Title: Learning in Quantum Control: High-Dimensional Global Optimization for Noisy Quantum Dynamics

Date: Aug 08, 2016 02:30 PM

URL: http://pirsa.org/16080004

Abstract: Quantum control is valuable for various quantum technologies such as high-fidelity gates for universal quantum computing, adaptive quantum-enhanced metrology, and ultra-cold atom manipulation. Although supervised machine learning and reinforcement learning are widely used for optimizing control parameters in classical systems, quantum control for parameter optimization is mainly pursued via gradient-based greedy algorithms. Although the quantum fitness landscape is often compatible for greedy algorithms, sometimes greedy algorithms yield poor results, especially for large-dimensional quantum systems. We employ differential evolution algorithms to circumvent the stagnation problem of non-convex optimization, and we average over the objective function to improve quantum control fidelity for noisy systems. To reduce computational cost, we introduce heuristics for early termination of runs and for adaptive selection of search subspaces. Our implementation is massively parallel and vectorized to reduce run time even further. We demonstrate our methods with two examples, namely quantum phase estimation and quantum gate design, for which we achieve superior fidelity and scalability than obtained using greedy algorithms.

Machine Learning for $\mathscr{Q}$ Metrology & $\mathscr{Q}$ Control

Barry C. Sanders

8 August 2016

Proc ITNG 2010: 506–511; *PRA* **90**, 032310 (2014); Proc ESANN 2016: 327–332
PRL **104** 063603 (2010) **107** 233601 (2011) **110** 220501 (2013) **114** 200502 (2015)

University of Calgary
中国科学技术大学 University of Science and Technology of China
NSERC CRSNG
CIFAR Canadian Institute for Advanced Research
Alberta Innovates Technology Futures
New Journal of Physics The open access journal at the forefront of physics

Table of contents

@bcsanders
in barry-sanders-69b5519
http:// iqst.ca/people/peoplepage.php?id=4
ORCID 0000-0002-8326-8912

Problem

Develop policies to control $\mathscr{Q}$ dynamics, subject to constraints (e.g., run-time T, # particles N, # control parameters K) when greedy algorithms fail.

Applications

- $\mathscr{Q}$ gate design
- Adaptive $\mathscr{Q}$ metrology
  - Coherent control of molecular dynamics
  - $\mathscr{Q}$ measurement trajectories & readout

Control

Dynamically change parameters so system follows closely reference or optimal trajectory.

Policy

Set of instructions that determine the control parameters during the system's evolution.

Claim

Heuristics-based reinforcement-learning algorithms can generate policies for designing desired $\mathscr{Q}$ channels when greedy alternatives fail.

Greedy Algorithms

Optimize using *local* conditions, e.g., local increments perhaps with steepest gradient.

Hill climbing	Quasi-Newton
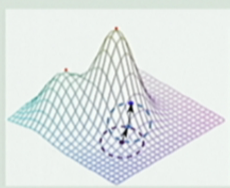	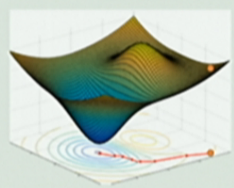
goo.gl/bcfol0	goo.gl/fUjwVT

Reinforcement Learning

Exploit knowledge & *explore* terrain to seek better policy based on reward/punishment.

Non-convex landscape	Particle swarm optimization
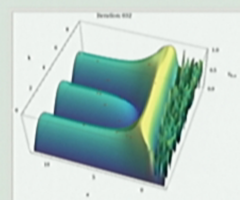	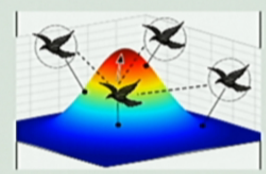

Claim

Heuristics-based reinforcement-learning algorithms can generate feasible policies for designing desired $\mathscr{Q}$ channels when greedy alternatives fail.

Adaptive Phase Estimation

Policies for single-shot adaptive phase estimation scheme with precision exceeding the standard $\mathscr{Q}$ limit (S$\mathscr{Q}$L) up to $N = 100$ particles including noise & loss.

$\mathscr{Q}$-Gate Design Policies

Policies for single-shot high-fidelity three-qubit gates for an architecture of three linearly coupled 4-level transmons.

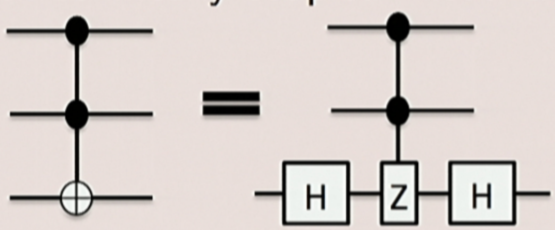

$\mathcal{Q}$-Enhanced Metrology ($\mathcal{Q}$EM) for Interferometric Phase

S$\mathcal{Q}$L

Measure reaction or interference on particles without quantum correlation: $\Delta\varphi \sim \frac{1}{\sqrt{N}}$.

Entanglement/Squeezing

Exploit multi-partite entangled state or squeezed collective uncertainty relations: $\Delta\varphi \sim \frac{1}{N}$.

