Proposals
Prepare & Submit Proposals
Demo Site: Prepare Proposals
Proposal Status
Reviews & Meetings
Reviews, Panels, and Other Meetings
Reviewing for the National Science Foundation
Awards & Reporting
Project Reports
Demo Site: Project Reports
Notifications & Requests
Award Documents
Supplemental Funding Requests
Demo Site: Supplemental Funding Requests
Continuing Grant Increments Reports
PAR Research Products
Fellowships & Opportunities
Graduate Research Fellowship Program (GRFP) (Applicants, Fellows)
Graduate Research Fellowship Program (GRFP) (Officials)
Education and Training Application (ETAP)
Manage Financials
ACM$ (Award Cash Management $ervice)
Program Income Reporting
Foreign Financial Disclosure Report (FFDR)
Individual Banking
Grant Post-Award Financial Contacts
Administration
User Management

Preview of Award 2211841 - Annual Project Report

Cover |
Accomplishments |
Products |
Participants/Organizations |
Impacts |
Changes/Problems
| Special Requirements

Cover

Federal Agency and Organization Element to Which Report is Submitted:	4900
Federal Award or Other Identifying Number Assigned by Agency:	2211841
Project Title:	FET: Medium: A Quantum Computing Based Approach to Undirected Generative Machine Learning Models
PD/PI Name:	Samee U Khan, Principal Investigator Joseph Picone, Co-Principal Investigator
Recipient Organization:	Mississippi State University
Project/Grant Period:	10/01/2022 - 09/30/2025
Reporting Period:	10/01/2023 - 09/30/2024
Submitting Official (if other than PD\PI):	Samee U Khan Principal Investigator
Submission Date:	09/04/2024
Signature of Submitting Official (signature shall be submitted in accordance with agency specific instructions)	Samee U Khan

Accomplishments

*** What are the major goals of the project?**

The proposed research investigated the fundamental nature of quantum annealing (QA) and its potential applications in machine learning (ML). The primary objectives were twofold:

1:- Characterize Quantum Annealing Behavior: To ascertain the unique ability of QA to identify “difficult-to-find” local valleys (LVs) within a diverse range of configuration options (Markov Random Fields) trained with various application-relevant datasets.

2:- Harness Quantum Annealing for Deep Graphical Models: To exploit the LV-finding capabilities of QA to enhance the accuracy of classification and pattern reconstruction in deep graphical models.

Research Plan

To achieve these objectives, the research plan outlined three specific tasks:

1:- Expand Quantum Annealing Confirmation: To extend the confirmation of LV-related QA behavior to more intricate graphs and datasets, ensuring its applicability to a broader spectrum of ML problems.

2:- Develop a Hybrid Classical Quantum Sampling Algorithm: To create a hybrid algorithm that synergizes the strengths of classical and quantum computing, potentially leading to improved performance and efficiency.

3:- Apply Hybrid Classical Quantum Sampling to ML Problems: Evaluate the efficacy of the hybrid algorithm in addressing contemporary ML challenges, such as classification and pattern recognition.

Focus on Challenging ML Tasks

To demonstrate the potential of QA and hybrid algorithms in addressing real-world ML challenges, the proposal focused on two particularly demanding tasks:

1:- Seizure Detection in EEG Signals: This task involves learning complex temporal and spatial dependencies within electroencephalogram (EEG) data to detect seizures accurately. Identifying subtle patterns within noisy data is crucial for this application.

2:- Automatic Interpretation of Digital Pathology Images: This task requires analyzing huge high-resolution images (50K x 50K pixels) to identify anomalies within tissue samples. The ability to integrate local and global dependencies is essential for accurate interpretation.

By concentrating on these demanding ML problems, the research aimed to showcase the potential of QA and hybrid algorithms to address real-world challenges and contribute to advancements in the field.

*** What was accomplished under these goals and objectives (you must provide information for at least one of the 4 categories below)?**

Major Activities:

1:- Identify Ground States and Local Minima: Employ D-Wave's quantum annealing process to efficiently search for the global minimum (ground state) and local minima of optimization problems. This involves exploiting the quantum system's ability to explore multiple states simultaneously and to tunnel through energy barriers.

