

Harnessing imperfect quantum data through domain adaptation

Kosuke Ito^{1,2}, Akira Tanji³, Hiroshi Yano^{4,2}, Yudai Suzuki^{5,6,2}, and Naoki Yamamoto^{2,3}

This abstract combines the 144-qubit experiment of [1] with the UDA method and numerical results of [2].

Problem. Classical supervised learning with measurement data from quantum systems can extract useful physical information [3–5]. However, obtaining reliable teacher labels (e.g., quantum phases for phase classification tasks) is often difficult in the target regime of interest. In addition, algorithmic and experimental imperfections degrade the quality of input data. While unlabeled measurement data from imperfectly prepared target states may still be available, labeled data may be accessible only from restricted easier regimes and thus follow a different distribution. This raises the following question:

Can unlabeled measurement data of imperfectly prepared target quantum states be used together with differently distributed labeled data from more accessible regime to learn a useful predictor for the target regime?

We address this question by constructing a learning pipeline [2] with unsupervised domain adaptation (UDA) techniques (Sec. I) [6–9], and showing its efficacy through experimental [1] and numerical [2] demonstrations.

Main result. [1, 2] A key contribution is a 144-qubit real-device demonstration on the IBM Quantum processor `ibm_kawasaki` [1] (see Sec. II). Fig. 1 depicts the learning problem in this experiment. The task is to classify whether a measured state belongs to the topological or trivial phase on 144-qubit rotated surface code. Following Refs. [3, 4], each state is generated by applying a random circuit composed of geometrically local layers containing non-Clifford gates, to either the surface-code state $|0_L\rangle$ for the topological phase or the product state $|0\rangle^{\otimes 144}$ for the trivial phase. The learner has no access to the labels and receives only classical shadows of the *target* states prepared and measured on the hardware. Nevertheless, the learner can classically generate labeled measurement data for analogous examples by restricting the random circuits to Clifford gates; we refer to these as the *source* data.




Fig. 1. Setup for the learning task of the 144-qubit experiment.

Thus, the source and target data are differently distributed because the source data are generated by exactly simulated Clifford circuits, whereas the target data are obtained from non-Clifford circuits executed on noisy hardware. In fact, a model trained only on the source data fails to classify the target data, yielding a macro-F1 score of only 0.322. In contrast, our UDA approach achieves a macro-F1 score of 0.958, even though only unlabeled classical shadows of target data are available for training. Fig. 2 visualizes the learned latent representations using t-SNE [10, 11] for (a) source-only training and (b) our UDA method.




Fig. 2. 2D t-SNE visualization of the learned latent representations for (a) training with source data only and (b) UDA (proposal).

TABLE I. Macro-F1 scores on the unseen target data. Target labels are not used in training or model selection.

Task	Source $\rightarrow$ target shift	Source only	UDA
144-qubit real-device experiment			
Topological phase classification	Clifford, simulation $\rightarrow$ Non-Clifford, experimental data	0.322	0.958
Numerical simulations			
1D cluster model phase classification	Exact, solvable lines $\rightarrow$ Noisy, approximated, broad region	0.678	0.896
1D ANNNI model phase classification	Exact, low-field region $\rightarrow$ Noisy, approximated, broad region	0.531	0.830
Multipartite separable/entangled classification	Noise-free, fixed six-partition $\rightarrow$ Noisy, random tripartition	0.382	0.940

Fig. 2 (a) shows that a latent representation learned only from the source data does not separate the two phases consistently across the source and target domains. In the source domain, samples from the trivial phase form a separate cluster, whereas in the target domain trivial-phase samples largely overlap with the topological-phase samples. Consequently, most of the target samples from the trivial phase are misclassified as the topological phase. In contrast, our UDA method learns a representation in which distributions of source and target samples are well aligned within each phase, while the two phases remain separated (Fig. 2 (b)). This cross-domain alignment allows the source labels to guide classification of the unlabeled target data.

We also demonstrate the efficacy of our pipeline in other quantum phase classifications and multipartite-entanglement classification through numerical simulations [2], as summarized in Table I (see also Sec. III).