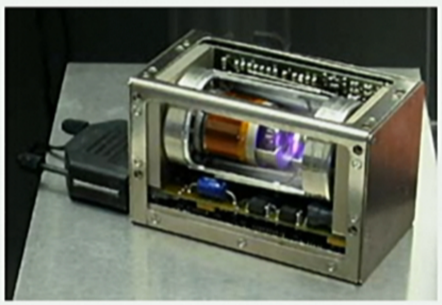

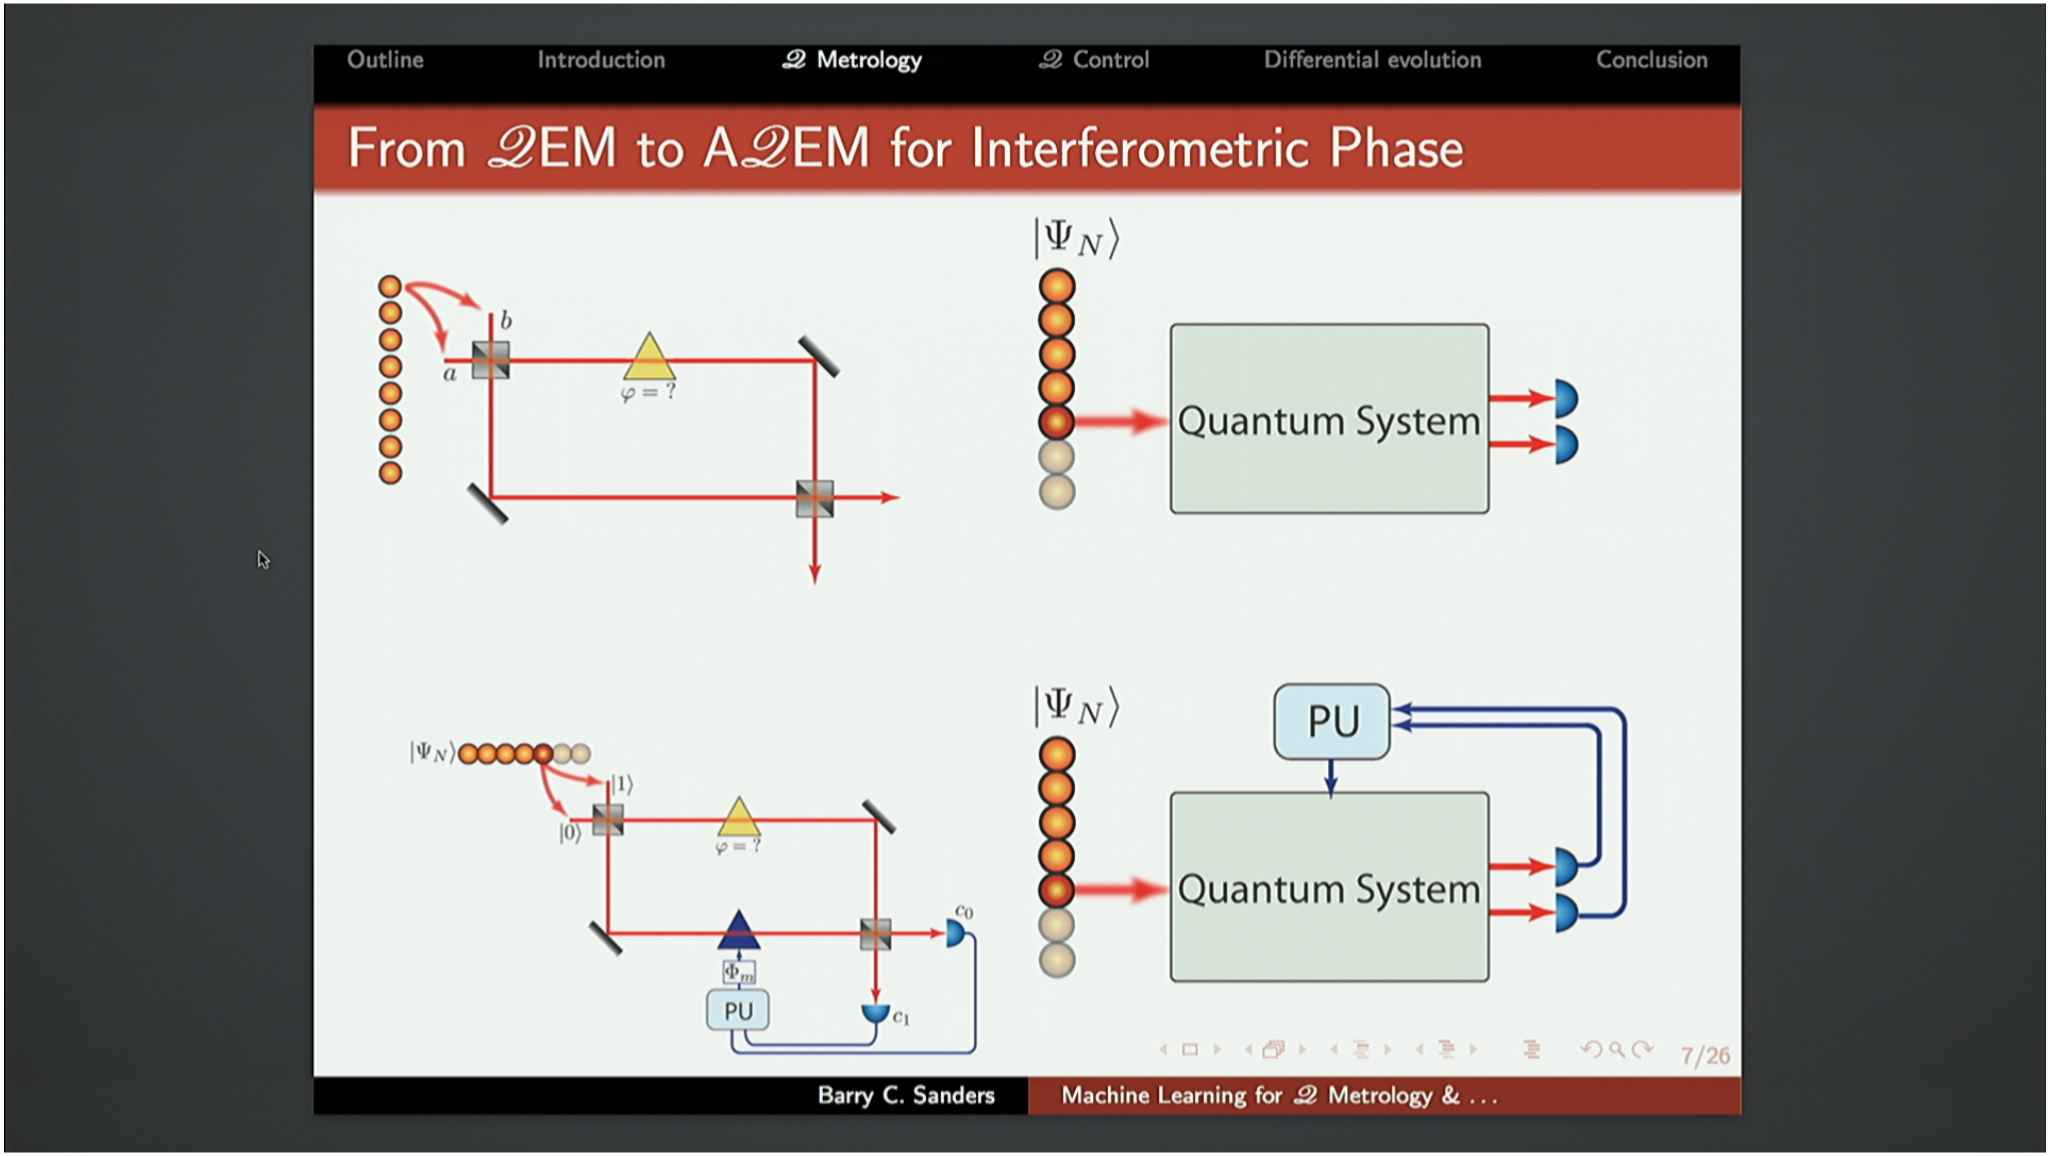
Outline
Introduction
𝒬 Metrology
𝒬 Control
Differential evolution
Conclusion
From 𝒬EM to A𝒬EM for Interferometric Phase
$|\Psi_N\rangle$
b
a
$\varphi = ?$
Quantum System
$|\Psi_N\rangle$
PU
$|1\rangle$
$|0\rangle$
$\varphi = ?$
c_0
Φ_m
PU
c_1
Quantum System
7/26
Barry C. Sanders
Machine Learning for 𝒬 Metrology & . . .

