SATHYAWAGEESWAR SUBRAMANIAN

+44 (0) 74887 02818 • ss2310@cam.ac.uk

PROFESSIONAL EXPERIENCE

Oct. 2023 - present	**Senior Research Associate**[1] (1851 Fellowship) University of Cambridge (UK)	*Department of Computer Science*
Sep. 2022 - Sep. 2023	**1851 Research Fellow**[2] Project title: Quantum algorithms for testing and learning	*DIMAP, Warwick*
Sep. 2020 - Aug. 2022	**Post Doctoral Research Associate** Division of Theory and Foundations & DIMAP University of Warwick, Coventry (UK)	*Department of Computer Science* Hosted by Prof. Tom Gur

EDUCATION

Oct. 2016 - Aug. 2020	**University of Cambridge - PhD in Applied Maths and Theoretical Physics** Supervisor: Prof. Richard Jozsa Thesis: *Quantum Algorithms for Matrix Problems and Machine Learning* Examiners: Dr. Iordanis Kerenidis & Dr. Sergii Strelchuk	
Oct. 2015 - Jun. 2016	**University of Cambridge - Part III of the Mathematical tripos** MASt in Applied Mathematics & Theoretical Physics First Class (Merit) Essay: *Verifying a Quantum Computation with Interactive Proof Systems*	
Aug. 2011 - Jun. 2015	**Indian Institute of Science, Bangalore** BSc. (Research)—Major: Physics, Minor: Mathematics First Class with Distinction Thesis: *Rotating and Magnetised White Dwarfs in General Relativity*	 GPA: 7.2/8

FUNDING AND SCHOLARSHIPS

2022 - 2025	Royal Commission for the Exhibition of 1851 Research Fellowship
2016 - 2020	Cambridge-India Ramanujan Scholarship for PhD at Cambridge
2015 - 2016	Full-cost Pemanda Monappa Scholarship for MASt. at Cambridge
2013	DAAD-WISE fellowship (Germany)
2010 - 2015	KVPY fellowship (DST, Govt. of India)

OTHER AWARDS

Sep. 2024	Invited as a Young Researcher to the 11$^{\text{th}}$ Heidelberg Laureate Forum
Sep. 2024	Invited to Dagstuhl Seminar 24381 on Computational Complexity
2023 - 2025	College Research Associateship, Sidney Sussex College, Cambridge
Jun. 2023	Warwick Faculty of Science, Engg. and Medicine PostDoc Prize (+cash award)
2019	Internship at NASA's *Quantum Artificial Intelligence Laboratory* (QuAIL, 3 months)
2018 - 2019	Visiting Student Researcher at University of Technology Sydney, Australia (2 months)
2012 - 2013	NIUS Physics fellowship (HBCSE - TIFR, India)
Jul. 2011	Harry Messel International Science School (University of Sydney)
Mar. 2009	C. Subramaniam award for Excellence in Character (Bharatiya Vidya Bhavan)
2006 - 2010	National Talent Search Scheme scholarship (NTSE - NCERT, Govt. of India)

[1] *Unestablished Assistant Professor* [2] *Assistant Professor (Research-Focused)* from 1 Aug. 2023.

TEACHING AND SUPERVISION

Lecturing (Cambridge)

- Quantum Computation, *Part III (year 4) of the Mathematical Tripos*
  - Lent term 2023-24 & Michaelmas term 2024-25 (16 lectures / 24 contact hours)
  - Setter and assessor for two Part III essays on Quantum Computation (3 students)

Supervisions (Cambridge)

- *Quantum Computing* : 2016-19 (20 students/year, Part II/year 3 CS)
- *Quantum Computing*: 2017-19 (10 students/year, Part II Maths)
- *Coding and Cryptography*: 2016-19 (6 students/year, Part II Maths)

PhD students

- Marcel Dall'Agnol, Jan. 2021-Oct. 2022
- Hugo Aaronson, Jan. 2023-present

Both co-supervised as PhD students of Prof. Tom Gur (Warwick & Cambridge)

Masters & Bachelors research projects

- Jeffrey Kam (MPhil Computer Science, Cambridge 2024-25)
- Ranjit Singh (MSci. Physics, Cambridge 2024-25)
- Shreyas Pandit (Part II Maths, Cambridge 2024)
- Léa Cassé (MS, University of Montpellier 2023)
- Sajib Sarker (BRAC University, Bangladesh 2019)