2:- Develop a D-Wave-Based RBM Training Algorithm: Create a novel training algorithm for Restricted Boltzmann Machines (RBMs) incorporating D-Wave's quantum sampling capabilities. Leveraging the quantum annealer's ability to sample from the Boltzmann distribution could accelerate RBM training and enhance their performance on various tasks.

3:- Apply D-Wave-Trained RBMs to Contemporary ML Datasets: Collaborate with Mississippi State University and Temple University researchers to prepare and curate contemporary machine learning datasets. These datasets would then be used to evaluate the effectiveness of the D-Wave-trained RBMs in solving real-world problems, such as image classification, natural language processing, and recommendation systems.

Specific Objectives:

Objective 1: Confirm the Behavior of D-Wave for More Complex Graphs and Datasets

Task 1 (Completed): Utilize D-Wave to identify unfamiliar local valleys that would be elusive even to an impractically lengthy classical search.

Task 2 (Completed): Discern the most promising local valleys through statistical analysis using a predefined set of criteria.

Task 3 (Continued): Statistically compare the basins of attraction for otherwise similar local valleys concerning their relative sizes.

Objective 2: Develop a Hybrid Classical Quantum (HCQ) Sampling Algorithm for RBM Training

Task 1 (Completed): Develop an approach for embedding larger graphs within the QAC hardware.

Task 2 (Postponed until Year 3): Develop an Approach for Embedding Deep Graphical Models (DGMs) within the QAC Hardware.

Task 4 (Initiated): Develop the Training Algorithm for Deep Belief Networks (DBMs).

Objective 3: Apply Quantum Annealing Computing to Contemporary Machine Learning Problems

We have conducted an extensive survey of long-term contextual models in machine learning based on transformer architectures and are developing baseline implementations in preparation for quantum computing implementations. Transformer-based architectures have emerged as a powerful method for modeling long-term context in time series. The temporal dependencies in time series data are crucial in many applications, such as EEG signals, digital pathology, and cardiology – the three domains we use as test cases in this project. The transformer architecture leverages self-attention (scaled dot-product attention) as its core mechanism. We have examined baseline implementations of several popular architectures to attempt to determine which is most suitable to the quantum computing paradigm:

· Transformer-XL: a foundational architecture that introduced recurrence mechanisms and relative positional encoding for improved long-range dependency modeling;

· Informer: a general-purpose model that utilizes ProbSparse self-attention and distillation to improve efficiency for long sequence forecasting;

· Autoformer: a model that enhances efficiency through a decomposition architecture and auto-correlation mechanism;

· Pyraformer: a model that introduces a pyramidal structure for multi-scale attention to capture both long- and short-term dependencies;

· Probabilistic Transformer: incorporates probabilistic modeling to quantify uncertainty in predictions;

· Non-stationary Transformers: models that address the challenges posed by non-stationary time series data.

We have also examined some specialized variations of these, namely LogTrans, InParformer, and Sageformer, which incorporate long sequences, personalized predictions, and external knowledge, respectively. Multivariate models such as Crossformer and Temporal Fusion Transformers (TFT) were also examined. Transformers designed for specific data representations like W-Transformers and privacy-preserving learning like FEDformer were also considered.

The Temporal Fusion Transformer (TFT) integrates several components to effectively handle different data types and temporal relationships. The core components include Gated Residual Networks (GRN), Variable Selection Networks, LSTM encoders, Multi-Head Attention, and Quantile forecasts. TFT uses recurrent layers for local processing and interpretable self-attention layers for long-term dependencies. The architecture includes specialized components to select relevant features and gating layers to suppress unnecessary components, enabling high performance in various scenarios. The architectural innovations include gating mechanisms that allow the model to adaptively manage its depth and complexity adaptively, enabling efficient information processing across different scenarios without overfitting to less relevant data components.

We believe this architecture is well-suited to problems such as EEG and EKG processing, where the signal consists of multichannel temporal signals. Long-term context in identifying key events like seizures plays a significant role in achieving high performance.