A𝒬EM Policy as Decision Tree

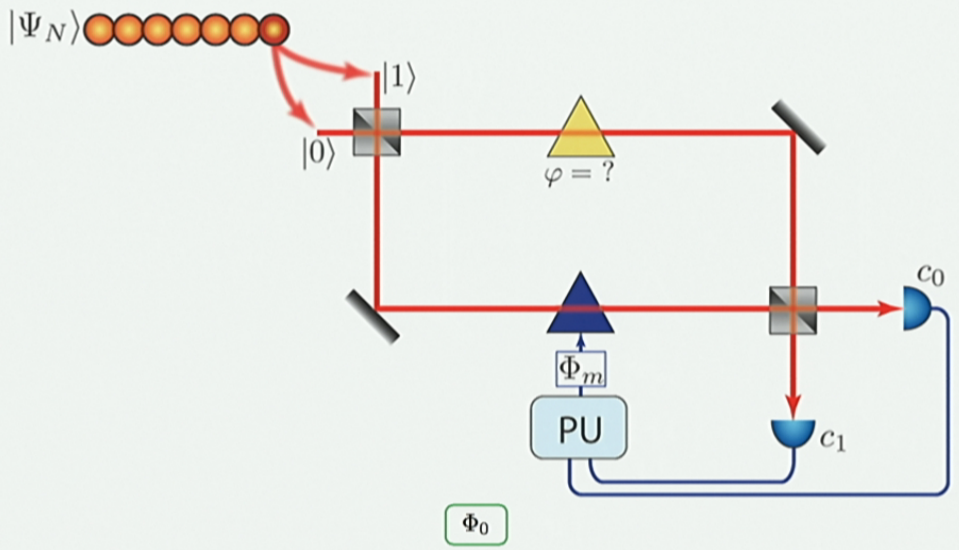

A$\mathscr{Q}$EM Policy as Decision Tree

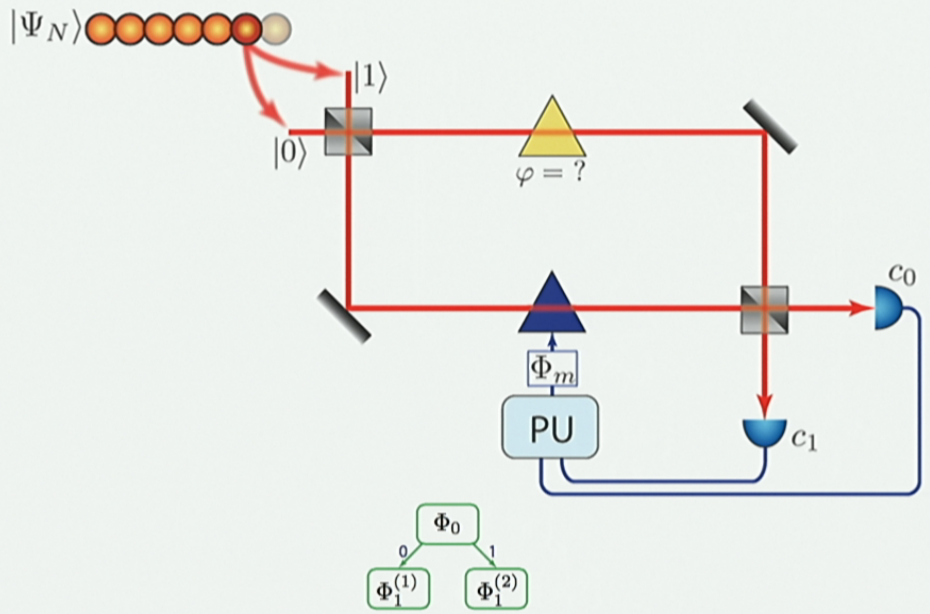

A𝒬EM Policy as Decision Tree

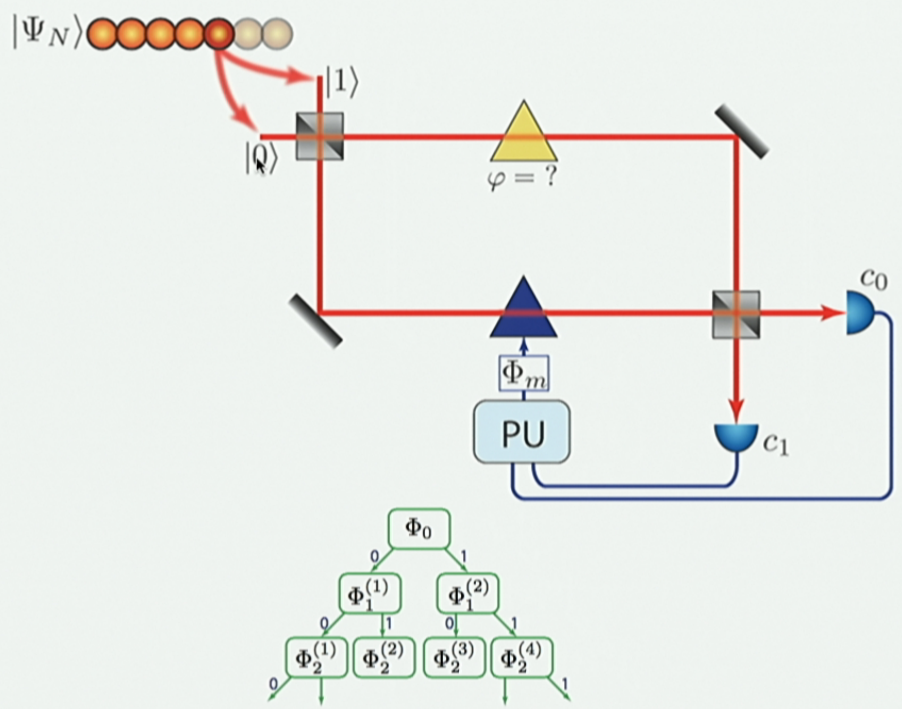

A𝒬EM Policy as Decision Tree

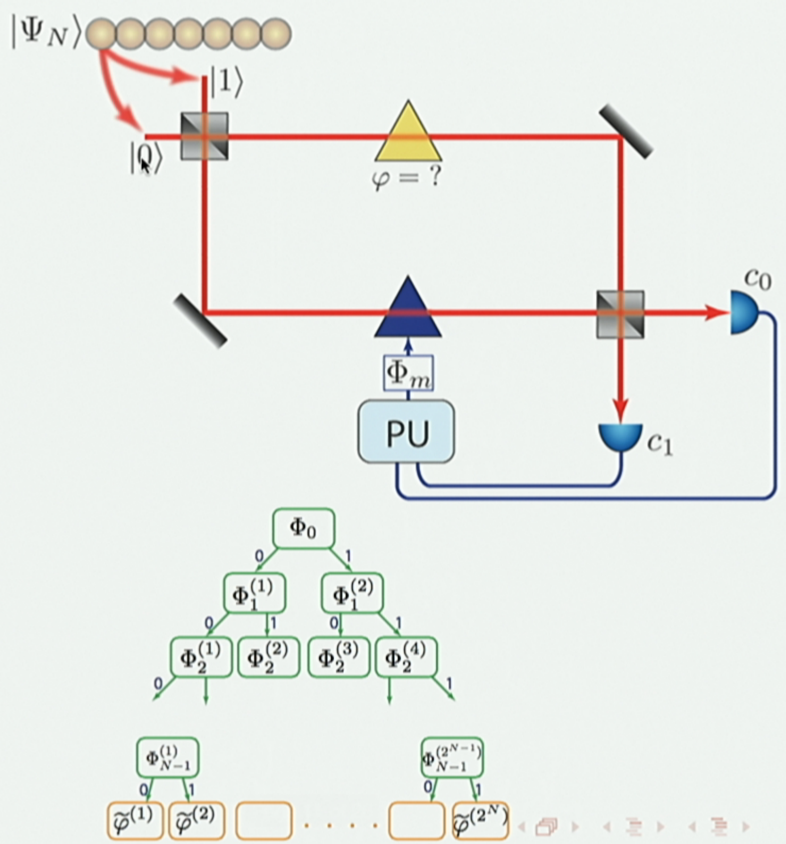

Heuristic: Generalized Logarithmic Search

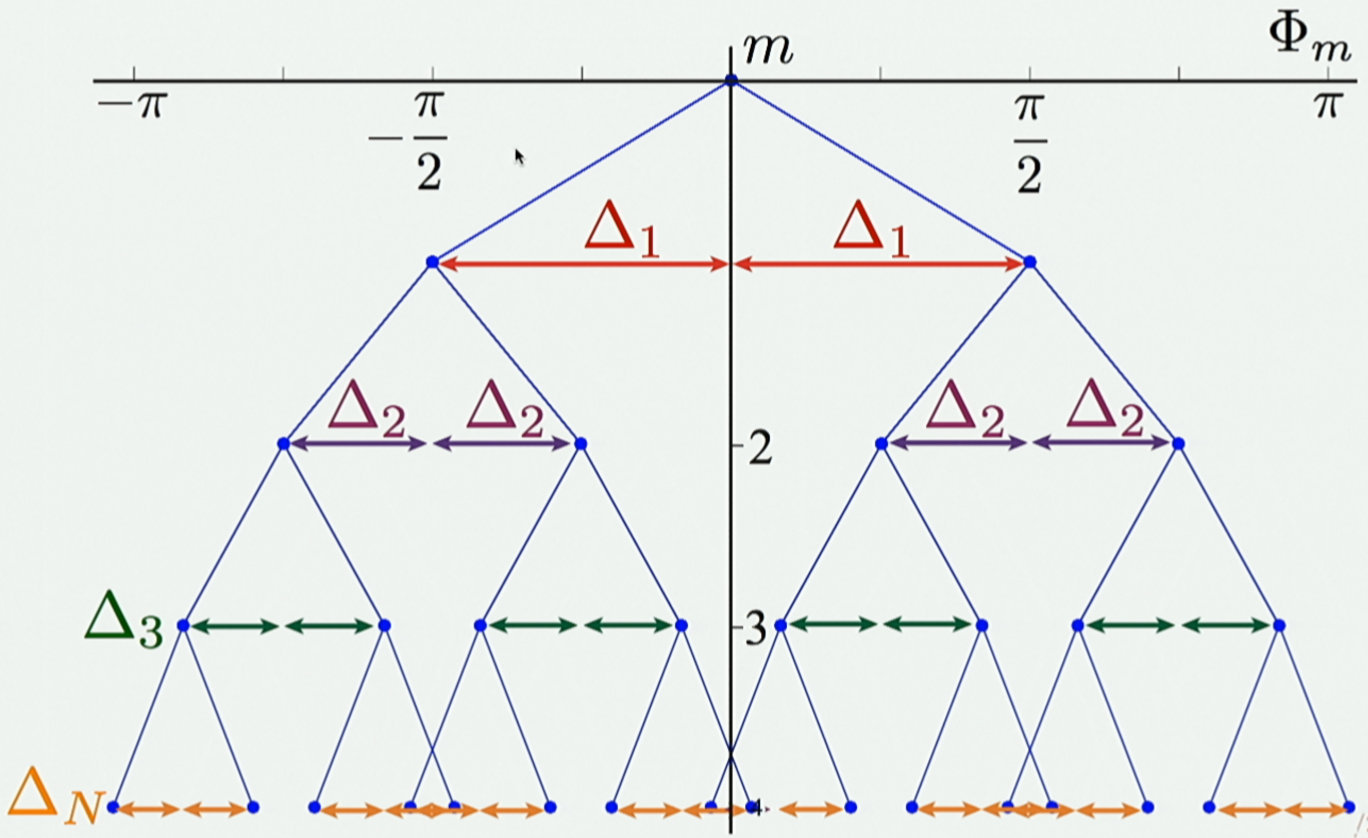

Evaluating Policy ϱ over $K = 10N^2$ Trials with Noise

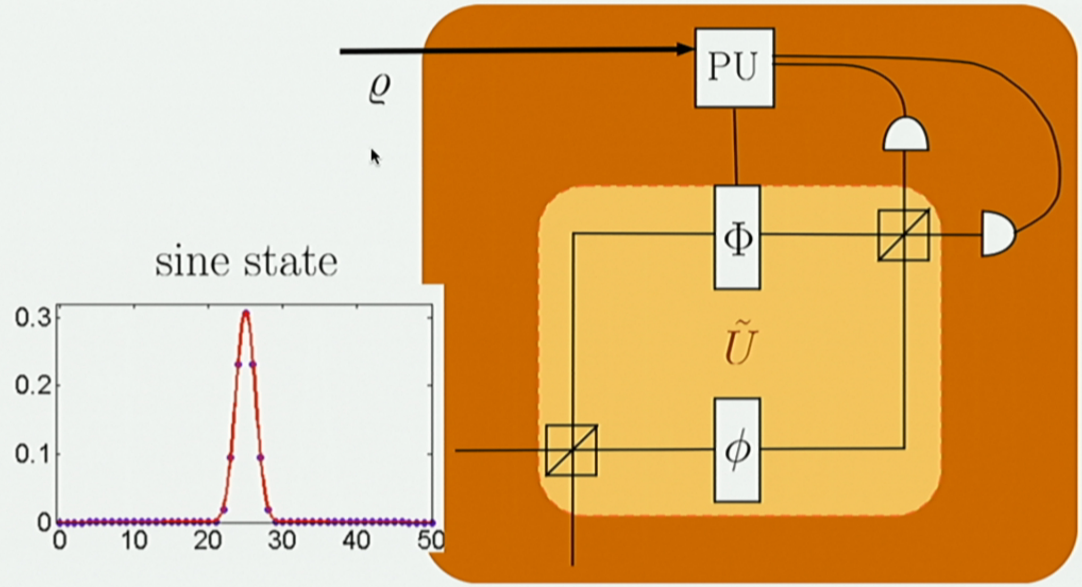

Fitness function

$$S(\varrho) = \left| \sum_{k=1}^{K} \frac{e^{i(\phi - \tilde{\phi})_k}}{K} \right|, \; V_H = S(\varrho)^{-2} - 1$$

Evaluating Policy ϱ over $K = 10N^2$ Trials with Noise

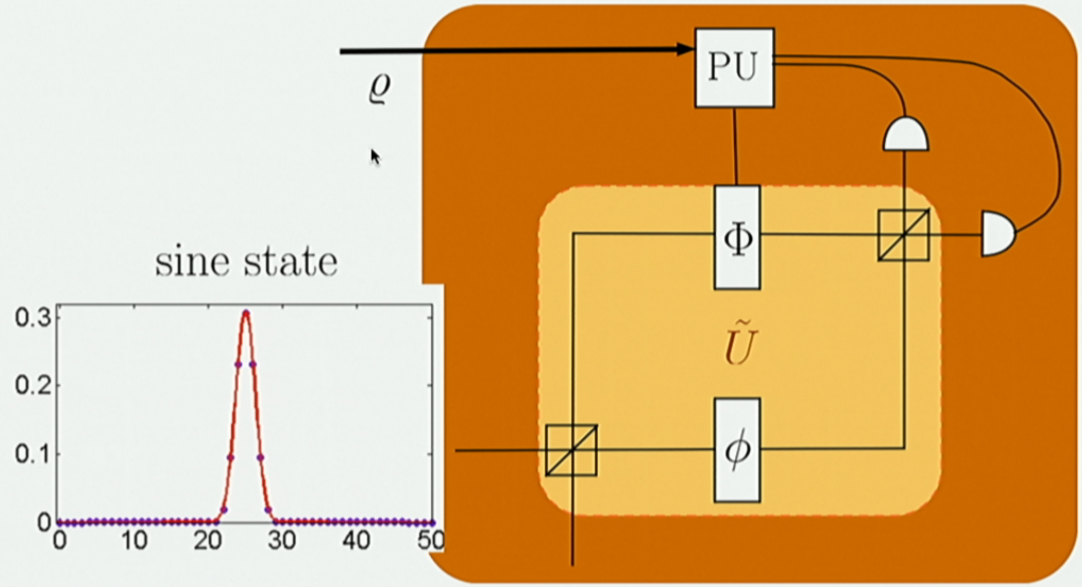

Fitness function

$$S(\varrho) = \left| \sum_{k=1}^{K} \frac{e^{i(\phi-\tilde{\phi})_k}}{K} \right|, \; V_H = S(\varrho)^{-2} - 1$$

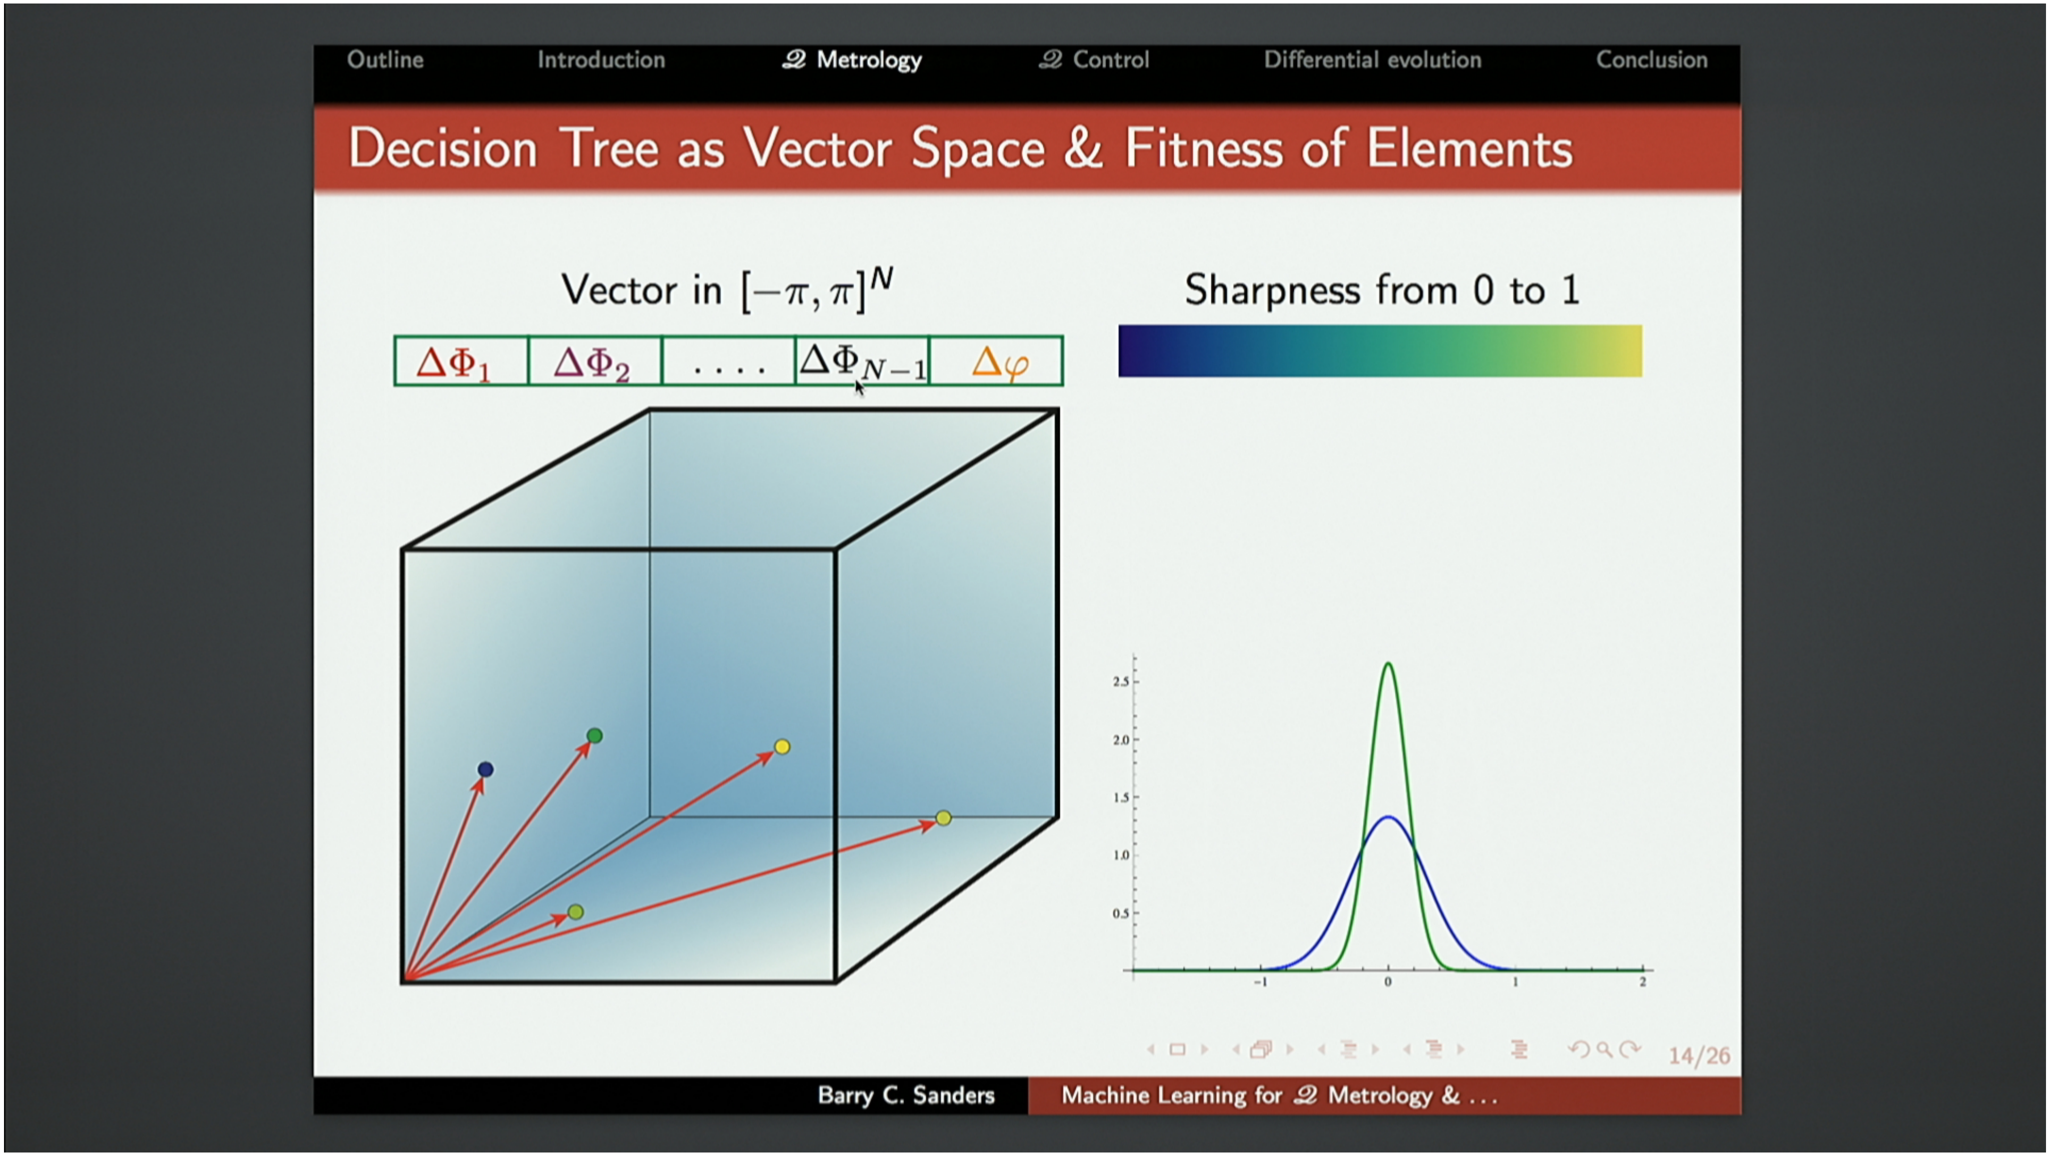
Outline
Introduction
𝒬 Metrology
𝒬 Control
Differential evolution
Conclusion
Decision Tree as Vector Space & Fitness of Elements
Vector in $[-\pi, \pi]^N$
$\Delta\Phi_1$
$\Delta\Phi_2$
....
$\Delta\Phi_{N-1}$
$\Delta\varphi$
Sharpness from 0 to 1
2.5
2.0
1.5
1.0
0.5
-1
0
1
2
14/26
Barry C. Sanders
Machine Learning for 𝒬 Metrology & ...

