i\})$ since the loss function converged before the new query and the gradient of the norm of the loss function based on the last training set is almost zero. Empirical studies have shown the effectiveness of this approach, but it can be computationally demanding when dealing with large feature spaces and label sets.

Similar strategies, such as expected error reduction and variance reduction, are also applicable to quantum ML, but it is essential to conduct theoretical analyses on the loss function. In the case of expected error reduction, which is another decision theoretic approach, the goal is to estimate how much the generalization error of the model is likely to be reduced. This strategy calculates the expected future error of a model taught on the training set that includes each candidate individually and selects the sample with the least risk from the remaining unlabeled pool. It is important to note that directly minimizing the expectation of a quantum loss function is even more computationally expensive than in the classical case. Variance reduction, on the other hand, reduces generalization error by minimizing output variance, and in some cases, it can be solved analytically with a closed-form solution. Interestingly, the variance reduction method has already found applications in the field of material science, as we discussed earlier. These strategies offer valuable avenues for exploration and advancement in the field of quantum ML, but their implementation requires careful consideration of computational costs, theoretical analysis, and the specific requirements of the learning task at hand.

4.2. Density-Weighted Methods

The concepts of estimated error and variance reduction in AL offer a distinct approach by focusing on the entire input space rather than individual instances. This characteristic sets them apart from simpler query strategies like uncertainty sampling, query by committee, and expected gradient length, which are more susceptible to querying outliers. In certain cases, the least certain sample may lie on the classification boundary but may not be the most informative in terms of the overall data distribution. Query by committee and expected gradient length strategies may also exhibit similar behavior, spending time querying potential outliers due to their controversial nature or their expected impact on the model. Although estimated error and variance reduction strategies implicitly address these issues, they can be further improved by considering the density information during the query process. The key idea is that informative instances should not only be uncertain but also representative of the underlying distribution. By incorporating density information, a more com-

prehensive and effective AL strategy can be developed. By multiplying an additional term with density information to the basic query strategy $\phi(x)$, we write down the following criteria:

$$x = \operatorname{argmax}_x \phi(x) \left[\frac{1}{U} \sum \operatorname{sim}(x, x^{(u)}) \right]^\beta \quad (9)$$

which characterizes the average similarity of the sample to all other samples in the unlabeled pool, while its importance is controlled by β . The similarity can be characterized by various functions sim, resulting in different sampling behaviors.

The concept of density-weighted sampling in AL provides inspiration for incorporating physical insights, particularly those related to quantum information. When employing a variational quantum circuit for post-processing the outcomes of quantum experiments, one can consider the distance between experimental parameters producing the samples in the parameter space to guide the density-weighted method. Similarly, distances between classical unlabeled samples can be calculated prior to encoding them in a quantum system for the same purpose. However, in the realm of quantum mechanics, the infidelity between quantum states can be leveraged as a measure of information density, leading to a variant of quantum AL. Specifically, our main idea revolves around the notion that unlabeled quantum samples with higher average infidelity to the training set are more likely to introduce additional information when measured and labeled. By incorporating this infidelity-based density measure, we can enhance the effectiveness of quantum AL strategies. Thus, we have

$$x_{\text{IF}} = \operatorname{argmax}_x \phi(x) \left[\frac{1}{X} \sum_i \left(1 - |\langle x | x_i \rangle|^2 \right) \right]^\beta \quad (10)$$

which is in the language of quantum information theory.

5. Conclusion

In summary, this review has aimed to provide a comprehensive and accessible introduction to the concept of AL for physicists who may be interested in this field. We have highlighted the diverse range of applications of AL in various areas of physical science, spanning from quantum information to high-energy physics, nuclear physics, plasma physics, material science, and many-body physics. By showcasing these applications, we have demonstrated the wide-reaching potential of AL in advancing scientific research. Furthermore, we have presented our own perspective on the exciting intersection of AL and quantum ML based on our experience in both topics, reflected in two papers already published^[11,19] and ongoing research. In this context, where the cost of labeling becomes increasingly expensive, we believe that AL offers a promising approach to mitigate this challenge. We have proposed alternative strategies that specifically leverage the advantages of AL in the training of quantum neural networks. By combining the principles of AL with the power of quantum ML, researchers can harness the benefits of both fields and unlock new possibilities for tackling complex scientific problems. We hope that this review serves as a valuable resource and inspiration for physicists interested in exploring the potential of AL and its fusion with quantum ML. As this field continues to

evolve, we anticipate exciting advancements and fruitful collaborations between the realms of AL and quantum physics.

Acknowledgements

This work was financially supported by the EU FET Open Grant EPIQUS (899368); HORIZON-CL4-2022-QUANTUM-01-SGA project 101113946 OpenSuperQPlus100 of the EU Flagship on Quantum Technologies; the Basque Government through Grant No. IT1470-22; the project grant PID2021-126273NB-I00 funded by MCIN/AEI/10.13039/501100011033 and by “ERDFA way of making Europe” and “ERDF Invest in your Future”; the Nanoscale NMR and complex systems (PID2021-126694NB-C21) and QUANTEK project (Grant No. KK-2021/00070); the Valencian Government Grant with Reference Number CIAICO/2021/184; the Spanish Ministry of Economic Affairs and Digital Transformation through the QUANTUM ENIA project call—Quantum Spain project, and the European Union through the Recovery, Transformation and Resilience Plan—NextGenerationEU within the framework of the Digital Spain 2025 Agenda; the NSFC (12075145); STCSM (Grant No. 2019SHZDZX01-ZX04). X.C. acknowledges “Ayudas para contratos Ramón y Cajal”—2015–2020 (RYC-2017-22482). Y. D. acknowledges the sponsorship from Chinese Physical Society and MindSpore Quantum for quantum machine learning.