CrossFormer is an enhanced vision transformer leveraging cross-scale attention for improved performance in image classification, object detection, instance segmentation, and semantic segmentation tasks. It introduces a cross-scale embedding layer (CEL) and long-short distance attention (LSDA) for efficient feature processing across scales. Additionally, it addresses issues like self-attention map enlargement and amplitude explosion with progressive group size (PGS) and an amplitude cooling layer (ACL), respectively, in the improved version called Crossformer++.

Two additional innovations in CrossFormer++ are the Progressive Group Size (PGS) and the Amplitude Cooling Layer (ACL). PGS addresses the varying attention needs at different layers of the model by progressively adjusting the group size. This ensures that local features are emphasized in early layers while global ones are prioritized in deeper ones. ACL is introduced to manage the amplification of activation amplitudes across layers, which can destabilize training. By cooling down the amplitude, ACL helps maintain training stability and improve model performance.

CrossFormer++ models achieve a noticeable improvement in accuracy over existing vision transformers on image processing problems like our digital pathology data.

Significant Results: Objective 1: Confirmation of D-Wave Performance

We have confirmed that while the improved time-to-solution for finding the ground state (GS) in the latest version of D-Wave hardware is beneficial, it also leads to a slight degradation in performance when used to obtain a representative, diverse sample from probability distributions. We employed the same criterion for sample quality: the number of distinct local minima (LMs) found by D-Wave compared to classical Markov Chain Monte Carlo (MCMC).

Our research has unveiled a novel hypothesis that a reduced annealing time while hindering the ground state's discovery, produces a more diverse sample. This finding is of significant interest, particularly within the problem sizes we investigated. The probability of finding the GS versus annealing time exhibited an opposite trend of the number of LMs found by D-Wave. However, there was limited room to reduce the annealing time below the default value, resulting in only a minor (desirable) increase in the number of found LMs.

In parallel, we also conducted studies to confirm our understanding that D-Wave can be a platform for large-scale optimization problems. This is a crucial step because once the large-scale model of MCMC is deployed, a certain set of baseline expectations must be kept in mind.

We conducted a comprehensive and meticulous statistical comparison of D-Wave-found and classically-found LMs under conditions that yielded the most significant overlap between the two LM searches; over 80% of D-Wave LMs were missed by Gibbs sampling, and over 70% of Gibbs-found LMs were missed by D-Wave. However, there was nearly 100% overlap for LMs found by both techniques, with local minima states having the lowest energy (highest probability). The differences consistently resided in the medium- and high-energy (medium- and low-probability) regions of the LMs distribution. This thorough statistical comparison instills confidence in the rigor of our research.

Objective 2: Development of a Hybrid Classical Quantum (HCQ) Sampling Algorithm

The most promising approach, which was expected to elucidate the origin of the difference between D-Wave-based and classical sampling during RBM training, was to utilize D-Wave solutions as seeds for Gibbs sampling. To enable a more direct comparison with classical sampling during Contrastive Divergence (CD) RBM training, the following algorithm was developed:

In classical CD RBM training, each training pattern is used as a seed to initiate a Markov Chain. To employ a D-Wave-based sample comparably, a Markov Chain was commenced from a subset of LMs selected from the set of LMs found by D-Wave. This type of sampling was utilized in the "purely quantum" algorithm developed in this work.

Furthermore, we have initiated efforts to develop a hybrid quantum-classical algorithm, which holds great promise for the future of RBM training. This algorithm selects seeds for Gibbs sampling using training patterns and D-Wave-found LMs. Currently, we are experimenting with different approaches, all of which ensure the same number of seeds as in purely classical training at each training epoch. The central flexibility lies in handling cases where the number of D-Wave-found LMs exceeds the required number of seeds. Alternative approaches under investigation include selecting all the lowest-energy LMs, random selection, and, most promisingly, sampling the LMs according to the Boltzmann probability distribution concerning their energy. We expect the results of these experiments to be available in Year 3 of this project, offering a hopeful outlook for the future of quantum computing research.