Interviewing

- Undergrad admissions to Maths Tripos, 2024 (Sidney Sussex College, Cambridge)
- Two PDRA hiring interviews, 2023 (Dept. of CS, Warwick)

COLLEGIALITY, LEADERSHIP & MANAGEMENT

2023-2024	Organiser of joint CQIF-CS seminars (Cambridge Maths+CS)
Jul. 2023	Local organiser for CCC'23 (with Prof. T Gur and Prof. A Czumaj)
2023	Member of the EDI MathSci Forum for Equality, Diversity, and Inclusion (Warwick)
2022-2023	Member and Postdoc coordinator of Warwick Quantum
2021-2023	Co-organiser of bi-weekly Complexity seminar series and group meetings (Warwick)
2018 - 2019	Co-organiser of weekly seminar series and group meetings for CQIF (Cambridge)
Apr 2018	Organiser of a Sarangi concert at the Ancient India & Iran Trust (Cambridge)
Apr 2018	Co-organiser of the Heilbronn Quantum Algorithms Meeting 2018
2016 - 2017	Music Officer of the *Cambridge University Indian Classical Arts Society* (CUICAS)

SCIENCE COMMUNICATION, OUTREACH, & PROFESSIONAL SERVICE

Selected talks

- Physics Quantum Information Seminar, University of Minnesota, December 2024.
- QuSoft Seminar, October 2024.
- Quantinuum Friday Seminar (Cambridge), August 2024.
- 37th Conference on Learning Theory, COLT 2024.
- CQT CS Seminar (Centre for Quantum Technologies, Singapore), April 2023.
- Bangalore Theory Seminars, IISc and Microsoft Research Bangalore (2023).
- QIP 2022 (Caltech).
- HonHai Quantum Computing Research Centre, Taipei (2021)
- NCTS Annual Theory Meeting, Taipei (2021).

Refereeing

- **Journals**: Quantum Info. & Comp. (QIC), Quantum, ACM Transactions on Computing Theory (TOCT), IEEE Transactions on Info. Theory, J. Phy. A: Theoretical and Mathematical, Phys. Rev. A, Nature Comms., Physica Scripta
- **Conferences**: AQIS, QIP, TQC, QCTIP, FOCS, ITCS, Crypto, CCC, ICALP...
- **Grant reviews**: US Department of Energy (DoE) Office of Science High Energy Physics (HEP) - Quantum Information Science Enabled Discovery (QuantISED 2.0, 2024), Leverhulme Trust Research Project Grants (2023), EPSRC Quantum Technology Career Development Fellowship (2022)

Outreach

- Invited talk to a popular audience at the Cambridge Commonwealth European and International Trust (April 2018)
- Book reviews for plus.maths.org (2015 - 2017)

- Quantum Channel Testing in Average-Case Distance;
  G Rosenthal, H Aaronson, *SS*, A Datta, T Gur. arXiv:2409.12566.

- Quantum Circuits surpass Biased Threshold Circuits in Constant-Depth;
  MH Hsieh, L Mendes, M Oliveira, *SS*.
  19th Theory of Quantum Computation, Communication and Cryptography – TQC 2024
  21th International Conference on Quantum Physics and Logic – QPL 2024
  Accepted for publication in Nature Communications.

- Information-theoretic generalization bounds for learning from quantum data;
  M Caro, T Gur, C Rouzé, D Stilck França, *SS*.
  19th Theory of Quantum Computation, Communication and Cryptography – TQC 2024
  37th Annual Conference on Learning Theory – COLT 2024 (Edmonton).

- Quantum complexity of the Kronecker coefficients[3];
  S Bravyi, A Chowdhury, D Gosset, V Havlicek, C Ikenmeyer, *SS*, G Zhu.
  27th Conference on Quantum Information Processing – QIP 2024 (Taipei).

- Quantum Ridgelet Transform: Winning the Neural Network Lottery Ticket;
  H Yamasaki, *SS*, S Hayakawa, S Sonoda.
  40th International Conference on Machine Learning – ICML 2023 + IBIS 2022.