I. FORMULATION AND UDA PIPELINE

We formulate the learning problem within the UDA framework. Domains may differ in data-generation conditions such as the physical parameter region, state-preparation procedure, system size, or device. The target domain is the regime of interest, whereas the source domain is a related, more accessible regime with available labels. For each domain $d \in \{s, t\}$, let Q^d denote the joint distribution of a quantum state and label, with state marginal Q_ρ^d and conditional label distribution $Q_{Y|\rho}^d$. Domain shift, $Q^s \neq Q^t$, may involve $Q_\rho^s \neq Q_\rho^t$, $Q_{Y|\rho}^s \neq Q_{Y|\rho}^t$, or both. For quantum data, such shifts can reflect changes in the system regime and target-state imperfections, including hardware noise and algorithmic errors from limited quantum resources. In particular, a measured target state may be an imperfect realization of the ideal state that defines its label, so training only on labeled source data need not generalize to the target domain. UDA instead uses labeled source and unlabeled target data to learn a representation in which a source-supervised decision rule is better aligned with the target data [6–9].

The full pipeline is detailed in the technical version [2]. As a concrete measurement data, we use Pauli classical shadows [12–14]: for each state, T randomized Pauli measurements produce a classical shadow $S_T(\rho)$. We post-process it into a task-dependent feature tensor $z = \Phi(S_T(\rho))$ consisting of local Pauli expectation values, following shadow-based many-body learning [3, 4]. Our model uses a convolutional neural network feature extractor $G_{\theta_f}^d$ and a linear classifier C_{θ_Y} that outputs class probabilities $p_{\theta_f, \theta_Y}(\cdot|z) = C_{\theta_Y}(h)$, with $h = G_{\theta_f}^d(z)$. For UDA, we use a conditional domain adversarial network [8]. A domain discriminator learns to distinguish source from target using the label-conditioned representation $h \otimes p_{\theta_f, \theta_Y}(\cdot|z)$, while the feature extractor and classifier minimize the source-label loss and adversarially confuse it. The resulting representation is encouraged to remain class-discriminative while suppressing domain-specific spurious features.

II. REAL-DEVICE EXPERIMENT

As described above, we conduct a real-device experiment for our learning problem on the IBM Quantum Heron R2 processor `ibm_kawasaki`, accessed through IBM Quantum services [15]. The processor is composed of 156 fixed-frequency transmon qubits arranged in a heavy-hexagonal lattice, where each qubit is connected to two or three neighboring qubits via tunable couplers [16]. The experiment is designed to test whether our UDA method works for measurement data of real quantum hardware at scale. We use a 12×12 rotated surface-code patch with 144 physical qubits and classify the trivial and topological phases. The setup for the learning task is described above and in Fig. 1. The surface-code state $|0_L\rangle = 2^{-72} \prod_p (I + B_p) |0\rangle^{\otimes 144}$, with $B_p = \prod_{i \in p} X_i$, is prepared using a circuit similar to Ref. [17], while the trivial-phase initial state is the product state $|0\rangle^{\otimes 144}$. Since finite-depth geometrically local unitary circuits preserve intrinsic topological

order [18, 19], target samples are generated from each initial state by applying five-layer non-Clifford random circuits. In each layer, CNOT gates can act on either horizontal or vertical nearest-neighbor qubit pairs on the square lattice, with the two directions alternating layer by layer. For each such pair, a CNOT gate is applied independently with probability 0.5, and its control-target direction is chosen at random. The selected CNOT gates within a layer need not have disjoint supports. The target dataset contains 60 states per class, of which 40 per class are used as unlabeled target data for UDA, whereas the remaining 20 per class are excluded from the train data and used only for evaluating the final classification score.

In the hardware implementation, we do not redesign the square-lattice circuits to follow the processor connectivity. We simply use Qiskit [20] to map the original theoretical circuits to the device and transpile it to the native gate set, in which CZ is the only two-qubit gate. After transpilation, the average single- and two-qubit gate counts and two-qubit depth are (7715, 1890, 183) for trivial-phase circuits and (11097, 3108, 378) for topological-phase circuits, where the depth counts parallel layers of mutually disjoint CZ gates. For error mitigation, we use only dynamical decoupling and Pauli twirling.

The labeled source data are generated by classically simulating analogous five-layer Clifford circuits with 500 measurement shots per data point. Thus, the target and source data differ in both circuit family and data generation: target non-Clifford circuits are mapped, transpiled, and executed on the processor, whereas source Clifford circuits are simulated classically. As summarized in Table I and visualized in Fig. 2, the hardware target data yield a macro-F1 score of only 0.322 for a model trained only on labeled source data, whereas UDA reaches 0.958 without target labels.