Objective 3: Apply Quantum Annealing Computing to Contemporary Machine Learning Problems

Transformer-based architectures have shown great promise for time series analysis, but they also present challenges. One major challenge is the interpretability of these models, particularly understanding the attention mechanisms that drive their decision-making. While some models like Informer, Autoformer, and Pyraformer offer insights into feature importance, there is a need for more transparent and explainable methods, especially in models like Transformer-XL and the Probabilistic Transformer. Another challenge is scalability, as the efficiency of models like Transformer-XL and the Probabilistic Transformer can become a bottleneck when dealing with extremely long time series data. This calls for research into more efficient attention mechanisms or model architectures.

Incorporating domain knowledge is another area for improvement in models like Informer and the Probabilistic Transformer. While some models like Sageformer have started integrating external knowledge, there is potential for more sophisticated methods to leverage domain-specific information, such as features, constraints, or prior distributions. Additionally, real-world time series data often presents challenges like missing values and irregular sampling intervals, which most current models like Informer and Transformer-XL do not adequately address. Developing robust methods to handle such data is crucial.

Transfer learning and adaptability are also areas where further research is needed, particularly for specialized models like InParformer, Sageformer, and W-Transformer. While some models show promise in adapting to different domains, enhancing their ability to transfer knowledge and generalize across tasks would be valuable. Moreover, many models like Transformer-XL focus on offline forecasting, but real-time forecasting is essential in many applications. Adapting transformer architectures for real-time prediction with low latency and high accuracy is an open challenge. Ensuring the reliability of uncertainty estimates in probabilistic models like the Probabilistic Transformer and improving the robustness of all these models against adversarial attacks and data perturbations are important considerations for their deployment in critical applications.

Our baseline implementations of these algorithms are not yet providing the level of performance we expected. Work is underway to better understand and tune these complex algorithms. We have also begun examining how the training process for these algorithms can be accelerated using a search process that leverages quantum computing. Progress has been slower than expected.

Key outcomes or Other achievements:

*** What opportunities for training and professional development has the project provided?**

Graduate and undergraduate students were mentored, and work was reported at various venues.

*** Have the results been disseminated to communities of interest? If so, please provide details.**

Our is currently being disseminated through publications and open source code offerings on our project web site.

*** What do you plan to do during the next reporting period to accomplish the goals?**

The team will continue working on the objectives outlined above and start reporting out major studies at premier venues related to quantum computing.

Supporting Files

Filename	Description	Uploaded By	Uploaded On
fig_01.pdf	A typical transformer architecture used my many systems today.	Joseph Picone	09/03/2024
fig_02.pdf	A comparison of architectures	Joseph Picone	09/03/2024
fig_03.pdf	A summary of popular transformer architectures	Joseph Picone	09/03/2024

Products

Books

Book Chapters

Thundiyil, Saneesh Picone, Joseph (2024). Time Series Analysis from Classical Methods to Transformer-Based Approaches: A Review. *Machine Learning Applications in Medicine and Biology 1.* 1. Ahmed, Ammar Picone, Joseph. Springer. New York, New York, USA. 1. Status = AWAITING_PUBLICATION; Acknowledgement of Federal Support = Yes ; Peer Reviewed = Yes ; OTHER: N/A.

Inventions

Journals or Juried Conference Papers

View all journal publications currently available in the NSF Public Access Repository for this award.

The results in the NSF Public Access Repository will include a comprehensive listing of all journal publications recorded to date that are associated with this award.

Licenses

Other Conference Presentations / Papers

K. Ellenberger, D. Couch, J. Greer, N. Gregory, L. Sanchez, K. Love, Y. Koshka, and S. U. Khan (2024). *Quantum Annealing Task Mapping for Heterogeneous Computing Systems*. SPIE Photonics for Quantum. Waterloo, CA. Status = PUBLISHED; Acknowledgement of Federal Support = Yes

K. Ellenberger, D. Couch, J. Greer, N. Gregory, L. Sanchez, K. Love, Y. Koshka, and S. U. Khan (2023). *Quantum Task Mapping for Distributed Heterogeneous Computing Systems*. IEEE Quantum Week. Bellevue, WA, USA. Status = PUBLISHED; Acknowledgement of Federal Support = Yes

Other Products

Other Publications

Patent Applications

Technologies or Techniques

Thesis/Dissertations

Websites or Other Internet Sites

Participants/Organizations

What individuals have worked on the project?