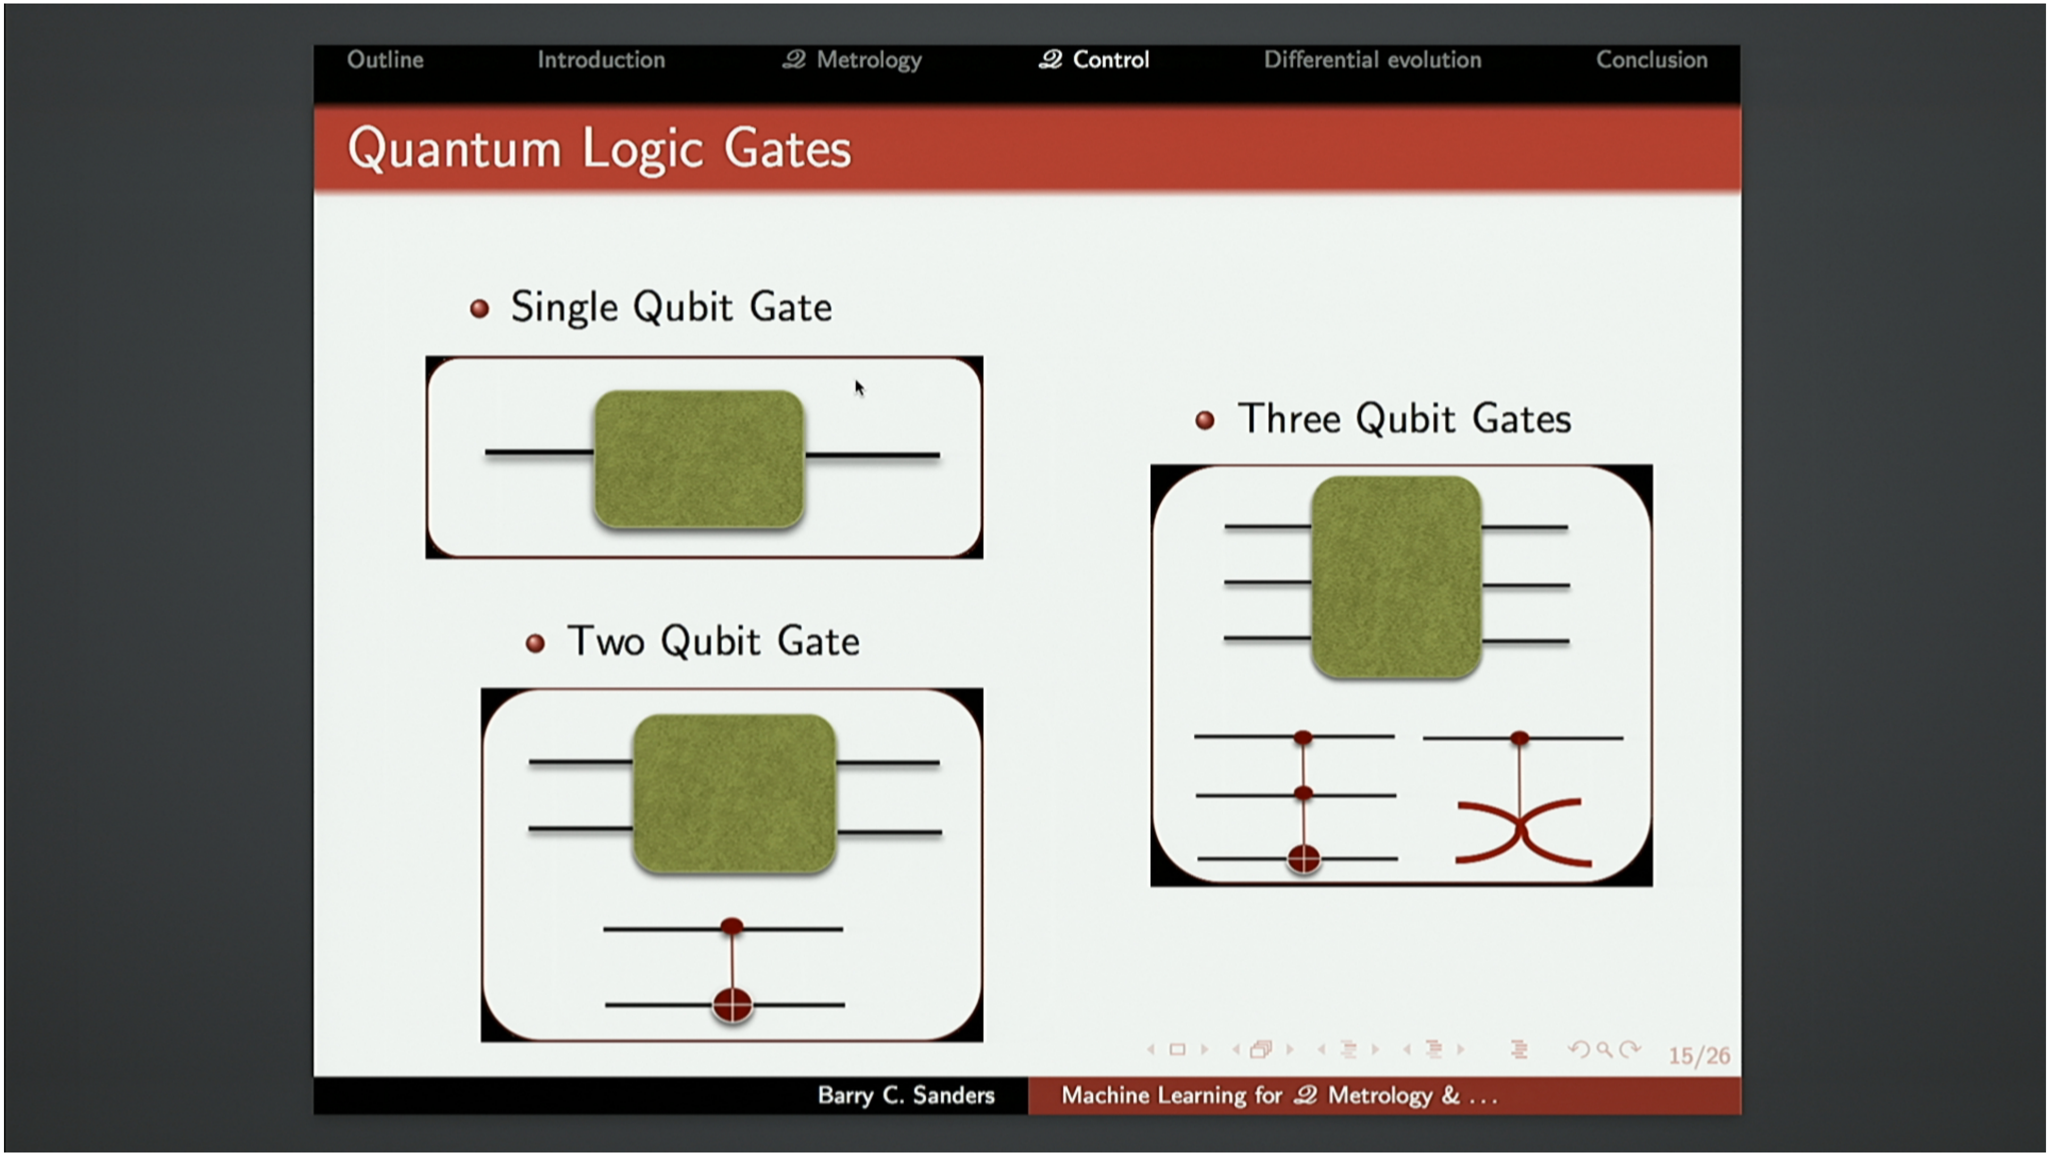
Outline
Introduction
Metrology
Control
Differential evolution
Conclusion
Quantum Logic Gates
Single Qubit Gate
Two Qubit Gate
Three Qubit Gates
15/26
Barry C. Sanders
Machine Learning for Metrology & ...

Controlled-Z gate – Avoided-crossing-level gate

$$\mathrm{X} = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & \sqrt{2} & 0 \\ 0 & \sqrt{2} & 0 & \sqrt{3} \\ 0 & 0 & \sqrt{3} & 0 \end{pmatrix}, \mathrm{Y} = \begin{pmatrix} 0 & -i & 0 & 0 \\ i & 0 & -\sqrt{2}i & 0 \\ 0 & \sqrt{2}i & 0 & -\sqrt{3}i \\ 0 & 0 & \sqrt{3}i & 0 \end{pmatrix}.$$

$\|0\rangle$	$\|0\rangle$	$\|00\rangle$
$\|0\rangle$	$\|1\rangle$	$\|01\rangle$
$\|1\rangle$	$\|0\rangle$	$\|10\rangle$
$\|1\rangle$	$\|1\rangle$	$-\|11\rangle$

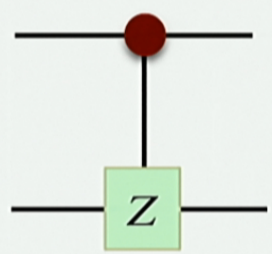

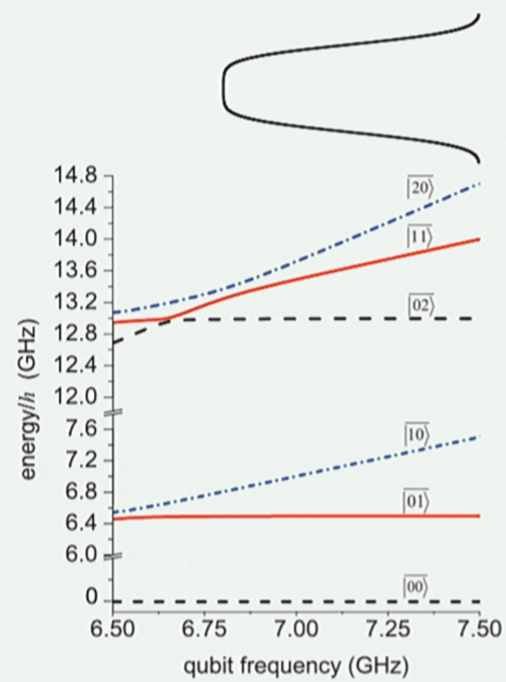

Controlled-Controlled-Z Gate

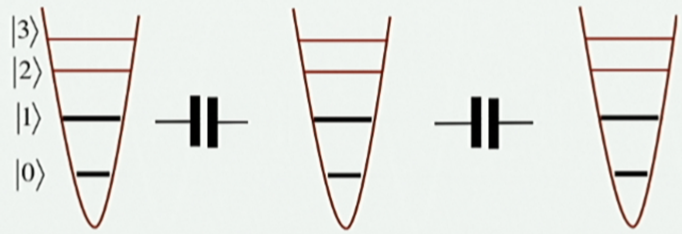

$\|0\rangle$	$\|0\rangle$	$\|0\rangle$	$\|000\rangle$
$\|0\rangle$	$\|0\rangle$	$\|1\rangle$	$\|001\rangle$
$\|0\rangle$	$\|1\rangle$	$\|0\rangle$	$\|010\rangle$
$\|0\rangle$	$\|1\rangle$	$\|1\rangle$	$\|011\rangle$
$\|1\rangle$	$\|0\rangle$	$\|0\rangle$	$\|100\rangle$
$\|1\rangle$	$\|0\rangle$	$\|1\rangle$	$\|101\rangle$
$\|1\rangle$	$\|1\rangle$	$\|0\rangle$	$\|110\rangle$
$\|1\rangle$	$\|1\rangle$	$\|1\rangle$	$-\|111\rangle$

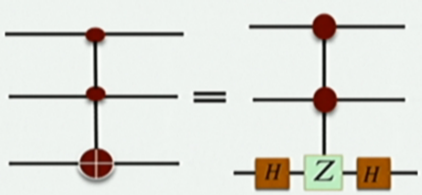

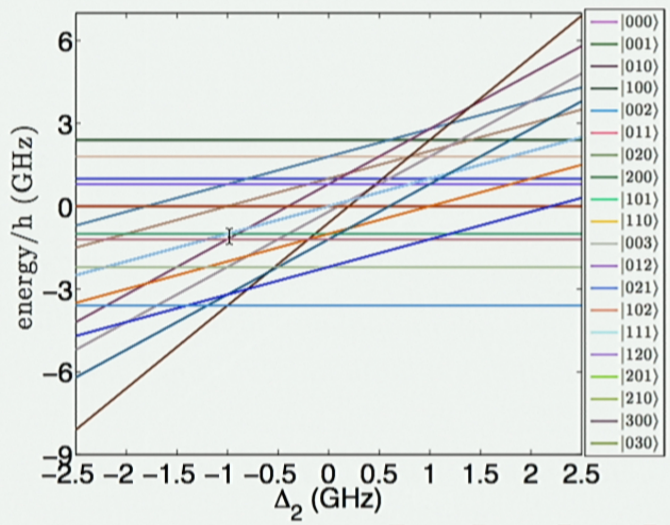

$\mathcal{Q}$ Control & Gate Design

Hamiltonian with drift & control

$$\hat{H}[\mathbf{\Delta}(t)] = \hat{H}^{\text{dr}} + \mathbf{\Delta}(t) \cdot \hat{\mathbf{H}}^{\text{c}} = \hat{H}^{\text{dr}} + \sum_{\ell=1}^{L} \Delta_\ell(t) \hat{H}_\ell^{\text{c}}$$

Resultant unitary evolution

$$\tilde{U}[\mathbf{\Delta}(\Theta); \Theta] = \mathcal{T} \exp\left\{-\mathrm{i} \int_0^{\Theta} \hat{H}(\mathbf{\Delta}(\tau))\mathrm{d}\tau\right\}$$

$\mathcal{M}$ time intervals

Equal time intervals $\Theta/\mathcal{M} \implies (K = L\mathcal{M})$ dimensions

$\mathcal{Q}$ control optimization objective function

Operation/gate: $\mathcal{F}(\Theta) = \text{Tr}\left[\tilde{U}[\mathbf{\Delta}(\mathbf{\Theta}); \Theta] CCZ^\dagger\right]/8$

Search-Space Dimension D, # Constraints Υ, # Agents Ξ

Adaptive Phase Estimation

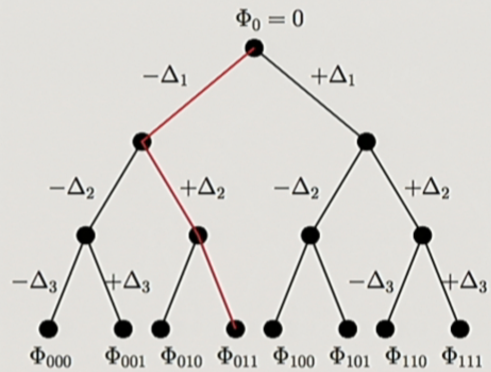

- $4 \le N = D \le 100$
- $\Upsilon = 2N + 1$
- $\Xi = 50$

Gate Design

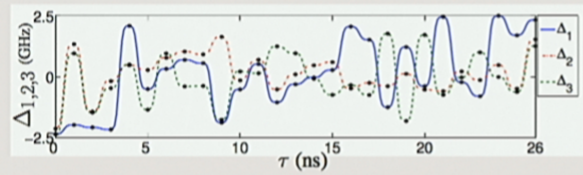

- $D = K \le 84$
- $\Upsilon = 2K$
- $\Xi = 4K + 1$

Heuristic-Based Machine-Learning Scheme

Differential Evolution

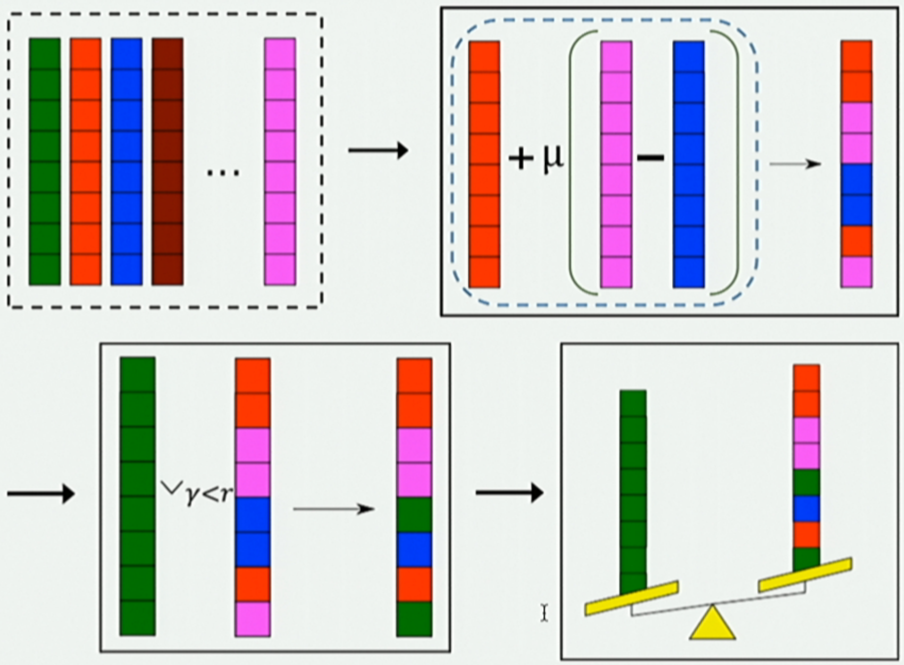

Iteration (i) up to Generation G

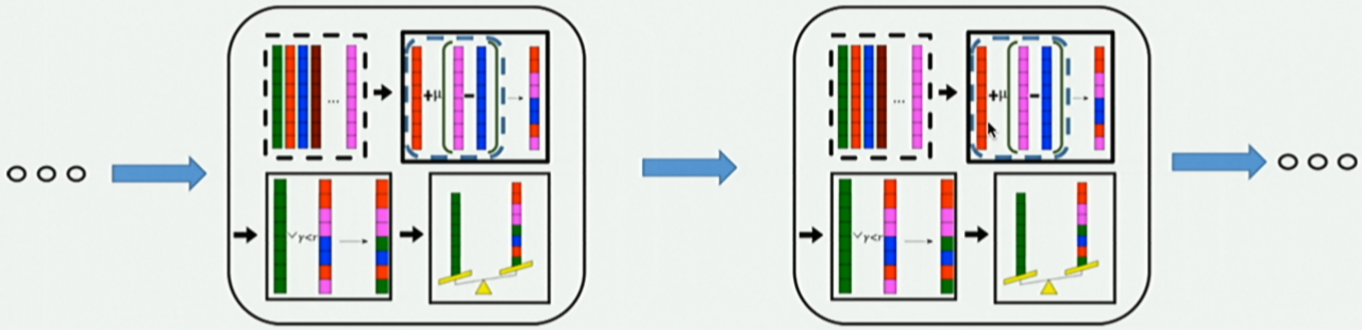

Heuristic-Based Machine-Learning Scheme

Differential Evolution

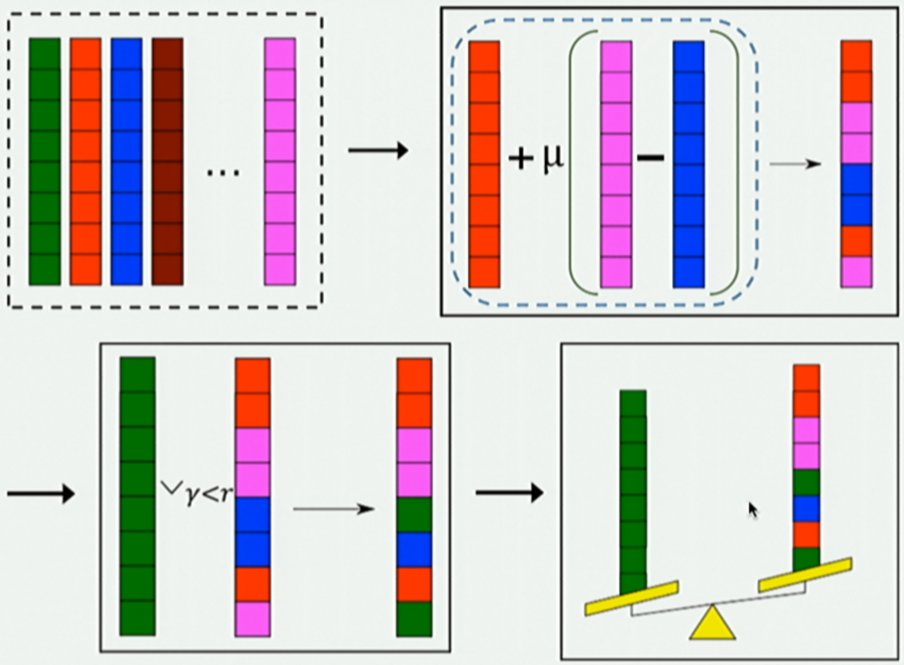

SuSSADE

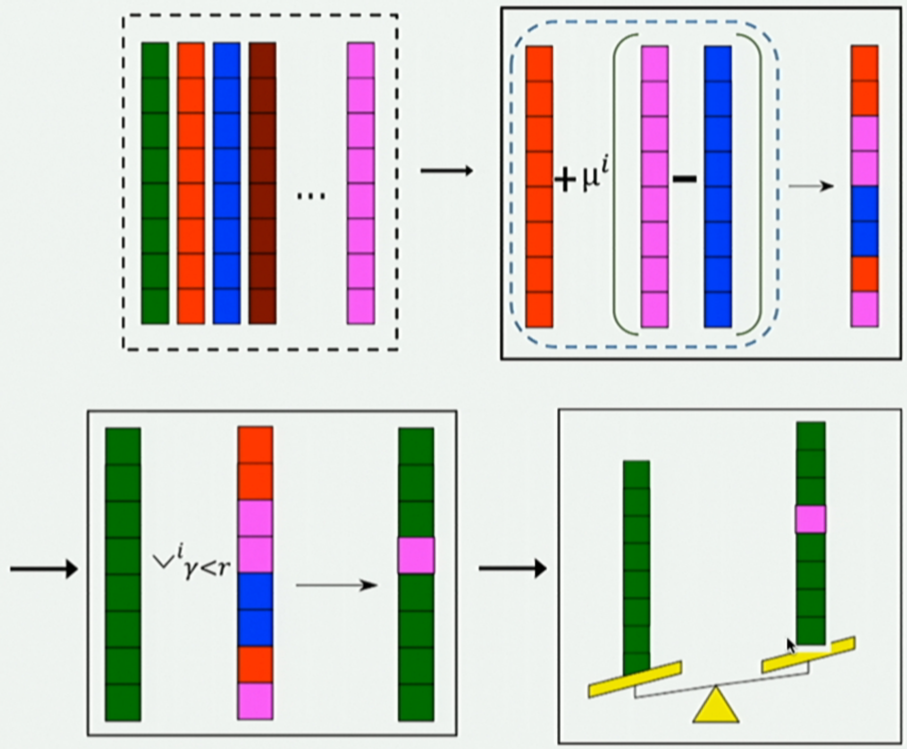

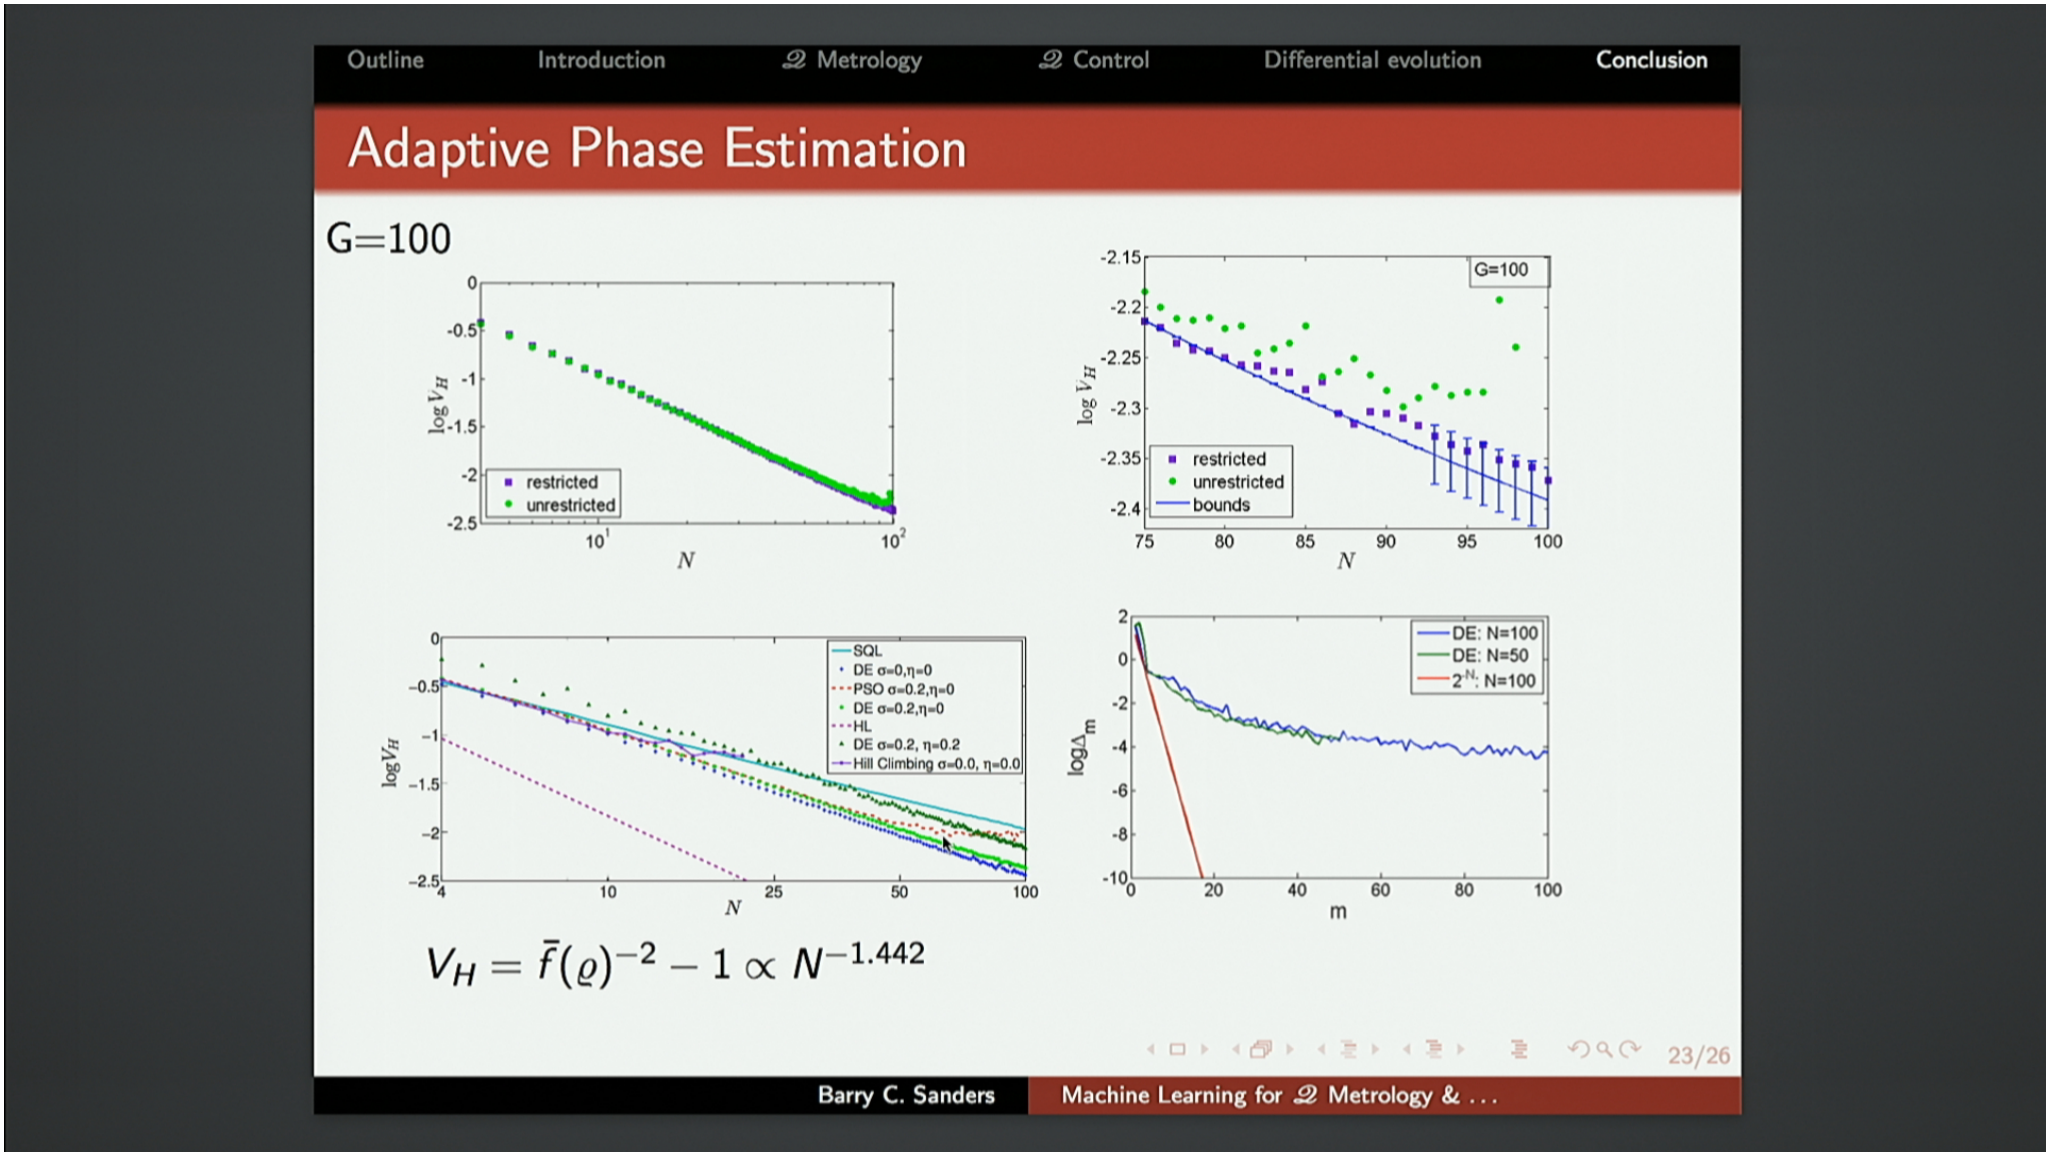
Outline
Introduction
Metrology
Control
Differential evolution
Conclusion
Adaptive Phase Estimation
G=100
log V_H
restricted
unrestricted
N
G=100
restricted
unrestricted
bounds
N
SQL
DE σ=0,η=0
PSO σ=0.2,η=0
DE σ=0.2,η=0
HL
DE σ=0.2, η=0.2
Hill Climbing σ=0.0, η=0.0
logV_H
N
DE: N=100
DE: N=50
2^-N: N=100
logΔ_m
m
V_H = f̄(ϱ)^−2 − 1 ∝ N^−1.442
23/26
Barry C. Sanders
Machine Learning for Metrology & ...

Adaptive Phase Estimation

G=100

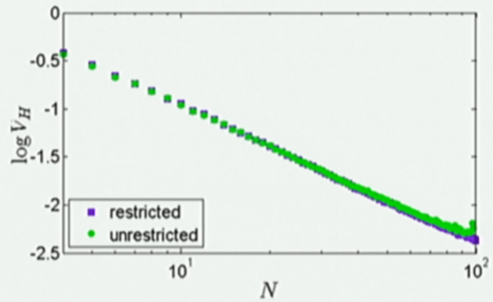

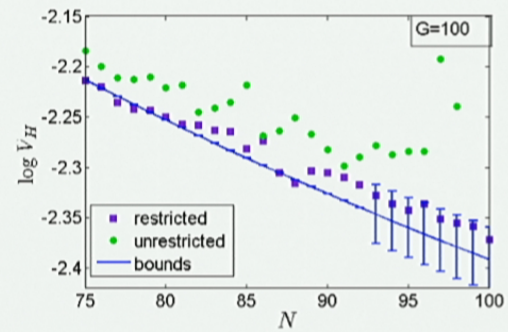

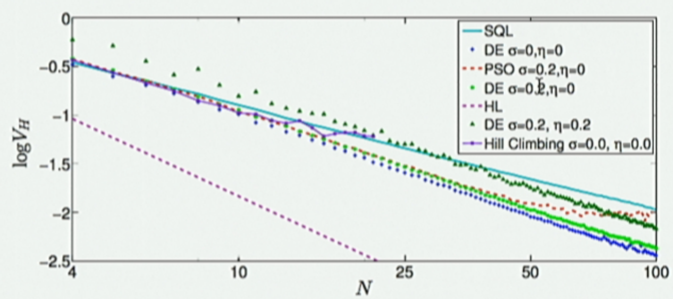

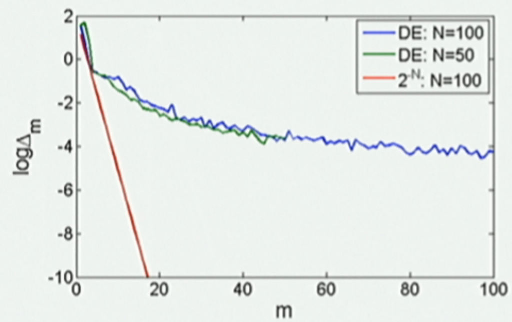

$$V_H = \bar{f}(\varrho)^{-2} - 1 \propto N^{-1.442}$$

Assessing Performance of CCZ Design Algorithms

G=200000

Method	$\mathcal{F}_{\text{best}}$
Quasi-Newton	0.9912
Simplex	0.9221
DE	0.9931
SuSSADE	0.9999

- $g = 30$MHz, $\Theta = 26$ns
- # runs for Quasi-Newton is 40
- # runs for DE & SuSSADE is 20

CCZ gate

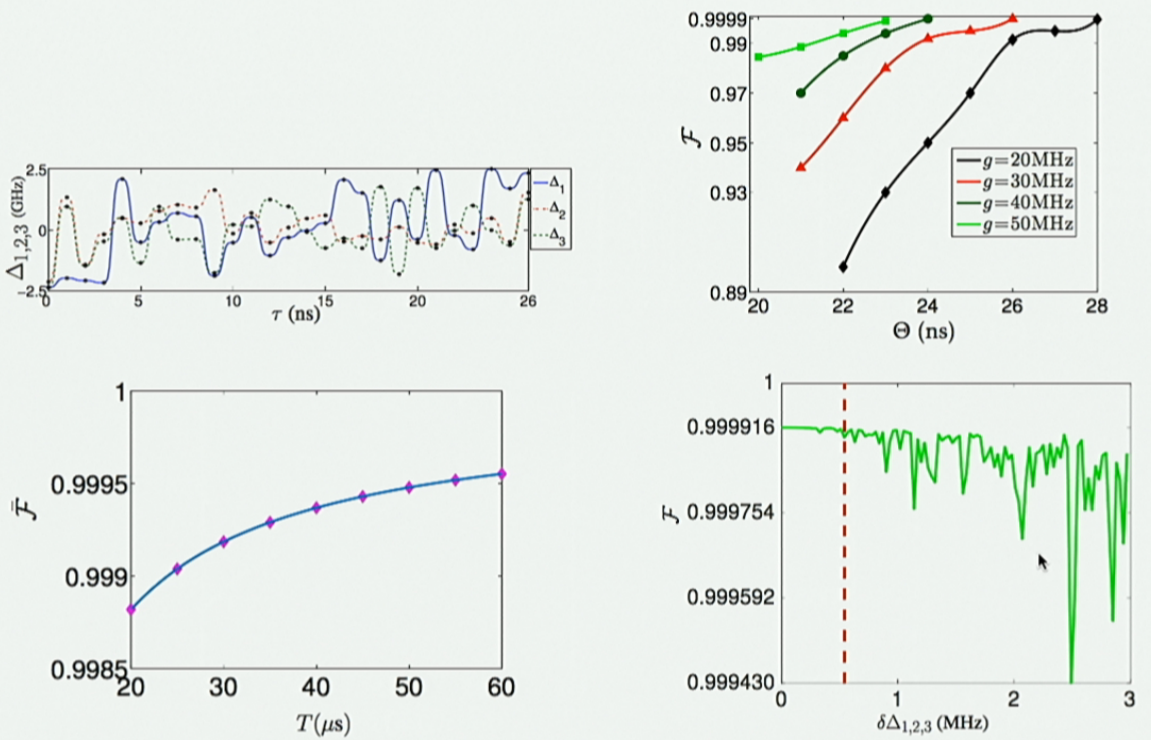
$\Delta_{1,2,3}$ (GHz)
τ (ns)
Δ_1
Δ_2
Δ_3
$\mathcal{F}$
Θ (ns)
g=20MHz
g=30MHz
g=40MHz
g=50MHz
$\bar{\mathcal{F}}$
$T(\mu s)$
$\mathcal{F}$
$\delta\Delta_{1,2,3}$ (MHz)

Summary

- Heuristics-based learning algorithms can generate feasible policies for achieving desired $\mathscr{Q}$ channels when standard algorithms fail
- We generate policies to design single-shot high-fidelity three-qubit gates comprising linearly coupled transmons
- We generate policies to design single-shot adaptive phase estimation surpassing the standard $\mathscr{Q}$ limit
- Our approach can be applied to other $\mathscr{Q}$ control problems for any desired target channel