- Quantum Worst-Case to Average-Case Reductions for All Linear Problems;
  VR Asadi, A Golovnev, T Gur, I Shinkar, *SS*.
  26th Conference on Quantum Information Processing – QIP 2023 (Ghent)
  35th ACM-SIAM Symposium on Discrete Algorithms – SODA 2024.

- Testing graph states with constant time and sample complexity;
  H Yamasaki & *SS*. arXiv:2201.11127.

- Sublinear quantum algorithms for estimating von Neumann entropy;
  T Gur, MH Hsieh, *SS*. arXiv:2111.11139.
  25th Conference on Quantum Information Processing – QIP 2022 (Caltech).

- Quantum algorithm for estimating α-Renyi entropies of quantum states;
  SS & MH Hsieh. *Physical Review A* **104**, 022428 (Aug. 2021).

- A quantum search decoder for Natural Language Processing;
  J Bausch, *SS*, S Piddock. *Quantum Machine Intelligence* **3**, 16 (Apr. 2021).

- Learning with optimized random features: Exponential Speedup by Quantum ML;
  H Yamasaki, *SS*, S Sonoda, M Koashi.
  Advances in Neural Information Processing Systems 33 – NeurIPS (Dec. 2020).

- Implementing smooth functions of a Hermitian matrix on a quantum computer;
  SS, S Brierley, R Jozsa. *Journal of Physics Communications*, 3(**6**), 65002 (Oct. 2019).

- Spectral sparsification of matrix inputs as a preprocessing step for quantum algorithms;
  SS & S Herbert. arXiv:1910.02861.

[3] Merged submission to QIP 2024; standalone preprint — arXiv:2307.02389; C Ikenmeyer and *SS*.

- Do black holes store negative entropy?;
  K Azuma & *SS*. arXiv:1807.06753
  Accepted for publication in PTEP.

- Stable non-spherical configurations of rotating magnetized white dwarfs;
  SS & B Mukhopadhyay.
  Monthly Notices of the Royal Astronomical Society 454(**1**), pp. 752-765 (Sept. 2015).

MISCELLANEOUS

Programming
- C/C++, Shell, GNU Octave, LaTeX, Python, CUDA, MPI
- Experience implementing machine learning algorithms (e.g. LSTMs for NLP)
- Linux (Arch), Mathematica, git (`sathynius2@github`)

Languages
- *Native/Bilingual*: Malayalam, Tamil; *Fluent*: English, Hindi, Sanskrit
- *Intermediate*: German, Italian, Japanese

Other
- *Workshop on Algebraic Complexity Theory* (WACT'23, Warwick)
- *Quantum Techniques in Machine Learning* Workshop (QTML 2017, Verona)
- *Code for Good* (London) - Hackathon at JP Morgan (Dec. 2015)
- *Radio Astronomy Winter School* (RAWS 2013, NCRA Pune)
- *CUDA Programming and GPU* Workshop (Jan. 2013, SERC, IISc. Bangalore)
- Trained in Indian classical music (carnatic)—violin—for 7 years

PROFESSIONAL REFERENCES

(available on request)

Richard Jozsa, FRS
Emeritus Leigh Trapnell Professor
of Quantum Physics
University of Cambridge, UK

Tom Gur
Associate Professor
Department of Computer Science
University of Cambridge, UK

Min-Hsiu Hsieh
Director
Foxconn Quantum Computing Center, Taiwan

Koji Azuma
Distinguished Researcher
NTT Basic Research Laboratories, Japan

Christian Ikenmeyer
Professor
Departments of Computer Science and Mathematics
University of Warwick, UK

Alexander Golovnev
Assistant Professor
Deparment of Computer Science
Georgetown University, USA

Sergii Strelchuk
Associate Professor
Department of Computer Science
University of Oxford, UK

Iordanis Kerenidis
Co-Founder and CTO of Quantum Signals
Research Director CNRS, France

SATHYAWAGEESWAR SUBRAMANIAN

+44 (0) 74887 02818 • ss2310@cam.ac.uk

My research has focused broadly on three main intertwined strands of inquiry into the power and utility of quantum computational models: (1) algorithms and complexity [query, sample, and circuit complexity] (2) learning theory (3) property testing. In addition, I have worked on quantum information (black hole entropy) and stellar structure (general relativity and magnetohydrodynamics) in the past, showing versatility, and demonstrating that I see myself primarily as a mathematician keen to work on a variety of interesting problems from Physics and Computer Science.