Name	Most Senior Project Role	Nearest Person Month Worked
Khan, Samee	PD/PI	1
Picone, Joseph	Co PD/PI	1
Babaei, Mahdi	Graduate Student (research assistant)	8
El Yazizi, Abdelmoula	Graduate Student (research assistant)	2
Ellenberger, Kenzie	Graduate Student (research assistant)	12
Purba, Sadia	Graduate Student (research assistant)	1
Seifi, Somayeh	Graduate Student (research assistant)	2
Gregory, Noah	Undergraduate Student	2
Kamana, Hemanth	Undergraduate Student	1
Meng, Phuykong	Undergraduate Student	1
Ramos-Leiva, Brian	Undergraduate Student	1

Full details of individuals who have worked on the project:

Samee U Khan
Email: skhan@ece.msstate.edu
Most Senior Project Role: PD/PI
Nearest Person Month Worked: 1

Contribution to the Project: Advised the graduate students and managed the project.

Funding Support: Departmental support

Change in active other support: No

International Collaboration: No
International Travel: No

Joseph Picone
Email: joseph.picone@gmail.com
Most Senior Project Role: Co PD/PI
Nearest Person Month Worked: 1

Contribution to the Project: machine learning theory and algorithm development

Funding Support: None

Change in active other support: No

International Collaboration: No
International Travel: No

Mahdi Babaei
Email: mahdi.babaei@temple.edu
Most Senior Project Role: Graduate Student (research assistant)
Nearest Person Month Worked: 8

Contribution to the Project: Development of machine learning algorithms that are conducive to implementation on a quantum computer; benchmarking these algorithms on standard tasks.

Funding Support: None

International Collaboration: No
International Travel: No

Abdelmoula El Yazizi
Email: ae897@msstate.edu
Most Senior Project Role: Graduate Student (research assistant)
Nearest Person Month Worked: 2

Contribution to the Project: quantum computing implementations of machine learning

Funding Support: Departmental support

International Collaboration: No
International Travel: No

Kenzie Ellenberger
Email: comets45@hotmail.com
Most Senior Project Role: Graduate Student (research assistant)
Nearest Person Month Worked: 12

Contribution to the Project: software development to support quantum computing implementations

Funding Support: None

International Collaboration: No
International Travel: No

Sadia Purba
Email: tut44757@temple.edu
Most Senior Project Role: Graduate Student (research assistant)
Nearest Person Month Worked: 1

Contribution to the Project: research into new machine learning algorithms that support quantum computing

Funding Support: None

International Collaboration: No
International Travel: No

Somayeh Seifi
Email: somayeh.seifi1361@gmail.com
Most Senior Project Role: Graduate Student (research assistant)
Nearest Person Month Worked: 2

Contribution to the Project: machine learning algorithms

Funding Support: None

International Collaboration: No
International Travel: No

Noah Gregory
Email: njg146@msstate.edu
Most Senior Project Role: Undergraduate Student
Nearest Person Month Worked: 2

Contribution to the Project: Developed understanding of IPLs and utilizing of D-Wave for programming.

Funding Support: None.

International Collaboration: No
International Travel: No

Hemanth Kamana
Email: tuo73589@temple.edu
Most Senior Project Role: Undergraduate Student
Nearest Person Month Worked: 1

Contribution to the Project: software support for using quantum computers

Funding Support: None

International Collaboration: No
International Travel: No

Phuykong Meng
Email: phuykong.meng@temple.edu
Most Senior Project Role: Undergraduate Student
Nearest Person Month Worked: 1

Contribution to the Project: Software developer responsible for learning how to run jobs and program the DWAVE and IBM quantum computers.

