Research Statement

Wei Zhan

My research centered around **computational complexity**, with a focus on **space-bounded** models of both **classical** and **quantum** computation. These models put constraints on the size of the accessible memories during the computation, usually in addition to the time constraints. This additional dimension of complexity measure opens up new possibilities for studying the power of such models, which could lead to results that are deemed improbable or out of reach for their time-bounded counterparts.

Besides theoretical interest, spaced-bounded models are traditionally studied with real-life applications such as streaming and delegating computation, where randomness is often the key ingredient. Nowadays, in the era of NISQ that near-term quantum devices are heavily limited on the number of controllable qubits, understanding the power of space-bounded quantum computation also shows practical importance. As a result, my research is mainly motivated by the following fundamental and largely open question: **What is the power of randomness and quantum mechanics in space-bounded computation?**

The first major part of my research [4, 7, 9, 11] tries to answer the question by studying the corresponding space-bounded complexity classes. I gave new characterizations to these classes, which provide alternative viewpoints besides their natural definitions. These allow us to prove useful properties that are otherwise not evident, helping towards the final goal of proving the classes separate or collapse.

The second major part of my research [1, 3, 8, 12] tries to answer the question by proving lower bounds and separations, especially the ones that could demonstrate quantum advantage. All the super-polynomial quantum advantage thus far has been proved either conditionally (assuming a certain problem is hard for P), or relative to an oracle (assuming a large black-box object). The best result to date is by Zhandry and Yamakawa [29], which is still relative to a random oracle. Fundamentally, this is because of the lack of strong unconditional classical lower bounds, and I believe adding space constraints makes a difference.

In the rest of this statement, I will explain these two complementary parts of my research in more details.

Characterizing Logspace Complexity Classes

Since the work of Savitch [14], it becomes clear that complexity classes defined with logarithmic space, whether nondeterministic (NL), random (RL and BPL) or quantum (BQL), all contains L (deterministic logspace) and are all contained in $\mathsf{L}^2 = \mathsf{DSPACE}(\log^2 n)$ (deterministic space $O(\log^2 n)$), as they all reduce to the problem of computing matrix determinant. But where do they lie exactly in the space hierarchy? My work tries to further our understanding in this question. For randomized computation, it is currently known that $\mathsf{BPL} \subseteq \mathsf{DSPACE}(\log^{1.5-o(1)} n)$ [24] and widely believed that actually $\mathsf{BPL} = \mathsf{L}$. For quantum computation, the study of the class BQL began much more recently and was only defined consistently through my work and others.

Unitary Quantum Computation

Previously, quantum computation models are often defined with only unitary operators in mind, instead of considering all physically possible quantum channels. For time-bounded classes, this is not a problem because of Stinespring dilation (that every quantum channel is just the effect of a unitary operator in a larger system). However with space constraints, this becomes an issue. For instance, it was not even known how to simulate BPL within unitary logspace because randomness is irreversible by nature. Another example is intermediate measurements, which is usually deferred to the end of the computation by adding ancilla qubits, resulting in a unitary computation with much higher space usage.

In my work with Uma Girish and Ran Raz [4], we resolved this discrepancy by showing that unital quantum channels, including the bit-flip channel and measurement channels, can be space-efficiently simulated with unitary computation. As a result, the class $\mathsf{BQ_UL}$ of unitary quantum logspace computation contains BPL, and can naturally perform error reduction. Independently with a different approach, Fefferman and Remscrim [23] proved that the result extends to all quantum channels, and thus $\mathsf{BQ_UL} = \mathsf{BQL}$. Combining with our techniques, we can show that they are equal in a strong sense: The output distribution of any BQL computation, not limited to decision problems, can be simulated unitarily with polynomial accuracy. Subsequently, our simulations are proved to be optimal by Zhandry [30].

Verifying Computation and Derandomization

In another project with Uma and Ran [7, 11], we considered the problem of verifying quantum computation: How could a quantum computer prove to classical devices efficiently that the computation is correct? For polynomial-time quantum computation, this task is highly involved and currently known protocols require multiple non-communicating provers or assumptions on hard problems. We found out that for logspace quantum computation, there exists a very simple protocol with no interaction needed, and many other desired properties:

- The proof is the averaged computation history, represented as a stream to the verifier;
- The logspace verifier can check the proof using only logarithmic many random coins.

This protocol is interesting even for verifying classical randomized computation, as it allows characterizing BPL with *trusted randomness*. Imagine the random bits provided to a probabilistic logspace machine are arbitrarily corrupted. But as long as there is access to $O(\log n)$ trustworthy perfect random bits, by verifying its own computation history the machine could still perform the computation correctly, or realize that the rest random bits are rigged and abort immediately.

In particular, this means that we can safely use any source of pseudorandomness, whether theoretically proved or not, to derandomize BPL computation without worrying about the answer being wrong. With Ted Pyne and Ran Raz [9] we worked on this premise and obtained numerous applications for derandomization. One of the most interesting ones is that, we explicitly wrote out a deterministic program that universally and optimally derandomizes BPL: If $\mathsf{BPL} \subseteq \mathsf{DSPACE}(\log^c n)$, then the program solves every problem in BPL within space $O(\log^c n)$. So if you believe in $\mathsf{BPL} = \mathsf{L}$, our program already accomplishes that!

Unconditional Quantum Advantage and Time-Space Lower Bounds

There has been extensive research in the models that concern query complexity and communication complexity, where polynomial separations can be shown between quantum and classical computation, or even stronger separations for partial functions. My work in [3] proved that such separations persist even if the classical computation just want to be a little bit better (inverse quasi-polynomial, to be exact) than random guessing. A line of works that I was involved in [2, 5, 6] studied parallel repetition of multi-player nonlocal games, which also can be used to boost the separation of success probabilities between players sharing quantum entanglements and classical players.

But ultimately, separations like these do not represent convincing quantum advantage in real life: The classical lower bounds are information-theoretical and could not exceed the size of the problem. Therefore, without having an exponential-sized object as a mysterious oracle to "lift" the lower bounds, these problems are still "easy" for classical computation. On the other hand, unconditional classical lower bounds of high magnitudes for problems in BQP are extremely hard to prove. For instance, the best lower bounds against boolean circuits and random-access machines for explicit decision problems, after decades of study, are still linear. The goal of my work in this direction is to tackle the dilemma by adding space constraints to the picture.

The Almost-Linear Barrier

There are in general two ways of proving time-space lower bounds for explicit decision problems. One is the approach initiated by Fortnow [18], based on diagonalization and time-hierarchy theorems. It is capable of proving time lower bounds of form $\Omega(n^{1+c})$ for sub-polynomial space, but is not designed to work on BQP problems. The other one, which is more versatile, is to directly analyze the *branching program*, a non-uniform model that captures space-bounded computation. Sadly, the best lower bound proved via this approach is $\Omega(n \log^2(n/S))$ for space S, which is almost linear, by Babai, Nisan and Szedgedy more than 30 years ago [17].

In my work with Ran Raz [1], we attempted to break this almost-linear barrier by introducing a new computation model called *random-query*. In this model the algorithm could access a uniformly random index of the input in each time step, instead of querying a specific index. When there are certain dependencies between the indices received in different steps, any time-space lower bound in this model could be translated to the same lower bound against branching programs, thus could potentially break the barrier. To initiate the study, we proved that when the indices are all independent, it must take $\Omega(n^2/\max\{S, \log n\})$ time to compute functions such as majority and XOR with zero-error. The follow-up work by Dinur [26] extended our lower bound into the bounded-error scenario.

My work with Huacheng Yu [12] looked at the problem from a different angle. Multi-output functions, where the output size is polynomial in the input size, is not subject to the above-mentioned barrier. In fact, time-space lower bounds that are polynomially better than linear have been successfully proved for various multi-output functions by applying a framework called the *Borodin-Cook method* [15]. These lower bounds, proved for randomized computation, are often not known to be tight against deterministic algorithms. We made a surprising connection in

[12]: For a number of natural multi-output functions such as collision finding, a polynomially better lower bound that beats the Borodin-Cook method would break the almost-linear barrier for decision problems. Additionally, we designed an artificial problem where we indeed beat the Borodin-Cook method, giving an evidence that the barrier may not be that impenetrable.

Learning and Exponential Lower Bounds

Even if we break the almost-linear barrier as planned, we are still expected to prove only polynomial lower bounds. The quantum advantage that we are craving requires super-polynomial lower bounds for classical problems, which seems strictly out of reach.

But the result by Raz [21] leaves us hope. It shows that unconditional exponential time lower bound with space restriction is possible for *learning* problems, where the hypotheses and samples are of polynomial length. So my goal becomes clear: to find a learning problem such that,

- A quantum computer that receives classical samples can solve it in polynomial time and some space S;
- Any classical learning algorithm with space $O(S)$ must take exponentially many samples.

We also need S to be reasonably large so that the classical lower bound is relevant in practice, as otherwise problems like the coin problem [22] would suffice. Specifically, we want S to be super-logarithmic: In my paper [4] mentioned before, we showed that when S is logarithmic, the existence of such a learning problem is actually equivalent to $\mathsf{BQL} \neq \mathsf{BPL}$, so it is not even easier than focusing on decision problems.

What about learning a parity function on an unknown subset of variables, the original problem studied in [21] where the classical lower bound holds for $S = o(n^2)$? It turns out that quantum computing is not too magical for this problem. In my work with Qipeng Liu and Ran Raz [8], we proved that a quantum computer, with a small linear quantum memory, still requires either quadratic-sized classical memory or exponentially many samples to learn parity. It leaves the possibility that a quantum memory of size, say $S = O(n \log n)$, could do all the magic and allow better algorithms than Gaussian elimination, but we conjecture that it is not the case. As such, the search for a learning problem that demonstrates quantum advantage continues.

Future Directions

I left several loose ends in previous sections, which I will be continuously thinking about: derandomizing BPL, breaking the almost-linear barrier of decision problems, and demonstrating quantum advantage via learning problems. Besides these, here are some of the related directions I am actively working on, or would like to think more about in the future:

Random Quantum Circuits and Pseudorandomness

Currently the most popular paradigm of showing conditional quantum advantage is *random circuit sampling*, because of the theoretical evidences that simulating the output distribution of a

random quantum circuit is hard classically. In fact, it is often conjectured that random quantum circuits of super-logarithmic depth are pseudorandom unitaries (defined in [20]), that their complexity is high enough to be indistinguishable from a Haar random unitary over the entire system.

My recent work with Bill Fefferman and Soumik Goush [10] started as an exploration of the conditions that a random quantum circuit distribution is pseudorandom. A pseudorandom circuit cannot be efficiently learned, and we conjectured that the reverse is also true. While trying to connect the dots, I accidentally found the first proper learning algorithm for logarithmic-depth random quantum circuits. Although we did not prove our original conjecture, the result serves as an strong evidence towards it. I would like to work more on conjecture, as very few constructions of pseudorandom unitaries have been proposed, and our general understanding of the object is still lacking.

Quantum Data Structures

Data structure complexity is another way of viewing space-bounded computation. It models data structures as fixed-sized bit strings with random access, and ask how many bits (or groups of bits as words) need to be read in individual updates or queries. Classical data structure complexity has been systematically studied, and I also contributed through my work with Huacheng Yu [13], where we proved data structure lower bounds using communication complexity of sampling problems.

In comparison, the emerging research of data structures on quantum computers is quite ad hoc. Interestingly, recent works on Fermion data structures [25, 28] suggest that we can actually extract classical data structure problems from the quantum settings. I plan to work on such problems in the near future, and more generally use the ideas and techniques from classical data structure complexity to prove meaningful results on quantum data structures.

Space-Bounded Quantum Interactive Proofs

Our work in [7, 11] can be viewed as a streaming Merlin-Arthur proof system. When we consider the more general interactive proof systems with logspace verifier, their computational powers were initially explored by Condon and Ladner [16], and only fully understood after the work of Goldwasser, Kalai and Rothblum [19].

What if we consider interactive proofs with quantum provers and logspace quantum verifiers? This model was recently considered in [27], and it was shown that whether verifier is unitary or not makes a difference, but their exact powers are not pinned down. With the help of my works on BQL, I plan to work towards a full characterization of space-bounded quantum interactive proofs. For example, the techniques in [4] can be used to convert the verifier in the GKR protocol into a unitary one. Yet, more works need to be done to make the protocol sound against quantum proofs.

References

[1] Ran Raz and **Wei Zhan**. "The Random-Query Model and the Memory-Bounded Coupon Collector". In: *11th Innovations in Theoretical Computer Science Conference, ITCS 2020*.

[2] Uma Girish, Justin Holmgren, Kunal Mittal, Ran Raz, and **Wei Zhan**. "Parallel Repetition for the GHZ Game: A Simpler Proof". In: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2021*.

[3] Uma Girish, Ran Raz, and **Wei Zhan**. "Lower Bounds for XOR of Forrelations". In: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2021*.

[4] Uma Girish, Ran Raz, and **Wei Zhan**. "Quantum Logspace Algorithm for Powering Matrices with Bounded Norm". In: *48th International Colloquium on Automata, Languages, and Programming, ICALP 2021*.

[5] Uma Girish, Justin Holmgren, Kunal Mittal, Ran Raz, and **Wei Zhan**. "Parallel repetition for all 3-player games over binary alphabet". In: *54th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2022*.

