Adaptively finding the model parameters for a single-qubit system

Chris Granade[1,2,*], Chris Ferrie[1,3] and D. G. Cory[1,4,5]

[1]Institute for Quantum Computing, University of Waterloo, [2]Department of Physics, University of Waterloo, [3]Department of Applied Mathematics, University of Waterloo, [4]Department of Chemistry, University of Waterloo, [5]Perimeter Institute for Theoretical Physics

Motivation

Given a system of interest, we require accurate knowledge of its Hamiltonian H in order to:

- Make accurate and honest predictions.
- Develop control sequences and error correcting codes for the system.

In order to develop high-fidelity control sequences, we must also learn how the dynamics of our system depend on our control knobs.

Statement of Problem

Let our model of a system depend on some model parameters $\vec{\omega}$, and suppose that we can perform experiments described by design parameters $\vec{e}$. Then, we wish to reduce the risk associated with estimating $\vec{\omega}$ by choosing and performing a sequence of experiments on the system.

Example System

An example of a system of interest is that of a system undergoing Rabi oscillations, given by the Hamiltonian $H = \frac{\omega}{2}\sigma_x$ acting on the initial state $\rho_0 = |0\rangle\langle 0|$, where ω is the parameter to be estimated. Experiments on this system consist of allowing it to evolve for an amount of time t, then measuring the state. The likelihood function for this model is then given by

$$\Pr(0|\omega,t) = 1 - \Pr(1|\omega,t), \quad \Pr(1|\omega,t) = \cos^2(t\omega/2).$$

Experiment Design

This example shows that, when modeling quantum mechanical systems, the likelihood depends strongly on the design of an experiment.

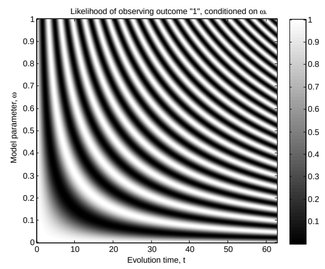

Utility

Because of the strong dependence of L on $\vec{e}$, not all experiments will be equally useful to perform. We capture this notion with utility functions such as the relative entropy or reduction in variance, written as:

$$U_{\text{re}}(\vec{e}) = \mathbb{E}_d\left[D\left(\Pr(\vec{\omega}|d,\vec{e}) \parallel \Pr(\vec{\omega})\right)\right]$$
$$U_{\text{var}}(\vec{e}) = \mathbb{E}_d\left[\text{Var}(\Pr(\vec{\omega})) - \text{Var}(\Pr(\vec{\omega}|d,\vec{e}))\right]$$

Greedy vs Global Optimization

As we gain more knowledge about $\vec{\omega}$, the utility landscape changes. The greedy strategy for selecting experiments maximizes the expected utility at each step.

We then can ask whether there is a better strategy that, by accepting lower utilities in intermediary steps, can achive a smaller risk for the same number of experiments.

Overview of Our Algorithm

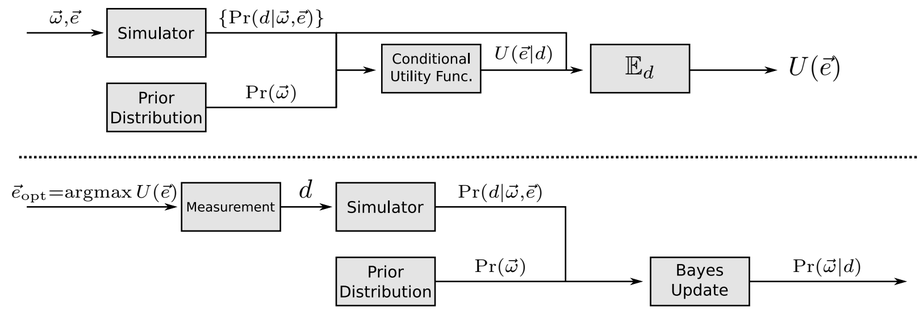

Results
Nyquist Considered Harmful

A common technique is to restrict to the experiments Nyquist sampling times, but we show that this is, on average, suboptimal.

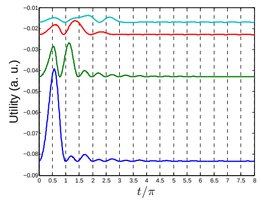

(a) Example of negative variance utility with Nyquist times indicated by vertical lines.

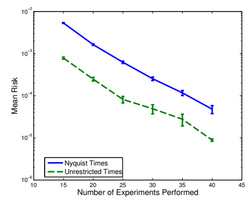

(b) Mean risk of Nyquist-restricted adaptive algorithm compared to unrestricted.

Crucially, the typical argument via Nyquist sampling fails in the strong-measurement case we considered, as we are not sampling a function, but a family of probability distributions. Thus, there is no "signal" to reconstruct, in the sense of the Nyquist sampling theorem.

Results
Greed and Ignorance

Choosing the globally optimal experiment is computationally intensive in general, as it requires averaging over the many branches of the decision tree shown below. On the other hand, the greedy strategy is much less computationally demanding.

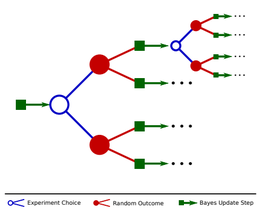

Example decision tree for experiment choice.

In the special case of a uniform prior, however, we found numerically that the greedy strategy performs nearly as well as the globally optimal strategy. Thus, in practice, we do well to consider only the next step when selecting experiments, making our algorithm far more attractive.

Conclusions

Our algorithm provides an exponential speedup over naive parameter estimation methods, such as Fourier estimation. We have also shown that Nyquist sampling is not optimal for a problem of interest, and that the greedy strategy is optimal only for a uniform prior.

Together, these results demonstrate the usefulness of an adaptive Bayesian algorithm for parameter estimation in quantum mechanical systems, especially in comparison with other algorithms in common use. In the presence of noise, this improvement becomes still more stark.

We expect that in more complicated systems, the Bayesian adaptive method will remain useful, especially in applications such as optimal control theory, where having a distribution over Hamiltonians is significantly more useful than a single best estimate.

Future Directions

There are many interesting questions left to answer, including:

- How does our algorithm work in multi-parameter systems?
- What is the appropriate loss function to use for multi-parameter systems?
- What utility function should be employed?
- Is this algorithm useful in the ensemble-measurement scenario?

Additionally, we plan on applying our algorithm to larger systems, as well as systems with more direct applicability to experimental physics.

Acknowledgments

The authors would like to thank NSERC and CERC for their financial support, their peers at the Institute for Quantum Computing for help and advice, and Miriam Diamond for her assistance in implementing the algorithm described here.

For More Information

Please visit `http://goo.gl/y406Y` for more information, or scan the code at right.

*: Author to whom correspondence should be addressed (cgranade@cgranade.com).