Funding Support: None

International Collaboration: No
International Travel: No

Brian Ramos-Leiva
Email: tup86891@temple.edu
Most Senior Project Role: Undergraduate Student
Nearest Person Month Worked: 1

Contribution to the Project: software support for the development of machine learning simulations; quantum computing software development on the DWAVE system

Funding Support: None

International Collaboration: No
International Travel: No

What other organizations have been involved as partners?
Nothing to report.

Were other collaborators or contacts involved? If so, please provide details.
Nothing to report

Impacts

What is the impact on the development of the principal discipline(s) of the project?

In Year 2, our understanding of the limitations of quantum annealers, particularly the D-Wave, for sampling from complex probability distributions has significantly deepened. The primary statistical disparities between classical and QA-based sampling, now extended to the newest version of D-Wave hardware, reside in sampling medium and lower probability states. Whether these differences are sufficiently substantial for RBM training remains to be determined. As previously mentioned, these findings support the rationale for hybrid sampling approaches.

We have also improved our understanding of deep learning models involving long-term context, something that is crucial to being able to apply the limited quantum computing hardware capabilities to these algorithms. We are attempting to produce scaled down reference implementations of several popular transformer-based architectures in an attempt to better understand the computational bottlenecks and how performance is impacted by limited amounts of data. The parameter search space is very large for these complex models, so there is hope that quantum computing can find better solutions with fewer data points and fewer iterations.

What is the impact on other disciplines?
Nothing to report.

What is the impact on the development of human resources?

Several graduate and undergraduate students were involved in this project and benefited tremendously from the knowledge base that was discovered and shared with the community.

What was the impact on teaching and educational experiences?

The MSU investigators involved in this project are leveraging their experience working with students on quantum computing (QC) activities to refine the scope of the Introductory QC course currently being developed in collaboration between the Electrical and Computer Engineering (ECE) and Physics departments at MSU. By drawing upon the insights gained from hands-on student interactions, the project team can identify key areas of emphasis, potential challenges, and effective pedagogical approaches for the course. This iterative process ensures that the Introductory QC course is well-aligned with the needs and interests of students, fostering a stimulating and engaging learning environment.

The Temple investigators have introduced quantum computing lectures into two graduate-level courses: Introduction to Machine Learning and Pattern Recognition and Engineering Computation IV. The latter is a course that focuses on parallel computing using GPUs. Both courses currently are available to undergraduates as split-level courses.

The Temple investigators have also added some machine learning algorithms relevant to quantum computing to their educational open source software that teaches fundamental concepts in quantum computing.

What is the impact on physical resources that form infrastructure?
Nothing to report.

What is the impact on institutional resources that form infrastructure?
Nothing to report.

What is the impact on information resources that form infrastructure?
Nothing to report.

What is the impact on technology transfer?
Nothing to report.

What is the impact on society beyond science and technology?
Nothing to report.

What percentage of the award's budget was spent in a foreign country?
Nothing to report.

Changes/Problems

Changes in approach and reason for change

In Year 2, the project continued to provide substantial evidence corroborating the primary hypotheses outlined in the proposal concerning the behavior of the newer-generation quantum hardware, specifically the Pegasus D-Wave quantum annealer. The experimental findings consistently aligned with the anticipated characteristics of the quantum system, reinforcing the validity of the research directions and methodologies. No significant deviations from the original research plan indicated a strong alignment between the proposed objectives and the actual outcomes. These confirmations further solidified the confidence in the potential of quantum annealing for addressing complex optimization problems.

Actual or Anticipated problems or delays and actions or plans to resolve them

We had an issue recruiting students with the right skill sets in the previous year. However, that has been resolved, and we should be able to speed up our efforts in the right direction.

Changes that have a significant impact on expenditures

In previous phases of the project, the team relied heavily upon the complimentary monthly time allocation provided by D-Wave, which was restricted to one minute per user. However, as the project progressed and the demand for quantum annealing resources intensified, it became imperative to transition from the free time allocation to purchasing more substantial D-Wave time. This shift in resource allocation was driven by the increasing complexity of the research tasks and the necessity for extended computational time to achieve meaningful outcomes.