III. NUMERICAL BENCHMARKS

We also test the same pipeline in numerical benchmarks, with the results summarized in Table I. First, we consider phase classification in the 15-qubit one-dimensional cluster and ANNNI spin chains. The source data are exact clean ground-state features from restricted Hamiltonian-parameter regions, whereas target data cover broader regions and include approximate state preparation, depolarizing and readout noise, and finite-shot classical shadows. Compared with training only on labeled source data, UDA improves the macro-F1 score from 0.678 to 0.896 for the cluster model and from 0.531 to 0.830 for the ANNNI model at 10^3 target shots per state.

Second, we consider classifying whether a 16-qubit state is separable or entangled with respect to some partitions of the qubits into subsystems. Among these numerical benchmarks, the largest observed gain over training only on labeled source data occurs in the following setting. In the target domain, the partitions are random tripartitions varying sample by sample. The target data are affected by depolarizing noise and measurement bit-flip error. The source domain instead uses a fixed six-partition across samples and no noise is added. Thus, the domains differ in both data quality and the partition structure defining the classes. With 500-shot classical shadows in both domains, training only on labeled source data yields a median macro-F1 of 0.382, whereas UDA reaches 0.940.

IV. DISCUSSION

Our results give a positive empirical answer to the question raised at the outset of this abstract: across our benchmarks, combining unlabeled measurements of imperfect target states with differently distributed labeled data from a more accessible regime can enable accurate target prediction. Our practical end-to-end pipeline takes experimentally feasible Pauli classical shadows as input, requires no target labels or post-measurement quantum processing, and succeeds on 144-qubit hardware data at an application-relevant scale.

These results motivate applying our UDA approach to extend the practical reach of imperfect experimental quantum data when reliable target labels are unavailable but related labeled data are accessible. Promising directions include data-driven quantum-device characterization and certification [21, 22] and learning phases or properties of strongly correlated systems from quantum-simulator or experimental measurements [23, 24], particularly when labeled data are difficult to obtain at scale or under target conditions, as in regimes where accurate classical calculations are difficult. Future work will pursue such applications and identify when quantum-data learning benefits from domain adaptation, given the need for transferable task-relevant structure across domains.