Conflict of Interest

The authors declare no conflict of interest.

Keywords

active learning, condensed matter physics, high energy physics, quantum information, quantum machine learning

Received: July 8, 2023
Revised: September 2, 2023
Published online:

- [1] B. Settles, Active Learning Literature Survey, University of Wisconsin–Madison, Madison, WI, USA **2009**.
- [2] Y. Sverchkov, M. Craven, *PLoS Comput. Biol.* **2017**, *13*, e1005466.
- [3] D. Tuia, M. Volpi, L. Copa, M. Kanevski, J. Munoz-Mari, *IEEE J. Sel. Top. Signal Process.* **2011**, *5*, 606.

- [4] J. Zhu, H. Wang, B. K. Tsou, M. Ma, *IEEE Trans. Audio Speech Lang. Process.* **2010**, *18*, 1323.
- [5] H. S. Seung, M. Oppen, H. Sompolinsky, Query by committee, presented at Proceedings of the ACM Workshop on Computational Learning Theory, Pittsburgh, PA, USA, July **1992**.
- [6] A. A. Melnikov, H. Poulsen Nautrup, M. Krenn, V. Dunjko, M. Tiersch, A. Zeilinger, H. J. Briegel, *Proc. Natl. Acad. Sci. USA* **2018**, *115*, 1221.
- [7] M. Krenn, M. Malik, R. Fickler, R. Lapkiewicz, A. Zeilinger, *Phys. Rev. Lett.* **2016**, *116*, 090405.
- [8] Y. Wu, Z. Meng, K. Wen, C. Mi, J. Zhang, H. Zhai, *Chinese Phys. Lett.* **2020**, *37*, 103201.
- [9] G. D. Paparo, V. Dunjko, A. Makmal, M. A. Martin-Delgado, H. J. Briegel, *Phys. Rev. X* **2014**, *4*, 031002.
- [10] P. A. M. Casares, M. A. Martin-Delgado, *New J. Phys.* **2020**, *22*, 073026.
- [11] Y. Ding, J. D. Martín-Guerrero, M. Sanz, R. Magdalena-Benedicto, X. Chen, E. Solano, *Phys. Rev. Lett.* **2020**, *124*, 140504.
- [12] R. Zhu, C. Pike-Burke, F. Mintert, Active Learning for Quantum Mechanical Measurements, *arXiv*, 2212.07513.
- [13] H. Lange, M. Kebrič, M. Buser, U. Schollwöck, F. Grusdt, A. Bohrdt, Adaptive Quantum State Tomography with Active Learning, *arXiv*, 2203.15719.
- [14] S. Caron, T. Heskes, S. Otten, B. Stienen, *Eur. Phys. J. C* **2019**, *79*, 944.
- [15] J. Rocamonde, L. Corpe, G. Zilgalvis, M. Avramidou, J. Butterworth, *SciPost Phys.* **2022**, *13*, 002.
- [16] D. Mroczek, M. Hjorth-Jensen, J. Noronha-Hostler, P. Parotto, C. Ratti, R. Vilalta, *Phys. Rev. C* **2023**, *107*, 054911.
- [17] A. Diaw, K. Barros, J. Haack, C. Junghans, B. Keenan, Y. W. Li, D. Livescu, N. Lubbers, M. Mckerns, R. S. Pavel, D. Rosenberger, I. Sagert, T. C. Germann, *Phys. Rev. E* **2020**, *102*, 023310.
- [18] J. Yao, Y. Wu, J. Koo, B. Yan, H. Zhai, *Phys. Rev. Res.* **2020**, *2*, 013287.
- [19] Y. Ding, J. D. Martín-Guerrero, Y. Song, R. Magdalena-Benedicto, X. Chen, *Phys. Rev. Res.* **2022**, *4*, 013213.
- [20] L. Zhang, D.-Y. Lin, H. Wang, R. Car, E. Weinan, *Phys. Rev. Mater.* **2019**, *3*, 023804.
- [21] A. Sahinovic, B. Geisler, *Phys. Rev. Mater.* **2021**, *3*, L042002.
- [22] Y. Tian, D. Xue, R. Yuan, Y. Zhou, X. Ding, J. Sun, T. Lookman, *Phys. Rev. Mater.* **2021**, *5*, 013802.
- [23] Y. Lysogorskiy, A. Bochkarev, M. Mrovec, R. Drautz, *Phys. Rev. Mater.* **2023**, *7*, 043801.
- [24] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, S. Lloyd, *Nature* **2017**, *549*, 195.
- [25] C. Ding, X.-Y. Xu, Y.-F. Niu, S. Zhang, W.-S. Bao, H.-L. Huang, *Quantum Sci. Technol.* **2023**, *8*, 035030.