PUBLICATIONS

+full list on arXiv:subramanian_s & G

1. Quantum Channel Testing in Average-Case Distance;
   G Rosenthal, H Aaronson, *S Subramanian*, A Datta, T Gur. arXiv:2409.12566

2. Quantum Circuits surpass Biased Threshold Circuits in Constant-Depth;
   MH Hsieh, L Mendes, M Oliveira, *S Subramanian*.
   19th Theory of Quantum Computation, Communication and Cryptography – **TQC 2024**
   21th International Conference on Quantum Physics and Logic – **QPL 2024**

3. Information-theoretic generalization bounds for learning from quantum data;
   M Caro, T Gur, C Rouzé, D Stilck França, *S Subramanian*.
   19th Theory of Quantum Computation, Communication and Cryptography – **TQC 2024**
   37th Annual Conference on Learning Theory – **COLT 2024** (Edmonton)

4. Quantum complexity of the Kronecker coefficients[1];
   S Bravyi, A Chowdhury, D Gosset, V Havlicek, C Ikenmeyer, *S Subramanian*, G Zhu.
   27th Conference on Quantum Information Processing – **QIP 2024** (Taipei)
   [[We apply insights from representation theory to show that the Kronecker coefficients, which arise as the multiplicities of of irreps in the decomposition of tensor product of two irreps of the permutation group, can be related to the dimension of the subspace of states that makes the verifier in a QMA protocol accept. Hence we obtain a QMA protocol to decide if a Kronecker coefficient is positive. Our work refines prior work of Bravyi et al. to show that kronecker is in #BQP (it is a long standing open problem in combinatorics whether Kronecker is in #P). This has applications to quantum approximate counting, and Kronecker coefficients have applications in areas such as the quantum marginal problem for distinguishable particles. We obtain similar results for plethysm coefficients (which have applications to the quantum marginal problem for indistinguishable particles).]]

5. Quantum Ridgelet Transform: Winning the Neural Network Lottery Ticket;
   H Yamasaki, *S Subramanian*, S Hayakawa, S Sonoda.
   40th International Conference on Machine Learning – **ICML 2023** + IBIS 2022.

6. Quantum Worst-Case to Average-Case Reductions for All Linear Problems;
   VR Asadi, A Golovnev, T Gur, I Shinkar, *S Subramanian*.
   26th Conference on Quantum Information Processing – **QIP 2023** (Ghent)
   35th ACM-SIAM Symposium on Discrete Algorithms – **SODA 2024**

7. Testing graph states with constant time and sample complexity;
   H Yamasaki & *S Subramanian*. arXiv:2201.11127.
   [[We bring insights from property testing and sublinear algorithms to the question of checking whether large, highly entangled resource states (called cluster states or graph states) used for measurement based quantum computation (MBQC, implemented for example on photonic platforms) are good enough (in the sense of having low enough physical error across lattice sites) to

[1] Merged submission to QIP 2024; standalone preprint — arXiv:2307.02389; C Ikenmeyer and *S Subramanian*.

be usable for fault-tolerant MBQC. Our results are novel and elegantly simple, achieving constant sample and time complexity for a task that has practical value in benchmarking MBQC platforms.]]

8. Sublinear quantum algorithms for estimating von Neumann entropy;
T Gur, MH Hsieh, *S Subramanian*. arXiv:2111.11139.
25th Conference on Quantum Information Processing – **QIP 2022** (Caltech).

9. Quantum algorithm for estimating α-Renyi entropies of quantum states;
S Subramanian & MH Hsieh. *Physical Review A* **104**, 022428 (Aug. 2021).
[[The key message of this paper is threefold - (1) that quantum matrix functions implementation techniques that I had worked on (item 12 below) could be used for property estimation of quantum states; (2) that query complexity in the standard circuit model can be traded off for a slightly (quadratically in epsilon) larger number of measurements on a weaker model (namely DQC1, using 1 clean qubit and n 'dirty' qubits in a maximally mixed state); and (3) additive estimates can be iteratively refined to obtain multiplicative estimates for the Renyi entropies of unknown quantum states. This work has since spurred work on entropy estimation in different quantum input models, and applications in a range of quantum parameter estimation tasks.]]

10. A quantum search decoder for Natural Language Processing;
J Bausch, *S Subramanian*, S Piddock. *Quantum Machine Intelligence* **3**, 16 (Apr. 2021).
[[We adapt and apply a result of Montanaro which shows super-quadratic (and even exponential) speedups in expected/average query complexity of searching when a prior distribution is available (in the form of a non-uniform initial state) to the case of the beamsearch subroutine used to decode the output of natural language processing models. This application is practically valuable due to the empirically well-established fact that inputs in NLP follow power-law distributions, which are useful priors for which super-polynomial speedups in expected complexity are obtained by quantum search.]]

11. Learning with optimized random features: Exponential Speedup by Quantum ML;
H Yamasaki, *S Subramanian*, S Sonoda, M Koashi.
Advances in Neural Information Processing Systems 33 – **NeurIPS (2020)**.
[[I jointly conceived and developed this project with Prof. Masato Koashi's PhD student and subsequent postdoc Dr. Yamasaki when I was a PhD student. At the time both of us were visiting Prof. Min-Hsiu Hsieh at UTS, Sydney. We gave a quantum algorithm with exponential speedups over best known classical algorithms for a class of machine learning tasks that can be solved by sampling suitably weighted random Fourier features of the data; essentially an efficient quantum algorithm for the optimised random feature sampling method presented in Francis Bach (JMLR 2017). The novelty of our algorithm includes the ability to deal with dense and full-rank inputs, a bottleneck that previous algorithms were unable to handle. We overcame this bottleneck in our setting by a clever use of the quantum Fourier transform (qFT) to diagonalise an exponentially large discretised integral operator that could then be inverted by HHL before using the QFT to switch back to the computational basis. The HHL algorithm is a naturally suited for this problem since the end goal is to *sample* Fourier features.]]

12. Implementing smooth functions of a Hermitian matrix on a quantum computer;
S Subramanian, S Brierley, R Jozsa.
Journal of Physics Communications, 3(**6**), 65002 (Oct. 2019).
[[We generalised the work of Childs, Kothari, Somma on matrix inversion using the linear combination of unitaries method to the case of arbitrary (smooth) functions of input matrices, using insights about Chebyshev polynomial approximations. This was contemporaneous with and contributed to the line of work that culminated in the quantum singular value transformation.]]

13. Spectral sparsification of matrix inputs as a preprocessing step for quantum algorithms;
S Subramanian & S Herbert. arXiv:1910.02861.

14. Do black holes store negative entropy?;
K Azuma & *S Subramanian*. arXiv:1807.06753.

15. Stable non-spherical configurations of rotating magnetized white dwarfs;

S Subramanian & B Mukhopadhyay.
Monthly Notices of the Royal Astronomical Society 454(**1**), pp. 752-765 (Sept. 2015).

HIGHLIGHTS

Out of the works listed above, I would like to highlight the following five (listed one per page, starting from the next page). In all cases, I initiated and led the project, carrying out its execution through research, writing, and publication of the manuscript with the support and collaboration of my coauthors. Items 1,2,3,5 follow the standard practice in Computer Science of alphabetically ordered author list. I am the sole lead author of 5, the joint lead in 2 and 4, and the joint lead (with the PhD students whom I co-supervised during the respective projects) and corresponding author of 1 and 3. Authors marked by "*" made equal contribution to research, writing of manuscript, and publication.

1. Quantum Circuits surpass Biased Threshold Circuits in Constant-Depth;
MH Hsieh, L Mendes, M Oliveira*, *S Subramanian**.
19th Theory of Quantum Computation, Communication and Cryptography – **TQC 2024**
21th International Conference on Quantum Physics and Logic – **QPL 2024**

Highlights. Landmark quantum algorithms such as integer factorisation or search are beyond the constraints of existing quantum hardware, requiring extensive resources including qubit counts, coherence times and fault-tolerant operations for their implementation. Motivated by this, recent work has drawn attention to the comparison between shallow-depth quantum circuits their classical shallow-depth counterparts, which have been studied extensively in classical computer science, and continue to be a vibrant area of research at the frontier of complexity theory.