-
- [1] K. Ito, A. Tanji, H. Yano, Y. Suzuki, and N. Yamamoto, in preparation.
 - [2] K. Ito, A. Tanji, H. Yano, Y. Suzuki, and N. Yamamoto, Learning from imperfect quantum data via unsupervised domain adaptation with classical shadows, arXiv:2603.28294 (2026).
 - [3] H.-Y. Huang, R. Kueng, G. Torlai, V. V. Albert, and J. Preskill, Provably efficient machine learning for quantum many-body problems, *Science* **377**, eabk3333 (2022).
 - [4] G. Cho and D. Kim, Machine learning on quantum experimental data toward solving quantum many-body problems, *Nature Communications* **15**, 7552 (2024).
 - [5] B. Jaderberg, F. Shah, M. Jeon, M. E. Sahin, C. Zoufal, and K. Sharma, Learning ground state observables from quantum computing experiments, arXiv:2606.15983 (2026).
 - [6] S. Ben-David, J. Blitzer, K. Crammer, A. Kulesza, F. Pereira, and J. W. Vaughan, A theory of learning from different domains, *Machine Learning* **79**, 151 (2010).
 - [7] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky, Domain-adversarial training of neural networks, *Journal of Machine Learning Research* **17**, 1 (2016).
 - [8] M. Long, Z. Cao, J. Wang, and M. I. Jordan, Conditional adversarial domain adaptation, in *Advances in Neural Information Processing Systems*, Vol. 31 (2018) pp. 1647–1657.
 - [9] G. Wilson and D. J. Cook, A survey of unsupervised deep domain adaptation, *ACM Transactions on Intelligent Systems and Technology* **11**, 1 (2020).
 - [10] L. van der Maaten and G. Hinton, Visualizing data using t-sne, *Journal of Machine Learning Research* **9**, 2579 (2008).
 - [11] M. Wattenberg, F. Viégas, and I. Johnson, How to use t-sne effectively, *Distill* 10.23915/distill.00002 (2016).
 - [12] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nature Physics* **16**, 1050 (2020).
 - [13] A. Elben, S. T. Flammia, H.-Y. Huang, R. Kueng, J. Preskill, B. Vermersch, and P. Zoller, The randomized measurement toolbox, *Nature Reviews Physics* **5**, 9 (2023).
 - [14] T. Zhang, J. Sun, X.-X. Fang, X.-M. Zhang, X. Yuan, and H. Lu, Experimental quantum state measurement with classical shadows, *Physical Review Letters* **127**, 200501 (2021).
 - [15] IBM Quantum, <https://quantum.cloud.ibm.com/> (2026), accessed: 2026-06-09.
 - [16] IBM Quantum, Processor types, <https://quantum.cloud.ibm.com/docs/guides/processor-types> (2026), accessed: 2026-06-09.
 - [17] K. J. Satzinger, Y.-J. Liu, A. Smith, C. Knapp, M. Newman, C. Jones, Z. Chen, C. Quintana, X. Mi, A. Dunsworth, C. Gidney, I. Aleiner, F. Arute, K. Arya, J. Atalaya, R. Babbush, J. C. Bardin, R. Barends, J. Basso, A. Bengtsson, A. Bilmes, M. Broughton, B. B. Buckley, D. A. Buell, B. Burkett, N. Bushnell, B. Chiaro, R. Collins, W. Courtney, S. Demura, A. R. Derk, D. Eppens, C. Erickson, L. Faoro, E. Farhi, A. G. Fowler, B. Foxen, M. Giustina, A. Greene, J. A. Gross, M. P. Harrigan, S. D. Harrington, J. Hilton, S. Hong, T. Huang, W. J. Huggins, L. B. Ioffe, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, T. Khattar, S. Kim, P. V. Klimov, A. N. Korotkov, F. Kostritsa, D. Landhuis, P. Laptev, A. Locharla, E. Lucero, O. Martin, J. R. McClean, M. McEwen, K. C. Miao, M. Mohseni, S. Montazeri, W. Mruczkiewicz, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Y. Niu, T. E. O’Brien, A. Opremcak, B. Pató, A. Petukhov, N. C. Rubin, D. Sank, V. Shvarts, D. Strain, M. Szalay, B. Villalonga, T. C. White, Z. Yao, P. Yeh, J. Yoo, A. Zalcman, H. Neven, S. Boixo, A. Megrant, Y. Chen, J. Kelly, V. Smelyanskiy, A. Kitaev, M. Knap, F. Pollmann, and P. Roushan, Realizing topologically ordered states on a quantum processor, *Science* **374**, 1237 (2021).
 - [18] X. Chen, Z.-C. Gu, and X.-G. Wen, Local unitary transformation, long-range quantum entanglement, wave function renormalization, and topological order, *Phys. Rev. B* **82**, 155138 (2010).
 - [19] S. Bravyi, M. B. Hastings, and F. Verstraete, Lieb-robinson bounds and the generation of correlations and topological quantum order, *Phys. Rev. Lett.* **97**, 050401 (2006).
 - [20] A. Javadi-Abhari, M. Treinish, K. Krsulich, C. J. Wood, J. Lishman, J. Gacon, S. Martiel, P. D. Nation, L. S. Bishop, A. W. Cross, B. R. Johnson, and J. M. Gambetta, Quantum computing with Qiskit, arXiv:2405.08810 (2024).
 - [21] T. Xiao, J. Huang, H. Li, J. Fan, and G. Zeng, Intelligent certification for quantum simulators via machine learning, *npj Quantum Information* **8**, 138 (2022).
 - [22] É. Genois, J. A. Gross, A. Di Paolo, N. J. Stevenson, G. Koolstra, A. Hashim, I. Siddiqi, and A. Blais, Quantum-tailored machine-learning characterization of a superconducting qubit, *PRX Quantum* **2**, 040355 (2021).
 - [23] K. Ch’ng, J. Carrasquilla, R. G. Melko, and E. Khatami, Machine learning phases of strongly correlated fermions,

Physical Review X **7**, 031038 (2017).

- [24] A. Bohrdt, C. S. Chiu, G. Ji, M. Xu, D. Greif, M. Greiner, E. Demler, F. Grusdt, and M. Knap, Classifying snapshots of the doped hubbard model with machine learning, Nature Physics **15**, 921 (2019).

Author affiliations.

¹Advanced Material Engineering Division, Toyota Motor Corporation, Shizuoka 410-1193, Japan.

²Quantum Computing Center, Keio University, Yokohama 223-8522, Japan.

³Department of Applied Physics and Physico-Informatics, Keio University, Yokohama 223-8522, Japan.

⁴Toyota Central R&D Labs., Inc., 1-4-14 Koraku, Bunkyo-ku, Tokyo 112-0004, Japan.

⁵Institute of Physics, École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland.

⁶Centre for Quantum Science and Engineering, École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland.