

Research Statement

Yi-Kai Liu

November 14, 2009

1 Introduction

My research is in quantum information and theoretical computer science. In particular, I am interested in quantum algorithms and complexity, and in compressed sensing and machine learning.

The first of these topics, quantum algorithms and complexity, is about understanding the capabilities and limitations of quantum computers. Can large-scale quantum computers be built? Small ones have been demonstrated, but the existing technology is not scalable. Can quantum computers solve problems that are hard for classical computers? They can factor numbers efficiently, breaking the RSA public-key cryptosystem; they also have uses in quantum chemistry and combinatorial optimization [24, 1, 17, 13].

These are important questions, for both practical and fundamental reasons. On the practical side, small quantum computers, with only 100 qubits, would already be able to perform certain calculations in quantum chemistry, for which classical methods perform poorly. Larger quantum computers would break most of the public-key cryptosystems that currently protect communications on the Internet. This has motivated an effort to design new cryptosystems that are secure against quantum attack [3]. There are some promising directions, but at present, there is no satisfactory alternative that could replace RSA.

At a more fundamental level, quantum computation seeks to answer some basic questions at the interface between computer science and physics. What is the correct model of computation, in our physical world? If large-scale quantum computers can be built, this would give strong evidence that classical Turing machines are *not* the correct model. On the other hand, quantum computation represents a radical test of the validity of quantum mechanics. If large-scale, fault-tolerant quantum computation is not possible, this would give us new insight into the transition from quantum to classical physics at macroscopic scales.

I have worked on several different aspects of quantum computation: designing quantum algorithms based on wavelet transforms, using quantum algorithms to test graph properties, and proving the hardness of certain problems arising in quantum chemistry [20, 9, 21, 19]. These research directions have led to interesting insights, and I think they are worth pursuing further.

My second topic, compressed sensing, is about methods for finding “hidden structure” in high-dimensional data. This is a common problem in many areas of science and engineering. Essentially, one wants to identify an unknown object that lives in a high-dimensional space, but has some kind of low-dimensional structure (though its precise form is unknown). For instance, one may want to learn an unknown vector in \mathbb{R}^n that has only a few nonzero entries; however, one does not know the locations of the nonzero entries. Can one learn such an object efficiently?

Compressed sensing is a collection of techniques for solving these problems, using special kinds of measurements, together with efficient algorithms for reconstructing the unknown object from the measurement results. These techniques can be applied to a surprising range of problems, most recently, matrix completion, which is closely related to the “Netflix problem” of learning people’s preferences for movies [5, 7]. They also have a beautiful underlying theory, which is based on concentration of measure phenomena in high-dimensional spaces.

My own work has focused on using compressed sensing to do quantum state tomography [16]. This is an important and challenging task in experimental investigations of quantum devices, and our new

methods offer substantial improvements over existing techniques. However, this work has also led to new results on matrix completion [16, 15], and I am interested in developing these ideas further.

Finally, in addition to my main research, I have some interests in cryptography and machine learning. These include lattice-based cryptography, novel tasks such as homomorphic encryption and program obfuscation, and semi-supervised learning (e.g., selective sampling and active learning).

2 Quantum Algorithms and Complexity

2.1 Quantum Algorithms Using Wavelet Transforms: In many quantum algorithms, the crucial “quantum” step consists of a unitary transformation on a high-dimensional vector space \mathbb{C}^{2^n} , that can be performed efficiently (using only $\text{poly}(n)$ elementary quantum gates). A good example is the quantum Fourier transform over \mathbb{Z}_N ($N = 2^n$), which is used in Shor’s algorithm for factoring and discrete logarithms [24]. This approach has since been generalized to non-Abelian groups, to try to solve graph isomorphism and certain lattice problems that have cryptographic significance (see [22, 23, 18, 2], as well as many other works). But this has not led to efficient quantum algorithms thus far, due to a variety of technical obstacles.

My work has concentrated on extending the \mathbb{Z}_N quantum Fourier transform in a different way, using ideas from the theory of wavelets [20]. A particularly useful tool is the curvelet transform over $(\mathbb{Z}_N)^d$, which produces quantum superpositions over the wavefront set of a function [8, 4]. (Roughly speaking, for a function f that is discontinuous along a smooth surface, the wavefront set of f is the set of pairs $(\vec{b}, \vec{\theta})$, where \vec{b} is a point on the surface of discontinuity, and $\vec{\theta}$ is the direction of the discontinuity at \vec{b} .) My work proposed a new approach to designing fast quantum algorithms, using the curvelet transform. In particular, I showed how a quantum curvelet transform can be implemented efficiently, and used this to obtain quantum speed-ups for two oracle problems involving finding the center of a radial function.

The obvious open question is whether the quantum curvelet transform can be used to solve a problem of practical interest. There are a few problems where curvelets might be useful, including lattice problems [23], and hidden polynomial problems over finite fields [10]. Many of these problems are hard for classical computers, so there is hope of finding an *exponential* quantum speed-up. Also, it may prove useful to study other kinds of wavelet transforms, such as diffusion wavelets, which are defined over graphs [11]. This may have some connection to other known quantum algorithms that use quantum walks on graphs.

2.2 Testing Graph Properties: I have some very recent results on quantum algorithms for testing graph properties [9]. In particular, we show quantum algorithms for testing bipartiteness and expansion in bounded-degree graphs, that run in time $O(N^{1/3} \text{poly}(\log N))$; this is polynomially faster than the best possible classical algorithm [14]. These problems had not been studied before in the quantum case, so there is a possibility to discover something new here. In particular, it is an open problem to prove a nontrivial quantum lower bound for either of these problems.

2.3 Quantum Chemistry: I have also worked on hardness results for simulating quantum systems [19, 21]. Such results provide a new perspective on problems in condensed matter physics and quantum chemistry. For instance, we showed that an important problem in quantum chemistry, N -representability, is QMA-complete. (QMA-completeness is a generalization of NP-completeness, that is suitable for problems involving quantum systems.) This result seems surprising at first, since several successful methods for electronic structure calculations are built around heuristics for solving N -representability [12]. Our result implies that these methods must all fail in the worst case. This suggests that real molecular wavefunctions have some hidden structure, that is not present in our worst-case instances. Such structure is known to occur in certain classes of spin chains, but it is not yet understood in the more complicated case of molecules.

3 Compressed Sensing and Quantum State Tomography

Another of my interests concerns compressed sensing, and its application to quantum state tomography and machine learning. The general idea is as follows. Compressed sensing shows that it is possible to learn a sparse vector $\vec{v} \in \mathbb{R}^d$ from a small number of linear measurements. For instance, one can measure a few Fourier coefficients of \vec{v} , then reconstruct \vec{v} by solving a convex program (minimizing the ℓ_1 norm). Matrix completion is a generalization of this idea, where one learns a low-rank matrix from a small number of matrix elements [5, 7]. (This is possible provided that the unknown matrix satisfies an incoherence assumption — its singular vectors must be aligned away from the coordinate axes.)

Matrix completion has many applications, such as learning people’s preferences for movies (the “Netflix problem,” or collaborative filtering), and learning distance metrics and similarity measures. My work adapts these techniques for a new use: quantum state tomography. This is the task of using measurements to determine the state of a quantum system, given many identical copies of that system. This has many experimental applications, such as studying entangled quantum states and characterizing quantum noise processes. However, it is also very resource-intensive: for a system of n qubits, standard techniques require roughly 4^n quantum measurements, and comparable amounts of classical post-processing.

Our work [16] shows that quantum state tomography can be done quadratically faster for states that are almost pure — a case which arises frequently in practice. At a high level, this works because states that are almost pure correspond to density matrices that are close to having low rank, so we can then apply the techniques of matrix completion. In fact, we go further than that: we develop a generalization of matrix completion, using a different set of measurements (Pauli measurements) that are experimentally feasible, and do not require the unknown matrix to satisfy any incoherence assumptions. We also introduce new proof techniques that greatly simplify the existing theory of matrix completion, and allow for further generalizations [15].

Regarding tomography, the next step is to demonstrate our method in a real experiment. Our method is quite different from standard tomography, so while the general theory looks promising, there are many details yet to be worked out. Another direction for future work is to develop the theory of matrix completion along the path taken by vector compressed sensing. For instance, do Pauli measurements satisfy the restricted isometry property? This would be a natural extension of one of the major results in vector compressed sensing [6]. Finally, one wonders how far this general paradigm of compressed sensing can be extended. What other kinds of “sparse” objects can be learned using these techniques? What other problems can be solved using this approach?

References

- [1] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon. Simulated quantum computation of molecular energies. *Science*, 309(5741):1704–1707, 2005.
- [2] D. Bacon, A. M. Childs, and W. van Dam. From optimal measurement to efficient quantum algorithms for the hidden subgroup problem over semidirect product groups. In *FOCS*, pages 469–478, 2005.
- [3] D. J. Bernstein, J. Buchmann, and E. Dahmen (Eds.). *Post-Quantum Cryptography*. Springer, 2009.
- [4] E. J. Candès and D. L. Donoho. Continuous curvelet transform: I. resolution of the wavefront set. *Appl. Comput. Harmon. Anal.*, 19:162–197, 2003.
- [5] E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Found. Comput. Math.* (to appear); arXiv:0805.4471, 2008.

- [6] E. J. Candes and T. Tao. Near-optimal signal recovery from random projections: universal encoding strategies. *IEEE Trans. Inform. Theory*, 52:5406–5425, 2004.
- [7] E. J. Candes and T. Tao. The power of convex relaxation: Near-optimal matrix completion. Submitted; arXiv:0903.1476, 2009.
- [8] E.J. Candès and D. L. Donoho. New tight frames of curvelets and optimal representations of objects with piecewise-c2 singularities. *Comm. Pure Appl. Math.*, 57:219–266, 2002.
- [9] A. M. Childs and Y.-K. Liu. Quantum algorithms for testing bipartiteness and expansion of bounded-degree graphs. In preparation, 2009.
- [10] A. M. Childs, L. J. Schulman, and U. V. Vazirani. Quantum algorithms for hidden nonlinear structures. In *FOCS*, pages 395–404, 2007.
- [11] R. R. Coifman and M. Maggioni. Diffusion wavelets. *Applied and Comput. Harmonic Analysis*, 21(1):53–94, 2006.
- [12] D. A. Mazziotti (Ed.). *Reduced-Density-Matrix Mechanics: With Application to Many-Electron Atoms and Molecules*. Wiley, 2007.
- [13] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser. Quantum computation by adiabatic evolution. arXiv:quant-ph/0001106, 2000.
- [14] O. Goldreich and D. Ron. Property testing in bounded degree graphs. *Algorithmica*, 32(2):302–343, 2002.
- [15] D. Gross. Recovering low-rank matrices from few coefficients in any basis. arXiv:0910.1879, 2009.
- [16] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert. Quantum state tomography via compressed sensing. Submitted; arXiv:0909.3304, 2009.
- [17] I. Kassal, S. P. Jordan, P. J. Love, M. Mohseni, and A. Aspuru-Guzik. Polynomial-time quantum algorithm for the simulation of chemical dynamics. *Proc. Nat. Acad. Sci.*, 105(48):18681–18686, 2008.
- [18] G. Kuperberg. A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. arXiv:quant-ph/0302112, 2003.
- [19] Y.-K. Liu. Consistency of local density matrices is QMA-complete. In *RANDOM 2006*, pages 438–449, 2006.
- [20] Y.-K. Liu. Quantum algorithms using the curvelet transform. In *STOC*, pages 391–400, 2009.
- [21] Y.-K. Liu, M. Christandl, and F. Verstraete. N-representability is QMA-complete. *Phys. Rev. Lett.*, 98:110503, 2007.
- [22] C. Moore, A. Russell, and P. Sniady. On the impossibility of a quantum sieve algorithm for graph isomorphism. In *STOC*, pages 536–545, 2007.
- [23] O. Regev. Quantum computation and lattice problems. *SIAM J. Comput.*, 33(3):738–760, 2004.
- [24] P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, 1997.