In this work that I initiated and led, we advance the state-of-the-art in unconditional, provable separations between the computational power of shallow-depth quantum and classical circuits. We study classes of constant-depth circuits with gates that compute restricted polynomial threshold functions, recently introduced (CCC 2023) as a family that strictly generalizes AC^0 and pushes towards TC^0. Denoting these circuit families $\mathsf{bPTF}^0[k]$ for *biased polynomial threshold circuits* parameterized by an integer-valued bias k, we prove three hardness results separating these classes from constant-depth quantum circuits (QNC^0): (1) We show that QNC^0 over qubits can solve a relational problem that polynomial size $\mathsf{bPTF}^0[k]$ circuits fail to solve with non-negligible advantage even on average; (2) by constructing a novel family of non-local games for p-dimensional systems for primes larger than 2, we extend our first result to separations between QNC^0 over p-dimensional quantum systems and polynomial size $\mathsf{bPTF}^0[k]$ circuits; (3) We show that both of these unconditional classical-quantum separations in computational power are robust to noise—i.e. even noisy QNC^0 circuits retain an advantage over noiseless $\mathsf{bPTF}^0[k]$ circuits.

In physical experiments that test computational separations, we need to pin down at what values of circuit depth d and input size n (i.e. number of input qubits) we observe a transition in the circuit size. That is, at what depths and input sizes do the quantum advantages kick in? We also prove tighter lower bounds on the size of $\mathsf{bPTF}^0[k]$ circuits that are required to solve the relational problem with *certainty*, which we leverage to significantly reduce the estimated resource requirements for potential demonstrations of quantum advantage against shallow-depth classical circuits.

$\mathsf{bPTF}^0[k]$ circuits can compute certain classes of Polynomial Threshold Functions (PTFs), which in turn serve as a natural model for neural networks and exhibit enhanced expressivity and computational capabilities. We also give an extensive discussion of recent advances in analysing transformers, a crucial neural network architecture underlying large language models, in terms of circuit complexity, and point out new connections to biased PTFs.

The main challenges we overcome include establishing subexponential classical average-case lower bounds, designing non-local games in all prime dimensions with quantum-classical gaps in winning probabilities, and developing noise-resilient non-Clifford quantum circuits necessary to extend beyond qubits to higher dimensions. We address the first challenge by developing new, tighter multi-switching lemmas for multi-output $\mathsf{bPTF}^0[k]$ circuits. For the second, we analyze a new family of non-local games defined in terms of $\mathsf{mod}\ p$ computations, characterized by an exponential difference between their classical and quantum success probabilities; prior to our work, non-local games in higher dimensions had received little attention. Lastly, we extend constant-depth error correction techniques from qubits to qupits, incorporating logical advice states and magic state-injection, and using the hard decision renormalisation decoder to enable encoding and decoding to be performed in constant-depth. The technical tools we develop in this work may be of wider independent interest.

2. Information-theoretic generalization bounds for learning from quantum data;
M Caro*, T Gur, C Rouzé, D Stilck França, *S Subramanian**.
19th Theory of Quantum Computation, Communication and Cryptography – **TQC 2024**
37th Annual Conference on Learning Theory – **COLT 2024** (Edmonton)

Highlights. Learning theory, the mathematical foundation underlying machine learning, has traditionally been dominated by dimension-based bounds that only depend on the worst-case behaviour of concept and hypotheses classes, and do not take into account the training data or the learning algorithm itself. This has been addressed in the last decade by a growing body of work on information-theoretic and data-dependent ways of analysing the generalisation error of learners, including methods based on differential privacy and conditional mutual information. Quantum learning theory has evolved rapidly in the last decade, with multiple directions developing separately. Learning tasks in quantum information and computation range from fundamental problems such as state discrimination, to the recently proposed shadow variants of state tomography.