[6] Uma Girish, Kunal Mittal, Ran Raz, and **Wei Zhan**. "Polynomial Bounds on Parallel Repetition for All 3-Player Games with Binary Inputs". In: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2022*.

[7] Uma Girish, Ran Raz, and **Wei Zhan**. "Is Untrusted Randomness Helpful?" In: *14th Innovations in Theoretical Computer Science Conference, ITCS 2023*.

[8] Qipeng Liu, Ran Raz, and **Wei Zhan**. "Memory-Sample Lower Bounds for Learning with Classical-Quantum Hybrid Memory". In: *55th Annual ACM Symposium on Theory of Computing, STOC 2023*.

[9] Edward Pyne, Ran Raz, and **Wei Zhan**. "Certified Hardness vs. Randomness for Log-Space". In: *64th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2023*.

[10] Bill Fefferman, Soumik Ghosh, and **Wei Zhan**. "Anti-Concentration for the Unitary Haar Measure and Applications to Random Quantum Circuits". In: *arXiv preprint arXiv:2407.19561* (2024).

[11] Uma Girish, Ran Raz, and **Wei Zhan**. "Quantum Logspace Computations are Verifiable". In: *2024 Symposium on Simplicity in Algorithms, SOSA 2024*.

[12] Huacheng Yu and **Wei Zhan**. "Randomized vs. Deterministic Separation in Time-Space Tradeoffs of Multi-Output Functions". In: *15th Innovations in Theoretical Computer Science Conference, ITCS 2024*.

[13] Huacheng Yu and **Wei Zhan**. "Sampling, Flowers and Communication". In: *15th Innovations in Theoretical Computer Science Conference, ITCS 2024*.

[14] Walter J. Savitch. "Relationships Between Nondeterministic and Deterministic Tape Complexities". In: *J. Comput. Syst. Sci.* 4.2 (1970).

[15] Allan Borodin and Stephen A. Cook. "A Time-Space Tradeoff for Sorting on a General Sequential Model of Computation". In: *SIAM J. Comput.* 11.2 (1982).

[16] Anne Condon and Richard E. Ladner. "Probabilistic Game Automata". In: *J. Comput. Syst. Sci.* 36.3 (1988).

[17] László Babai, Noam Nisan, and Mario Szegedy. "Multiparty Protocols, Pseudorandom Generators for Logspace, and Time-Space Trade-Offs". In: *J. Comput. Syst. Sci.* 45.2 (1992).

[18] Lance Fortnow. "Time-Space Tradeoffs for Satisfiability". In: *J. Comput. Syst. Sci.* 60.2 (2000).

[19] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. "Delegating Computation: Interactive Proofs for Muggles". In: *J. ACM* 62.4 (2015).

[20] Zhengfeng Ji, Yi-Kai Liu, and Fang Song. "Pseudorandom Quantum States". In: *38th Annual International Cryptology Conference, CRYPTO 2018*.

[21] Ran Raz. "Fast Learning Requires Good Memory: A Time-Space Lower Bound for Parity Learning". In: *J. ACM* 66.1 (2019).

[22] Mark Braverman, Sumegha Garg, and Or Zamir. "Tight Space Complexity of the Coin Problem". In: *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021*.

[23] Bill Fefferman and Zachary Remscrim. "Eliminating intermediate measurements in space-bounded Quantum computation". In: *53rd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2021*.

[24] William M. Hoza. "Better Pseudodistributions and Derandomization for Space-Bounded Computation". In: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2021*.

[25] Joseph Carolan and Luke Schaeffer. "Succinct Fermion Data Structures". In: *arXiv preprint arXiv:2410.04015* (2024).

[26] Itai Dinur. "Time-Space Lower Bounds for Bounded-Error Computation in the Random-Query Model". In: *Proceedings of the 2024 ACM-SIAM Symposium on Discrete Algorithms, SODA 2024*.

[27] François Le Gall, Yupan Liu, Harumichi Nishimura, and Qisheng Wang. "Space-bounded quantum interactive proof systems". In: *arXiv preprint arXiv:2410.23958* (2024).

[28] Brent Harrison, Mitchell Chiew, Jason Necaise, Andrew Projansky, Sergii Strelchuk, and James D Whitfield. "A Sierpinski Triangle Fermion-to-Qubit Transform". In: *arXiv preprint arXiv:2409.04348* (2024).

[29] Takashi Yamakawa and Mark Zhandry. "Verifiable Quantum Advantage without Structure". In: *J. ACM* 71.3 (2024).

[30] Mark Zhandry. "The Space-Time Cost of Purifying Quantum Computations". In: *15th Innovations in Theoretical Computer Science Conference, ITCS 2024*.