Despite the notable change in resource management, the team's focus on the project's core research questions and methodologies remained unwavering. They adapted their approach to optimize the utilization of the purchased quantum annealing resources. This shift in resource allocation allowed the team to delve more deeply into complex optimization problems and explore new avenues of research, ultimately contributing to the advancement of quantum computing and its applications, which is the goal of our project.

Significant changes in use or care of human subjects
Nothing to report.

Significant changes in use or care of vertebrate animals
Nothing to report.

Significant changes in use or care of biohazards
Nothing to report.

Change in primary performance site location
Nothing to report.

Special Requirements

Responses to any special reporting requirements specified in the award terms and conditions, as well as any award specific reporting requirements.

There were six cornerstones to our BPC plan: customer support, data annotation and development, software engineering, high school mentoring, participation in IEEE SPMB, summer outreach. Progress towards these goals is summarized below:

(1) **Customer Support:** we maintain a publicly accessible listserv, help@nedcdata.org, which we use to support our products. We typically handle about five external requests through this listserv per day. Most involve questions about the open-source data and software resources that we deliver. We usually resolve these requests within the same day they are received. We have not had many requests specifically for our quantum computing software. However, we have released packages, such as our ISIP Machine Learning Demo (IMLD), that include our baseline implementations of the algorithms described in our report.

(2) **Data Annotation and Development:** The two flagship applications we will use to validate our quantum approaches involve EEG and digital pathology data. We have added a third central application area – EKG signals used to diagnose cardiology problems. We acquired an open-source cardiology database containing 2.2M patients. We developed a baseline ML system for this task that delivers state-of-the-art performance. Concerning our EEG resources, we have released new versions of our major EEG corpora (e.g., v2.0.3 of our open-source TUH EEG Seizure Detection Corpus), mainly containing bug fixes (improved annotations). We have also completed the development of a new pathology corpus, the Fox Chase Cancer Center Breast Tissue subset, which contains 1,400+ annotated pathology images. This will allow some unique experiments about how quantum-derived models generalize to new corpora.

(3) **Software Engineering:** All personnel involved in the project have been trained on our software development process and are now developing software based on these guidelines. One of the more unique things about our lab is the strict software engineering process we use to build and release software. Students who learn this process and assimilate this coding style tend to successfully leverage these experiences into summer internships and permanent engineering positions. As described below, we have trained several high school students, two undergraduates, and one graduate student on our software engineering process.

(4) **High School Student Mentoring:** In Spring 2024, we hosted two high school students collaborating with Boys Latin High School (https://www.boyslatin.org/) in Philadelphia. In the Summer of 2024,, we hosted two high school students as part of the "Pathways to Temple" – Philly High School Summer STEM Research Program. Three of these four students were African American, and the fourth was an Asian female. Our undergraduate researchers mentored these students, which created an excellent synergy between these two groups, which are very close in age. We have also continued our relationship with a volunteer female high school student, Nidhi Ram, who has worked with us for three years. She graduated, and is

attending Columbia University to study mathematics and computer science. We began collaborating with another African-American female high school student who attends St. Francis Prep in New York. She will continue Nidhi's work on quantum computing applications.

(5) **The 2024 IEEE Signal Processing in Medicine and Biology Symposium (IEEE SPMB):** We are again hosting IEEE SPMB on December 7, 2024. We did not submit a paper on quantum computing to this year's conference, though we did submit papers in our three application areas. However, we completed the book chapter previously mentioned based on an abstract presented at IEEE SPMB 2023.

(6) **The Women in Engineering Summer Workshop:** The College of Engineering at Temple University has shifted its focus from this activity to a formal summer program where we host STEM students from underrepresented groups. These students typically come from high schools within the City of Philadelphia – both public and charter schools. It is a collaboration with Heights Philadelphia (https://heights.org/). We plan to continue hosting students every summer. These students are typically trained on Python programming and how to run machine learning experiments on a Linux cluster. In 2024, students were exposed to the fundamentals of quantum computing but did not progress far enough to run experiments.