I initiated and led this work jointly with Dr. Caro, with the intention of conceptualising a unifying theory of generalisation for quantum learning algorithms. We construct a mathematical framework that captures the most general quantum learning algorithms in terms of training on classical-quantum data and then testing how well the learned hypothesis generalizes to new data. Towards this end, we make novel definitions of quantum learners, loss observables, generalization error and change-of-measure arguments. In this framework, we can prove bounds on the expected generalization error of a quantum learner in terms of classical and quantum information-theoretic quantities that measure how strongly the learner's hypothesis depends on the data seen during training. We also demonstrate applications of our results to recover and even improve previously known bounds on PAC learning quantum states, quantum parameter estimation, and other problems. We draw on tools from quantum optimal transport and quantum concentration inequalities to establish non-commutative, quantum versions of decoupling lemmas that underlie recent information-theoretic generalization bounds in classical learning theory.

Our framework encompasses and gives intuitive generalization bounds for a variety of quantum learning scenarios such as quantum state discrimination, PAC learning quantum states, quantum parameter estimation, and quantumly PAC learning classical functions, laying a foundation for a unifying quantum information-theoretic perspective on quantum learning.

Invited talks: DPG Spring Meeting (Berlin), Quantinuum (Cambridge), QuSoft and CWI (Amsterdam).

3. Quantum Worst-Case to Average-Case Reductions for All Linear Problems;
VR Asadi*, A Golovnev, T Gur, I Shinkar, *S Subramanian**.
26th *Conference on Quantum Information Processing* – **QIP 2023**[2] (Ghent)
35th *ACM-SIAM Symposium on Discrete Algorithms* – **SODA 2024**

Highlights. Average-case algorithms play a central role in complexity theory. Indeed, even if a problem seems to be hard to solve efficiently due to pathological worst-case instances, we regularly have heuristics that work exceedingly well on the types of instances that appear often in real-life. On the flipside, if a problem is provably hard even on average, it finds use in cryptography.

Given an algorithm that has a small non-zero probability of answering correctly on an average or typical input, can we use it to design another algorithm that has non-zero probability of answering correctly even on worst-case inputs? In this work that I led, we were able to answer this question for quantum algorithms, and describe an explicit and efficient transformation that turns algorithms which are only correct on a small (even sub-constant) fraction of their inputs into ones that are correct on *all* inputs, for a large class of problems fundamental to all branches of science and engineering, known as linear problems. Intuitively, a linear problem takes as input a vector, and the solver must produce as output another vector that is obtained from the input by performing a matrix multiplication. This mathematical formalism is fundamental to all branches of science and engineering, and captures operations such as the discrete Fourier transform, rotations in 3D, and evaluating polynomials at a set of input points. Such fine-grained worst-case to average-case reductions for this general class of problems are still open for classical algorithms.

To do this, we introduce and develop a novel connection between additive combinatorics that has traditionally been studied in pure mathematics from the perspective of number theory, and quantum algorithms and complexity theory. This requires new ideas to overcome challenges on three fronts: (1) additive combinatorics techniques, (2) new quantum algorithms, and (3) complexity theoretic reduction techniques. We prove a robust, probabilistic version of Bogolyubov's lemma, and construct quantum algorithms for matrix-vector product verficiation and subspace learning that overcome classical bottlenecks for these tasks. Along the way, we also generalise Adcock and Cleve's (2001) fundamental quantum extension of the Fourier analytic Goldreich-Levin algorithm.

Our results offer an intriguing contrast to the classical setting, where similar results are only known for a small number of specific problems or restricted computational models, motivating questions of quantum-classical separations in average-case complexity. Thus our work solves the problem in the more complicated setting of quantum computing, while opening up a plethora of interesting new avenues of research.

Invited talks: Microsoft Research (India), the Indian Institute of Science (IISc Bangalore), Cambridge (CQIF, DAMTP), National University of Singapore (Centre for Quantum Technology).

[[Aspects of this work are related to the theme of average-case complexity that I propose to pursue during the URF, and demonstrate my prior expertise in tackling problems of a related flavour.]]

[2] The Annual Conference on Quantum Information Processing (QIP) is the largest and premier conference in the area of quantum computing and information theory; it was started in 1998 and is the most selective and competitive venue for presenting research, receiving nearly 700 talk submissions every year, with an acceptance rate of about 1 in 7.

4. Quantum Ridgelet Transform: Winning the Neural Network Lottery Ticket;
H Yamasaki*, *S Subramanian** , S Hayakawa, S Sonoda.
40^{th} International Conference on Machine Learning – **ICML 2023**.

Highlights. Ridgelet analysis has become a staple and important ingredient in classical signal processing and statistical learning theory, revolutionising image processing and computer vision. Wavelets and ridgelets formed a part of Yves Meyer's 2017 Abel Prize winning work. Its practical implementability in learning tasks has however been limited, since the best known numerical implementations by classical computation incur an exponential runtime overhead in the data dimension. In this paper, my coauthors and I introduced a novel discretised unitary version of the Ridgelet Transform and build an efficient quantum circuit implementing it. We showed how quantum computers can implement this transformation on the amplitudes of quantum states, with exponentially smaller circuit size than the best known classical algorithm, using deep relations to the quantum Fourier transform that we prove using functional analytic techniques.

We also demonstrate an application of our quantum ridgelet transform to explicitly give an algorithm that finds a sparse trainable sub-network in a dense, shallow neural network that generalises to the same accuracy as the original network. The phenomenon of the existence of a sparse subnetwork that can be trained to achieve the same accuracy as the original large and dense neural network has been dubbed the lottery ticket hypothesis, and has received intense attention in the last two years for its practical value in pruning ML implementations such as large language models. This allows us to explicitly compute neural network architectures that can be less than 10-20% of the size of several fully-connected and convolutional feed-forward architectures, which is considered of significant value in ML for reducing parameter counts, decreasing storage requirements and improving computational performance of inference without compromising accuracy. Our results have exciting applications to speeding up classical ML using quantum techniques, a topic that has received much attention in the last few years, due to genuinely quantum ML algorithms being resource-heavy and inoperable on noisy intermediate scale quantum devices of the near future.

[[This work takes inspiration from my own previous work on random Fourier features, that appeared in NeurIPS 2020 (item 10 in the list above)]]

Learning with optimized random features: Exponential Speedup by Quantum ML;
H Yamasaki*, *S Subramanian** , S Sonoda, M Koashi.
Advances in Neural Information Processing Systems 33 – **NeurIPS (2020)**.

[[I jointly conceived and developed this project with Dr. Yamasaki when we were both PhD students, and we succeeded in constructing a quantum algorithm with exponential speedups over the best known classical algorithms for a class of machine learning tasks. The novelty of our algorithm includes the ability to deal with dense and full-rank inputs, a serious bottleneck that previous algorithms were unable to handle. We overcame this bottleneck in our setting by a clever use of the quantum Fourier transform (qFT) to diagonalise an exponentially large discretised integral operator that could then be inverted before using the qFT to switch back to the computational basis.]]

5. Sublinear quantum algorithms for estimating von Neumann entropy;
T Gur, MH Hsieh, *S Subramanian*. arXiv:2111.11139.
25th Conference on Quantum Information Processing – **QIP 2022** (Caltech).
To appear in Quantum (2024)

Highlights. The entropy is a global property of both classical and quantum systems, and is a central quantity in myriad subjects ranging from physics to computer science. In this work, I addressed a two decade old open problem, resolving it positively by showing that it is possible to estimate the entropy of quantum states with an amount of resources that grows significantly slower than their dimension. More precisely, we constructed the first quantum algorithm that estimates the von Neumann entropy to within a multiplicative precision using a sublinear number of probes to the unknown state. This is of particular significance in the light of he exponential growth of the dimension of quantum systems as a function of the number of particles. Our algorithm combines unconventional function approximation techniques to get a multiplicative estimator for the entropy with powerful state-of-the-art quantum primitives,, such as the quantum singular value transformation toolkit, for implementing functions of input matrices.

In addition, we gave a quantum algorithm to estimate the entropy of classical random variables that provably beats the best conceivable classical algorithm in sample complexity. We also established the first non-trivial lower bounds on the query complexity of quantum algorithms in the very general purified query access input model that we studied in this work. This input model can capture the realistic scenario wherein one quantum algorithm's output state could be the input to another quantum algorithm, and is believed to be very powerful in admitting highly efficient property testing algorithms.

[[This paper continues a line of work that focuses on property estimation and testing algorithms started in my previous article on estimating Renyi entropies of quantum states that appeared in Physical Review A (item 8 in the above list). This in turn was a project which I independently proposed while I was a PhD student, and developed with the guidance of Prof. Hsieh who had at the time invited me to visit him at UTS Sydney for 2 months.]]